

Veritas Storage Foundation™ and High Availability Solutions 6.0.5 Installation Guide - Linux

6.0.5 Maintenance Release

Veritas Storage Foundation and High Availability Installation Guide

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

Product version: 6.0.5

Document version: 6.0.5 Rev 0

Legal Notice

Copyright © 2014 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.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

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

Support agreement resources

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

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

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>

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>

Contents

Technical Support	4	
Chapter 1	About Veritas Storage Foundation and High Availability Solutions	10
	About the installation script	10
	The installmr script options	11
	List of products	15
	Supported database software	15
Chapter 2	Installing the products for the first time	16
	Installing the Veritas software using the Install Bundles feature	16
	Installing the Veritas software using the script-based installer	17
	Installing the Veritas software using the Web-based installer	18
	Starting the Veritas Web-based installer	19
	Obtaining a security exception on Mozilla Firefox	19
	Installing 6.0.5 with the Veritas Web-based installer	20
Chapter 3	Upgrading to 6.0.5 from releases earlier than 6.0.1	21
	Planning to upgrade to SFHA Solutions 6.0.5	21
	Supported upgrade paths for SFHA Solutions 6.0.5 from releases earlier than 6.0.1	21
	Pre-upgrade planning for Volume Replicator	24
	Prepare upgrading to 6.0.5 with Install Bundles	27
	Performing a full upgrade with Install Bundles	28
	Performing a full upgrade on a VCS cluster using Install Bundles	28
	Performing a full upgrade of SFHA using Install Bundles	29
	Performing a full upgrade of SFCFSHA using Install Bundles	32
	Performing a full upgrade of SF Oracle RAC using Install Bundles	36
	Performing a phased upgrade using Install Bundles	42
	Performing a phased VCS upgrade using Install Bundles	43
	Performing a phased SFHA upgrade using Install Bundles	45

	Performing a phased SFCFSHA upgrade using Install Bundles	48
	Performing a phased upgrade of SF Oracle RAC using Install Bundles	49
	Performing an automated upgrade using response files with Install Bundles	60
	Performing an automated upgrade of VCS, SFHA, or SFCFSHA using response files with Install Bundles	60
	Upgrading SF Oracle RAC using a response file	61
	Performing a rolling upgrade using Install Bundles	62
	Supported rolling upgrade paths	62
	Performing a rolling upgrade of VCS, SFHA, and SFCFSHA with Install Bundles	62
	Performing a rolling upgrade of SF Oracle RAC with Install Bundles	64
Chapter 4	Upgrading to 6.0.5 from 6.0.1 or later	72
	Supported upgrade paths from 6.0.1 or later	72
	Prerequisites for upgrading to 6.0.5	73
	Downloading required software to upgrade to 6.0.5	73
	Performing a full upgrade to 6.0.5 on a cluster	73
	Performing a full upgrade to 6.0.5 on a Veritas Cluster Server	73
	Performing a full upgrade to 6.0.5 on an SFHA cluster	74
	Performing a full upgrade to 6.0.5 on an SFCFSHA cluster	75
	Performing a full upgrade to 6.0.5 on an SF Oracle RAC cluster	77
	Upgrading to 6.0.5 on a standalone system	81
	Performing a phased upgrade to SFCFSHA and SFRAC	83
	Performing a phased upgrade of SFCFSHA	83
	Performing phased upgrade of SF Oracle RAC to version 6.0.5	94
	Upgrading to 6.0.5 on a system that has encapsulated boot disk	106
	Performing a rolling upgrade using the installmr script	106
	About rolling upgrades	107
	Prerequisites for a rolling upgrade	107
	Performing a rolling upgrade using the <code>installmr</code> script	108
	Upgrading the operating system	117
	Upgrading SFHA with the Veritas Web-based installer	117
	Verifying software versions	118

Chapter 5	Removing Veritas Storage Foundation and High Availability Solutions	119
	About removing Veritas Storage Foundation and High Availability Solutions 6.0.5	119
	Uninstalling Veritas Storage Foundation Cluster File System High Availability 6.0.5	119
	Preparing to uninstall Veritas Storage Foundation Cluster File System High Availability	120
	Uninstalling Veritas Storage Foundation Cluster File System High Availability	121
	Uninstalling Veritas Storage Foundation for Oracle RAC	122
Appendix A	Using Oracle RAC High Availability IP (HAIP) with Oracle RAC 11.2.0.2 and later versions	125
	Using Oracle High Availability IP (HAIP)	125
	Configuring static IP addresses for HAIP using the SFHA script-based installer	126
	Configuring static IP addresses for HAIP manually	129

About Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [About the installation script](#)
- [List of products](#)
- [Supported database software](#)

About the installation script

Veritas Storage Foundation and High Availability Solutions 6.0.5 provides an installation and upgrade script. To install or upgrade the patches that are included in this release, you can use the `installmr` script. The `installmr` script lets you install or upgrade all the patches that are associated with the packages installed.

For more information regarding installation, See [“Installing the Veritas software using the script-based installer”](#) on page 17.

Symantec has introduced a new Install Bundles feature to help you install or upgrade directly to maintenance level with one execution. You can use the `-base_path` option to install or upgrade base and maintenance bundles. There are a few prerequisites for using Install Bundles feature for installation and upgrade of 6.0.5 mentioned below:

Scenario	Install and upgrade process
No product is installed on the target system	Run <code>installmr</code> with <code>-base_path</code> option to install 6.0.5
The product version before 6.0.1 is installed on the target system	Run <code>installmr</code> with <code>-base_path</code> option to upgrade to 6.0.5
The product version 6.0.1 or later is installed on the target system	Run <code>installmr</code> to upgrade to 6.0.5

For example, invoke the `installmr` script with `-base_path` option to install or upgrade 6.0.5.

For installation:

```
./installmr -base_path /tmp/sfha6.0.1/
```

For upgrade:

```
./installmr -base_path /tmp/sfha6.0.1/ -upgrade
```

For more information regarding installing 6.0.5 using Install Bundles feature, See [“Installing the Veritas software using the Install Bundles feature”](#) on page 16.

Symantec strongly recommends you to use the Install Bundles feature to install or upgrade Veritas Storage Foundation and High Availability solutions 6.0.5.

The installmr script options

The following table lists the command line options for the `installmr` and `upgrade` script:

Table 1-1 The command line options for the product `installmr` and `upgrade` script

Command Line Option	Function
<code>system1 system2...</code>	Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.
<code>-base_path</code>	The <code>-base_path</code> option is used to define the path of a base level release to be integrated with a maintenance level release in order for the two releases to be simultaneously installed.

Table 1-1 The command line options for the product installmr and upgrade script (*continued*)

Command Line Option	Function
<code>-precheck</code>	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
<code>-postcheck</code>	Checks any issues after installation or upgrading on the system.
<code>-responsefile <i>response_file</i></code>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
<code>-logpath <i>log_path</i></code>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.
<code>-tmppath <i>tmp_path</i></code>	Specifies a directory other than <code>/var/tmp</code> as the working directory for the installation scripts. This destination is where initial logging is performed and where RPMs are copied on remote systems before installation.
<code>-timeout <i>timeout_value</i></code>	The <code>-timeout</code> option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the <code>-timeout</code> option overrides the default value of 1200 seconds. Setting the <code>-timeout</code> option to 0 will prevent the script from timing out. The <code>-timeout</code> option does not work with the <code>-serial</code> option.
<code>-keyfile <i>ssh_key_file</i></code>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i <i>ssh_key_file</i></code> to every SSH invocation.
<code>-hostfile <i>full_path_to_file</i></code>	Specifies the location of a file that contains a list of hostnames on which to install.

Table 1-1 The command line options for the product installmr and upgrade script (*continued*)

Command Line Option	Function
<code>-patchpath</code>	The <code>-patchpath</code> option is used to define the complete path of a directory available to all install systems (usually NFS mounted) that contains all patches to be installed by installmr.
<code>-kickstart dir_path</code>	Produces a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of Symantec rpms in the correct order for installing, in a format that can be used for Kickstart installations. The <code>dir_path</code> indicates the path to the directory in which to create the file.
<code>-yumgroupxml</code>	The <code>-yumgroupxml</code> option is used to generate a yum group definition XML file. The <code>createrepo</code> command can use the file on Redhat Linux to create a yum group for automated installation of all rpms for a product. An available location to store the XML file should be specified as a complete path. The <code>-yumgroupxml</code> option is supported on Redhat Linux only.
<code>-require</code>	The <code>-require</code> option is used to specify a installer hot fix file.
<code>-serial</code>	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
<code>-rsh</code>	Specifies this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.
<code>-redirect</code>	Displays progress details without showing the progress bar.
<code>-pkgset</code>	Discovers and displays the RPM group (minimum, recommended, all) and RPMs that are installed on the specified systems.
<code>-pkgtable</code>	Displays product's RPMs in correct installation order by group.

Table 1-1 The command line options for the product installmr and upgrade script (*continued*)

Command Line Option	Function
-pkginfo	Displays a list of RPMs and the order of installation in a human-readable format.
-listpatches	The <code>-listpatches</code> option displays product patches in correct installation order.
-makeresponsefile	Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option.
-comcleanup	The <code>-comcleanup</code> option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
-version	Checks and reports the installed products and their versions. Identifies the installed and missed RPMs and patches where applicable for the product. Provides a summary that includes the count of the installed and any missed RPMs and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available.
-nolic	Allows installation of product RPMs without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
-rolling_upgrade	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
-rollingupgrade_phase1	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel RPMs get upgraded to the latest version

Table 1-1 The command line options for the product installmr and upgrade script (*continued*)

Command Line Option	Function
<code>-rollingupgrade_phase2</code>	The <code>-rollingupgrade_phase2</code> option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version."

List of products

Apply the patches for the following Veritas Storage Foundation and High Availability products:

- Veritas Dynamic Multi-Pathing (DMP)
- Veritas Volume Manager (VxVM)
- Veritas File System (VxFS)
- Veritas Storage Foundation (SF)
- Veritas Cluster Server (VCS)
- Veritas Storage Foundation and High Availability (SFHA)
- Veritas Storage Foundation Cluster File System and High Availability (SFCFSHA)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)
- Veritas Storage Foundation for Sybase ASE CE (SFSYBASECE)
- Symantec VirtualStore (SVS)

Supported database software

For the latest information on supported Oracle database versions, see the following TechNote:

<http://www.symantec.com/docs/DOC5081>

Support for minor database versions is also documented in the afore-mentioned TechNote.

Additionally, see the following Oracle support site for information on patches that may be required by Oracle for each release.

<https://support.oracle.com>

Installing the products for the first time

This chapter includes the following topics:

- [Installing the Veritas software using the Install Bundles feature](#)
- [Installing the Veritas software using the script-based installer](#)
- [Installing the Veritas software using the Web-based installer](#)

Installing the Veritas software using the Install Bundles feature

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product using the Install Bundles feature in one step.

To install the Veritas software 6.0.5 using Install Bundles:

- 1 Download Storage Foundation and High Availability Solutions 6.0.1 from <http://fileConnect.symantec.com>.

Note: For 6.0.5, the base release version can be 6.0.1, 6.0.2, or 6.0.4.

- 2 Extract the tar ball into the `/tmp/sfha6.0.1/` directory.
- 3 Download SFHA Solutions 6.0.5 from <http://sort.symantec.com/patches>.
- 4 Extract it to the `/tmp/sfha6.0.5` directory.
- 5 Change to the `/tmp/sfha6.0.5` directory by entering:

```
# cd /tmp/sfha6.0.5
```


- 6 Invoke the `installmr` script with `-base_path` option to install 6.0.1 and 6.0.5.

Enter:

```
./installmr -base_path /tmp/sfha6.0.1/
```

- 7 In the Task Menu, enter `i` to install a product.

See the 6.0.1 Installation Guide for configuration steps.

Installing the Veritas software using the script-based installer

This section describes how to install a 6.0.5 Veritas Storage Foundation and High Availability Solutions product for the first time on a host. Follow these instructions to make sure that you have the latest patches for the installer before you install or upgrade the product.

See the 6.0.1 *Installation Guide* and *Release Notes* for your product for more information.

See Upgrading to 6.0.5 from 6.0.1 or later for upgrade procedures.

To install the Veritas software for the first time

- 1 Download Storage Foundation and High Availability Solutions 6.0.1 from <http://fileConnect.symantec.com>.

Note: For 6.0.5, the base release version can be 6.0.1, 6.0.2, or 6.0.4.

- 2 Extract the tar ball into a directory called `/tmp/sfha6.0.1`.
- 3 Check <http://sort.symantec.com/patches> to see if there are any patches available for the 6.0.1 Installer. Download applicable P-patches and extract them to the `/tmp` directory.
- 4 Change to the `/tmp/sfha6.0.1` directory:

```
# cd /tmp/sfha6.0.1
```

- 5 Run the installer to install SFHA Solutions 6.0.1.

See the *Installation Guide* for instructions on installing the 6.0.1 version of this product.

```
#./installer -require complete_path_to_601_installer_patch
```

Note: If the P-patch is not available for 6.0.1 installer, use the `installer` script without `-require` option.

- 6 Download SFHA Solutions 6.0.5 from <http://sort.symantec.com/patches>.
- 7 Extract it to a directory called `/tmp/sfha6.0.5`.
- 8 Check <http://sort.symantec.com/patches> to see if there are patches available for the 6.0.5 installer. Download applicable P-patches and extract them to the `/tmp` directory.
- 9 Change to the `/tmp/sfha6.0.5` directory:

```
# cd /tmp/sfha6.0.5
```

- 10 Invoke the `installmr` script to install 6.0.5:

```
# ./installmr -require complete_path_to_605_installer_patch
```

Note: If the P-patch is not available for 6.0.5 installer, use the `installmr` script without `-require` option.

See “[About the installation script](#)” on page 10.

- 11 If you did not configure the product after the 6.0.1 installation, the installer prompts you to configure the product during MR installation. If you do not want to configure the product now, answer `n` when prompted. To configure the product in the future, run the product installation script from `/opt/VRTS/install` directory with the `-configure` option.

For configuration procedures, refer to *6.0.1 Installation Guide* for your product.

Installing the Veritas software using the Web-based installer

This section describes how to install a Veritas Storage Foundation and High Availability Solutions product for the first time on a host and then to install 6.0.5

using the Web-based installer. For detailed instructions on how to install 6.0.1 using the Web-based installer, follow the procedures in the 6.0.1 Installation Guide and Release Notes for your products.

See *Upgrading to 6.0.5 from 6.0.1 or later* for upgrade procedures.

You need to configure SFRAC before you upgrade it from 6.0.1 to 6.0.5. For more information, refer to Software Limitations in 6.0.5 *Release Notes*.

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

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

```
# ./webinstaller start
```

The `webinstaller` script displays a URL.

- 2 Start the Web browser on the system from which you want to perform the installation.
- 3 Navigate to the URL displayed from step 1.
- 4 The browser may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

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

Obtaining a security exception on Mozilla Firefox

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

To obtain a security exception

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

Installing 6.0.5 with the Veritas Web-based installer

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

To install SFHA

- 1 The 6.0.1 version of the Veritas product must be installed before upgrading to 6.0.5.

See [“Prerequisites for upgrading to 6.0.5”](#) on page 73.

Note: For 6.0.5, the base release version can be 6.0.1, 6.0.2, or 6.0.4.

- 2 On the **Select a task and a product** page, select **Install 6.0.5** from the **Task** drop-down list, and click **Next**.
- 3 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Next**.
- 4 You have the option to let the installer configure SSH or RSH communications between the systems. If you choose to allow this configuration, select the communication type and provide the root passwords for each system.
- 5 After the validation completes successfully, click **Next** to install 6.0.5 patches on the selected system.
- 6 Select the checkbox to specify whether you want to send your installation information to Symantec.

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

Click **Finish**.

Upgrading to 6.0.5 from releases earlier than 6.0.1

This chapter includes the following topics:

- [Planning to upgrade to SFHA Solutions 6.0.5](#)
- [Performing a full upgrade with Install Bundles](#)
- [Performing a phased upgrade using Install Bundles](#)
- [Performing an automated upgrade using response files with Install Bundles](#)
- [Performing a rolling upgrade using Install Bundles](#)

Planning to upgrade to SFHA Solutions 6.0.5

This section includes the following topics:

- [Supported upgrade paths for SFHA Solutions 6.0.5 from releases earlier than 6.0.1](#)
- [Pre-upgrade planning for Volume Replicator](#)
- [Prepare upgrading to 6.0.5 with Install Bundles](#)

Supported upgrade paths for SFHA Solutions 6.0.5 from releases earlier than 6.0.1

You can run the `installmr` script with Install Bundles to upgrade SFHA Solutions to 6.0.5 from releases earlier than 6.0.1.

For information on operating systems that are supported for 6.0.5, see *System requirements in Veritas Storage Foundation™ and High Availability Solutions 6.0.5 Release Notes*.

The following tables describe upgrading to 6.0.5.

Table 3-1 RHEL 5 x64 and RHEL 6 x64 upgrades from releases earlier than 6.0.1

Veritas product versions	RHEL 5	RHEL 6
5.0 MP3 5.0 MP4	Upgrade OS to RHEL 5 U5 or above. Upgrade to 6.0.5 using the <code>installmr</code> script with Install Bundles.	Not applicable.
5.1 5.1 RPs 5.1 PR1 5.1 SP1 5.1 SP1 RP1	Upgrade OS to RHEL 5 U5 or above. Upgrade to 6.0.5 using the <code>installmr</code> script with Install Bundles.	Not applicable.
5.1 SP1 PR2	Not applicable.	Upgrade OS to RHEL 6 U1 or above. Note: If VRTSvxvm is installed, upgrade VM P-patch for RHEL 6 U1. Upgrade the product to 6.0.5 using the <code>installmr</code> script with Install Bundles.
5.1 SP1 PR3 5.1SP1RP2 5.1SP1RP3 5.1SP1RP4	Upgrade OS to RHEL 5 U5 or above. Upgrade to 6.0.5 using the <code>installmr</code> script with Install Bundles.	Upgrade OS to RHEL 6 U1 or above. Note: If VRTSvxvm is installed, upgrade VM P-patch for RHEL 6 U1. Upgrade the product to 6.0.5 using the <code>installmr</code> script with Install Bundles.
6.0 6.0 RP1	Upgrade directly to 6.0.5 using the <code>installmr</code> script with Install Bundles.	Upgrade directly to 6.0.5 using the <code>installmr</code> script with Install Bundles.

Table 3-2 OL 5 x64 and OL 6 x64 upgrades from releases earlier than 6.0.1

Veritas product versions	OL 5 (RHEL Compatible mode)	OL 6 (RHEL Compatible mode)
5.0 MP3 5.0 MP4	Upgrade OS to OL 5 U5 or above. Upgrade to 6.0.5 using the <code>installmr</code> script with Install Bundles.	Not applicable.
5.1 5.1 RPs 5.1 PR1 5.1 SP1 5.1 SP1 RP1 5.1 SP1 RP2	Upgrade OS to OL 5 U5 or above. Upgrade to 6.0.5 using the <code>installmr</code> script with Install Bundles.	Not applicable.
5.1 SP1 PR2	Not applicable.	Upgrade OS to OL 6 U1 or above. Note: If VRTSvxvm is installed, upgrade VM P-patch for OL 6 U1. Upgrade the product to 6.0.5 using the <code>installmr</code> script with Install Bundles.
5.1 SP1 PR3 5.1 SP1 RP3 5.1 SP1 RP4	Upgrade OS to OL 5 U5 or above. Upgrade to 6.0.5 using the <code>installmr</code> script with Install Bundles.	Upgrade OS to OL 6 U1 or above. Note: If VRTSvxvm is installed, upgrade VM P-patch for OL 6 U1. Upgrade the product to 6.0.5 using the <code>installmr</code> script with Install Bundles.
6.0 6.0 RP1	Upgrade directly to 6.0.5 using the <code>installmr</code> script with Install Bundles.	Upgrade directly to 6.0.5 using the <code>installmr</code> script with Install Bundles.

Table 3-3 SLES 10 x64 and SLES 11 x64 upgrades from releases earlier than 6.0.1

Veritas product versions	SLES 10	SLES 11
5.0 MP3 5.0 RU4	Upgrade the operating system (OS) to SLES10 SP4. Use the <code>installmr</code> script with Install Bundles to upgrade the product to 6.0.5.	Not applicable
5.0 MP4	Upgrade OS to SLES10 SP4. Use the <code>installmr</code> script with Install Bundles to upgrade the product to 6.0.5.	Upgrade to SLES11 SP1 or above. Use the <code>installmr</code> script with Install Bundles to upgrade the product to 6.0.5.
5.0 RU1	Not applicable.	Upgrade to SLES11 SP1 or above. Use the <code>installmr</code> script with Install Bundles to upgrade the product to 6.0.5.
5.1 5.1 RPs 5.1 SP1 5.1 SP1 RPs	Upgrade OS to SLES10 SP4. Use the <code>installmr</code> script with Install Bundles to upgrade the product to 6.0.5.	Upgrade OS to SLES11 SP1 or above. Use the <code>installmr</code> script with Install Bundles to upgrade the product to 6.0.5.
5.1 PR1	Upgrade OS to SLES10 SP4. Use the <code>installmr</code> script with Install Bundles to upgrade the product to 6.0.5.	Not applicable
6.0 6.0 RP1	Use the <code>installmr</code> script with Install Bundles to upgrade the product directly to 6.0.5.	Use the <code>installmr</code> script with Install Bundles to upgrade the product directly to 6.0.5.

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.

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.

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.

If you want to manage a single cluster using Cluster Manager (Java Console), a version is available for download from:

<http://www.symantec.com/operations-manager/support>

You cannot manage the new features of this release using the Java Console. Symantec Cluster Server Management Console is deprecated.

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 3-4](#), if either the Primary or Secondary are running a version of VVR prior to 6.0.5, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 6.0.5, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

Table 3-4 VVR versions and checksum calculations

VVR prior to 6.0.5 (DG version <= 140)	VVR 6.0.5 (DG version >= 150)	VVR calculates checksum TCP connections?
Primary	Secondary	Yes
Secondary	Primary	Yes

Table 3-4 VVR versions and checksum calculations (*continued*)

VVR prior to 6.0.5 (DG version <= 140)	VVR 6.0.5 (DG version >= 150)	VVR calculates checksum TCP connections?
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.

Prepare upgrading to 6.0.5 with Install Bundles

The following procedure describes how to upgrade to 6.0.5 with Install Bundles from releases earlier than 6.0.1.

If you are upgrading from releases earlier than 6.0.1, Symantec suggest you upgrade with Install Bundles.

Note: For 6.0.5, the base release version can be 6.0.1, 6.0.2, or 6.0.4.

- 1 Download Storage Foundation and High Availability Solutions 6.0.1 from <http://fileConnect.symantec.com>.
- 2 Extract the tar ball into a directory called /tmp/sfha6.0.1
- 3 Download SFHA Solutions 6.0.5 from <http://sort.symantec.com/patches>.
- 4 Extract it to a directory called /tmp/sfha6.0.5

Note: SF Sybase CE 6.0.5 doesn't support Install Bundles.

Performing a full upgrade with Install Bundles

- [Performing a full upgrade on a VCS cluster using Install Bundles](#)
- [Performing a full upgrade of SFHA using Install Bundles](#)
- [Performing a full upgrade of SFCFSHA using Install Bundles](#)
- [Performing a full upgrade of SF Oracle RAC using Install Bundles](#)

Performing a full upgrade on a VCS cluster using Install Bundles

You can use the installer to upgrade VCS.

To upgrade VCS using the product installer

- 1 Log in as superuser.
- 2 Change to the `/tmp/sfha6.0.5` directory.
- 3 Invoke the `installmr` script with `-base_path` option to upgrade to 6.0.5:

```
# ./installmr -base_path /tmp/sfha6.0.1/
```

- 4 From the opening Selection Menu, choose: **G** for "Upgrade a Product."
- 5 Choose **1** for Full Upgrade.
- 6 Enter the names of the nodes that you want to upgrade. Use spaces to separate node names. Press the Enter key to proceed.

The installer runs some verification checks on the nodes.

- 7 When the verification checks are complete, the installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.

The installer lists the RPMs to upgrade.

- 8 The installer asks if you want to stop VCS processes. Press the Enter key to continue.

The installer stops VCS processes, uninstalls RPMs, installs or upgrades RPMs, and configures VCS.

The installer lists the nodes that Symantec recommends you to restart.

- 9 The installer asks if you would like to send the information about this installation to Symantec to help improve installation in the future. Enter your response.

The installer displays the location of log files, summary file, and response file.

Performing a full upgrade of SFHA using Install Bundles

This section describes how to perform a full upgrade of SFHA using Install Bundles.

- [Upgrading SFHA using with Install Bundles](#)

Upgrading SFHA using with Install Bundles

Use this procedure to upgrade SFHA with Install Bundles.

Note: The installer doesn't support upgrading with Install Bundles when the boot disk is encapsulated. If the boot disk is encapsulated, you need to un-encapsulate the boot disk first, and then run Install Bundles to perform upgrade.

To upgrade SFHA from previous versions to 6.0.5

- 1 Log in as superuser.
- 2 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

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

- 3 Unmount all Storage Checkpoints and file systems:

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

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

```
# echo "8192B.p S" | fsdb -t vxfs filesystem | grep clean  
flags 0 mod 0 clean clean_value
```

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

Perform the following steps in the order listed:

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

```
# fsck -t vxfs raw_device  
# mount -t vxfs block_device mountpoint  
# umount mountpoint
```

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

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

```
file system device busy
```

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

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

- If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large RPM clone can take several hours.
 - Repeat this step to verify that the unclean file system is now clean.
- 5 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
 - 6 Stop all the volumes by entering the following command for each disk group:

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

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

```
# vxprint -Aht -e v_open
```
 - 7 Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/fstab` file. You will need to recreate these entries in the `/etc/fstab` file on the freshly installed system.
 - 8 Perform any necessary preinstallation checks.
 - 9 To invoke the installer, run the `installmr` command as shown in this example:

```
# cd /tmp/sfha6.0.5  
# ./installmr -base_path /tmp/sfha6.0.1/
```
 - 10 Enter `g` to upgrade and press Return.

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

```
Enter the 64 bit <platform> system names separated  
by spaces : [q, ?] host1 host2
```

where <platform> is the platform on which the system runs, such as RHEL5.

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

During the system verification phase, the installer checks if the boot disk is encapsulated and the upgrade's path. If the upgrade is not supported, you need to un-encapsulate the boot disk.

- 12 The installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.
- 13 The installer lists the RPMs to upgrade. Then the installer asks if you want to stop SFHA processes, enter **y** to continue.

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

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

- 14 The installer stops, uninstalls, reinstalls, and starts specified RPMs.
- 15 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node that you recorded in step 7.
- 16 Restart all the volumes by entering the following command for each disk group:

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

- 17 Remount all VxFS file systems and Storage Checkpoints on all nodes:

```
# mount /filesystem  
# mount /checkpoint_name
```

- 18 You can perform the following optional configuration steps:
 - If you want to use features of Veritas Storage Foundation 6.0.5 for which you do not currently have an appropriate license installed, obtain the license and run the `vxlicinst` command to add it to your system.
 - To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.

Performing a full upgrade of SFCFSHA using Install Bundles

This section describes how to perform a full upgrade of SFCFSHA using Install Bundles.

- [Performing a full SFCFSHA upgrade with Install Bundles](#)

Performing a full SFCFSHA upgrade with Install Bundles

Performing a full upgrade involves the following tasks:

- Ensuring that the file systems are clean
- Performing the upgrade
- Updating the configuration and confirming startup

Ensuring the file systems are clean

Before upgrading to SFCFSHA 6.0.5, ensure that the file systems are clean. To ensure that the logs have been replayed and the file systems are marked clean:

To ensure the file systems are clean

- 1 Log in as superuser onto any node in the cluster.
- 2 Take the service group offline on each node of the cluster, which contains VxFS and CFS resources:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFMount resource.

Repeat this step for each SFCFSHA service group.

Note: This unmounts the CFS file systems.

- 3 Unmount all VxFS file systems not under VCS control:

```
# umount mount_point
```

- 4 Check and repair each VxFS file system:

```
# fsck -t vxfs /dev/vx/dsk/diskgroup/volume
```

The `fsck` command in `/opt/VRTS/bin` accepts either the block or character device (`/dev/vx/dsk/dg/vol`) or (`/dev/vx/rdisk/dg/vol`). The operating system version of `fsck` may limit the device types it accepts.

For more information, see the `fsck` and `fsck_vxfs` man pages.

Repeat this step for each file system.

Performing the upgrade

To perform the upgrade

- 1 Log in as superuser.
- 2 Change to the `/tmp/sfha6.0.5` directory:

```
# cd /tmp/sfha6.0.5
```

- 3 Change to the `/tmp/sfha6.0.5` directory. Invoke the `installmr` script with `-base_path` option to upgrade to 6.0.5:

```
# ./installmr -base_path /tmp/sfha6.0.1/
```

- 4 Verify there are no VxFS file systems mounted on the nodes being upgraded:

```
# mount -t vxfs
```

If any VxFS file systems are mounted, offline the group on each node of the cluster. For details, see step 2 and step 3 of [Ensuring the file systems are clean](#).

- 5 From the opening Selection Menu, choose: **G** for **Upgrade a Product**. Choose **1** for **Full Upgrade**.
- 6 You are prompted to enter the system names (in the following example, "sys1" and "sys2") on which the software is to be upgraded. Enter the system name or names and then press Return.

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

- 7 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA). Press **y** to agree and continue.
- 8 During the initial system check, the installer verifies that communication between systems has been set up.

If the installer hangs or asks for a login password, setup passwordless ssh or setup rsh from the system that run `installmr` to the system that need to be upgraded to 6.0.5. Then run the installer again.
- 9 After you accept EULA and the system checks complete, the installer displays a list of the RPMs that will be upgraded. Press Enter to continue with the upgrade.
- 10 Output shows information that SFCFSHA must be stopped on a running system. Enter **y** to continue.
- 11 The installer stops, uninstalls, reinstalls, and starts specified RPMs.
- 12 Press **Enter** again for summary information about logs and reboots.

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

Updating the configuration and confirming startup

Perform the following steps on each upgraded node.

To update the configuration and confirm startup

- 1 Remove the `/etc/VRTSvcs/conf/config/.stale` file, if it exists.

```
# rm -f /etc/VRTSvcs/conf/config/.stale
```

- 2 Verify GAB is configured:

```
# gabconfig -l | grep 'Driver.state' | \  
grep Configured  
Driver state : Configured
```

- 3 Verify VxVM daemon is started and enabled:

```
# /opt/VRTS/bin/vxdctl mode  
mode: enabled
```

- 4 Confirm all upgraded nodes are in a running state.

```
# gabconfig -a
```

- 5 After the configuration is complete, the CVM and SFCFSHA groups may come up frozen. To find out the frozen CVM and SFCFSHA groups, enter the following command:

```
# /opt/VRTS/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFSHA groups using the following commands for each group:

- Make the configuration read/write.

```
# /opt/VRTS/bin/haconf -makerw
```

- Unfreeze the group.

```
# /opt/VRTS/bin/hagrp -unfreeze group_name -persistent
```

- Save the configuration.

```
# /opt/VRTS/bin/haconf -dump -makero
```

- 6 If VVR is configured, and the CVM and SFCFSHA groups are offline, bring the groups online in the following order:

Bring online the CVM groups on all systems.

```
# /opt/VRTS/bin/hagrp -online group_name -sys sys1
```

```
# /opt/VRTS/bin/hagrp -online group_name -sys sys2
```

where *group_name* is the VCS service group that has the CVMVolDg resource.

Bring online the RVGShared groups and the virtual IP on the master node using the following commands:

```
# hagrp -online RVGShared -sys masterhost
```

```
# hares -online ip_name -sys masterhost
```

Bring online the SFCFSHA groups on all systems.

```
# /opt/VRTS/bin/hagrp -online group_name -sys sys1
```

```
# /opt/VRTS/bin/hagrp -online group_name -sys sys2
```

where *group_name* is the VCS service group that has the CFSSMount resource.

If the SFCFSHA service groups do not come online then your file system could be dirty.

- 7 Find out which node is the CVM master. Enter the following:

```
# vxdctl -c mode
```

- 8 On the CVM master node, upgrade the CVM protocol. Enter the following:

```
# vxdctl upgrade
```

Performing a full upgrade of SF Oracle RAC using Install Bundles

This section describes how to perform a full upgrade of SF Oracle RAC using Install Bundles.

- [Preparing to perform a full upgrade to 6.0.5 on an SF Oracle RAC cluster](#)
- [Upgrading SF Oracle RAC and operating system \(minor OS upgrade\)](#)

Preparing to perform a full upgrade to 6.0.5 on an SF Oracle RAC cluster

Perform the preparatory steps in this section if you are performing a full upgrade of the cluster. Before you upgrade, make sure that your systems meet the hardware and software requirements for this release.

To prepare to upgrade SF Oracle RAC

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `CRSResource.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```
# cp /etc/VRTSvcs/conf/config/main.cf \  
/etc/VRTSvcs/conf/config/main.cf.save
```

- 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 will be displayed after `installmr upgrade prechecks`.

Warning: Backup `/etc/vx/cbr/bk` directory.

4 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.

If the applications are under VCS control:

```
# hagrps -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

5 Stop all Oracle RAC resources.

- If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

```
# hagrps -offline oracle_group -any
```

- If the database instances are not managed by VCS, then run the following on one node:

- For Oracle RAC 11g:

```
$ srvctl stop database -d db_name
```

- For Oracle RAC 12c:

```
$ srvctl stop database -db db_name
```

6 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw  
# hagrps -modify oracle_group AutoStart 0  
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

- For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y manual
```

- For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy MANUAL
```

7 Stop VCS on all nodes:

```
# hastop -all
```

8 Unmount the VxFS file system, which is not under VCS control.

```
# mount |grep vxfs
```

```
# fuser -m /mount_point
```

```
# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

9 If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to 6.0.5.

See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 38.

10 If you plan to upgrade the operating system, stop all ports.

If you are running version 5.1 and later, stop the ports using the installer:

```
# /opt/VRTS/install/installsfrac -stop
```

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 RAC, you must migrate the SFDB repository database to 6.0.5.

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 Oracle RAC 6.0.5.

Perform the following before upgrading SF Oracle RAC.

To prepare to migrate the repository database

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

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

Warning: The Database Flashsnap clone database will not be able to be carried over after upgrading. You must create a new Database Flashsnap clone database after upgrading to 6.0.5.

Upgrading SF Oracle RAC and operating system (minor OS upgrade)

This section provides instructions for SF Oracle RAC and minor operating system upgrade.

Perform the steps in the following procedure if you plan to perform a minor upgrade of the operating system, for example from SLES 10 SP2 to SLES 10 SP4, along with SFRAC.

- If required, upgrade the operating system.
- Upgrade to SF Oracle RAC 6.0.5.
- Bring the SF Oracle RAC online.

Upgrading the operating system

If you want to upgrade the operating system, perform the following steps:

- 1 Rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system on all nodes in the cluster.

For instructions, see the operating system documentation.

Note: If reboot is required, use `shutdown -r now` command to reboot the nodes.

- 3 After the system restarts, restore the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

Upgrading SF Oracle RAC using Install Bundles

Use the `installmr` Veritas script-based installation programs to upgrade SF Oracle RAC.

The installer performs the following tasks to upgrade SF Oracle RAC:

- Verifies the compatibility of the systems before the upgrade.
- Stops the SF Oracle RAC processes before the upgrade.
- Uninstalls SF Oracle RAC.
- Installs the SF Oracle RAC 6.0.1 RPMs on the nodes.
- Installs the SF Oracle RAC 6.0.5 patches on the nodes.
- Displays the location of the log files, summary file, and response file.

To upgrade to SF Oracle RAC 6.0.5 using the `installmr` program

- 1 Log in as superuser.
- 2 Change to the `/tmp/sfha6.0.5` directory.
- 3 Invoke the `installmr` script with `-base_path` option to upgrade to 6.0.5:

```
# ./installmr -base_path /tmp/sfha6.0.1/
```

- 4 From the opening Selection Menu, choose: **G** for "Upgrade a Product."
- 5 Select 1 for **Full upgrade**.

The installer displays the copyright message and specifies the directory where the running logs are created.

The installer verifies the systems for compatibility.

Note: If `had` is stopped before upgrade, the installer displays the following warning:

VCS is not running before upgrade. Please make sure all the configurations are valid before upgrade.

If the configuration files are valid, you may ignore the message.

During the system verification phase, the installer checks if the boot disk is encapsulated and the upgrade path. If the upgrade is not supported, you need to un-encapsulate the boot disk.

Review the messages displayed and make sure that you meet the requirements before proceeding with the upgrade.

- 6 Press **Enter** to continue with the upgrade.

Enter **y** to agree to the End User License Agreement (EULA).

The installer displays the list of RPMs that will be uninstalled. Press **Enter** to view the list of RPMs that will be upgraded.

- 7 Enter **y** to stop the SF Oracle RAC processes.

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

The installer stops the processes and uninstalls SF Oracle RAC. After the uninstallation, the installer installs SF Oracle RAC 6.0.5 and starts 6.0.5 on all the nodes.

If you are upgrading from 5.0 releases that use regular license keys (not vxkeyless), then the installer shows the following warning. Select 'n' when prompted for additional licenses.

```
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not
updated to 6.0 on sys1. It's recommended to upgrade to a 6.0 key.
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not updated
to 6.0 on sys2. It's recommended to upgrade to a 6.0 key.
```

```
SFRAC is licensed on the systems
```

```
Do you wish to enter additional licenses? [y,n,q,b] (n) n
```

- 8 Relink the SF Oracle RAC libraries with Oracle:

The installer prompts a menu after upgrade. If you want the installer to relink the Oracle Database Binary, choose the option **Relink Oracle Database Binary** from the menu.

Complete the remaining tasks to finish the upgrade.

Bringing the Oracle database online

- 1 Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -any
```

- If the Oracle database is not managed by VCS:

- For Oracle RAC 11g:

```
# srvctl start database -d db_name
```

- For Oracle RAC 12c:

```
# srvctl start database -db db_name
```

- 2 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 3
 - If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:
 - For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y AUTOMATIC
```

- For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy AUTOMATIC
```

- 4 Complete other post-upgrade steps.

For instructions, see the chapter *Performing post-upgrade tasks* in *Veritas Storage Foundation for Oracle RAC 6.0.1 Installation and Configuration Guide*.

- 5 If you want to upgrade all application clusters to version 6.0.5, make sure that you upgraded CP server systems that use VCS or SFHA to 6.0.5. Then, upgrade all application clusters to version 6.0.5.

For instructions to upgrade VCS or SFHA on the CP server systems, see the 6.0.1 VCS or SFHA installation guide.

Performing a phased upgrade using Install Bundles

This section explains how to perform a phased upgrade of SFHA Solutions on four nodes with four service groups. Note that in this scenario, SFHA Solutions and the service groups cannot stay online on the second subcluster during the upgrade of the second subcluster. Do not add, remove, or change resources or service groups on any nodes during the upgrade. These changes are likely to get lost after the upgrade.

- [Performing a phased VCS upgrade using Install Bundles](#)

- [Performing a phased SFHA upgrade using Install Bundles](#)
- [Performing a phased SFCFSHA upgrade using Install Bundles](#)
- [Performing a phased upgrade of SF Oracle RAC using Install Bundles](#)

Performing a phased VCS upgrade using Install Bundles

You can perform a phased VCS upgrade with the following steps:

- 1 Moving the service groups to the second subcluster.
See *Veritas Cluster Server 6.0.1 Installation Guide*.
- 2 Upgrading the operating system on the first subcluster.
See *Veritas Cluster Server 6.0.1 Installation Guide*.
- 3 Upgrading the first subcluster.
See [“Step 3: Upgrading the first subcluster”](#) on page 43.
- 4 Preparing the second subcluster.
See *Veritas Cluster Server 6.0.1 Installation Guide*.
- 5 Activating the first subcluster.
See *Veritas Cluster Server 6.0.1 Installation Guide*.
- 6 Upgrading the operating system on the second subcluster.
See *Veritas Cluster Server 6.0.1 Installation Guide*.
- 7 Upgrading the second subcluster.
See [“Step 7: Upgrading the second subcluster”](#) on page 44.
- 8 Finishing the phased upgrade.
See *Veritas Cluster Server 6.0.1 Installation Guide*.

Step 3: Upgrading the first subcluster

After step 1 and step 2, you now navigate to the installer program and start it.

To start the installer for the phased upgrade

- 1 Confirm that you are logged on as the superuser and you mounted the product disc.
- 2 Make sure that you can ssh or rsh from the node where you launched the installer to the nodes in the second subcluster without requests for a password.

- 3 Start the `installvcs` program, specify the nodes in the first subcluster (`sys1` and `sys2`).

```
# cd /tmp/sfha6.0.5
# ./installmr -base_path /tmp/sfha6.0.1/ sys1 sys2
```

The program starts with a copyright message and specifies the directory where it creates the logs.

- 4 From the opening Selection Menu, choose: G for **Upgrade a Product**.
- 5 The installer performs a series of checks and tests to ensure communications, licensing, and compatibility.
- 6 Enter **y** to agree to the End User License Agreement (EULA).
- 7 The installer displays the list of RPMs that get removed, installed, or upgraded on the selected systems.
- 8 The installer ends for the first subcluster with the following output:

```
Configuring VCS: 100%

Estimated time remaining: 0:00

Performing VCS upgrade configuration ..... Done
```

Veritas Cluster Server Configure completed successfully

You are performing phased upgrade (Phase 1) on the systems. Follow the steps in install guide to upgrade the remaining systems.

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

The upgrade is finished on the first subcluster. Do not reboot the nodes in the first subcluster until you complete the Preparing the second subcluster procedure.

Step 7: Upgrading the second subcluster

After step 4 to step 6, perform the following procedure to upgrade the second subcluster (`sys3` and `sys4`).

To start the installer to upgrade the second subcluster

- 1 Confirm that you are logged on as the superuser and you mounted the product disc.

- 2 Navigate to the folder that contains the SFHA Solutions 6.0.5 binary.

```
# cd /tmp/sfha6.0.5
```

- 3 Confirm that VCS is stopped on sys3 and sys4. Specify the nodes in the second subcluster (sys3 and sys4).

```
# ./installmr -base_path /tmp/sfha6.0.1/ sys3 sys4
```

The program starts with a copyright message and specifies the directory where it creates the logs.

- 4 From the opening Selection Menu, choose: G for **Upgrade a Product**.
- 5 The installer performs a series of checks and tests to ensure communications, licensing, and compatibility.
- 6 Enter **y** to agree to the End User License Agreement (EULA).
- 7 The installer displays the list of RPMs that get removed, installed, or upgraded on the selected systems.
- 8 Monitor the installer program answering questions as appropriate until the upgrade completes.

Performing a phased SFHA upgrade using Install Bundles

You can perform a phased upgrade from SFHA 5.0 MP3 or other supported previous versions to SFHA 6.0.5.

Performing a phased upgrade involves the following tasks:

- 1 Moving the service groups to the second subcluster.
See *Veritas Storage Foundation and High Availability 6.0.1 Installation Guide*.
- 2 Upgrading the operating system on the first subcluster.
See *Veritas Storage Foundation and High Availability 6.0.1 Installation Guide*.
- 3 Upgrading the first subcluster.
See [“Upgrading the first subcluster”](#) on page 46.
- 4 Preparing the second subcluster.
See *Veritas Storage Foundation and High Availability 6.0.1 Installation Guide*.

- 5 Activating the first subcluster.
See *Veritas Storage Foundation and High Availability 6.0.1 Installation Guide*.
- 6 Upgrading the operating system on the second subcluster.
See *Veritas Storage Foundation and High Availability 6.0.1 Installation Guide*.
- 7 Upgrading the second subcluster.
See [“Upgrading the second subcluster”](#) on page 47.
- 8 Finishing the phased upgrade.
See *Veritas Storage Foundation and High Availability 6.0.1 Installation Guide*.

Upgrading the first subcluster

After step 1 and step 2, you now navigate to the installer program and start it.

To start the installer for the phased upgrade

- 1 Confirm that you are logged on as the superuser and you mounted the product disc.
- 2 Navigate to the folder that contains the SFHA Solutions 6.0.5 binary.

```
# cd /tmp/sfha6.0.5.
```
- 3 Make sure that you can ssh or rsh from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 4 Start the installmr program, specify the nodes in the first subcluster (sys1 and sys2).

```
# ./installmr -base_path /tmp/sfha6.0.1/ sys1 sys2
```

The program starts with a copyright message and specifies the directory where it creates the logs. It performs a system verification and outputs upgrade information.

- 5 From the opening Selection Menu, choose: G for **Upgrade a Product**.
- 6 Enter **y** to agree to the End User License Agreement (EULA).

- 7 The installer displays the list of RPMs that get removed, installed, and upgraded on the selected systems.
- 8 When you are prompted, reply **y** to stop appropriate processes.

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

The installer stops processes, uninstalls RPMs, and installs RPMs.

The upgrade is finished on the first subcluster. Do not reboot the nodes in the first subcluster until you complete the [Upgrading the second subcluster](#) procedure.

Upgrading the second subcluster

After step 4 to step 6, perform the following procedure to upgrade the second subcluster (sys3 and sys4).

To start the installer to upgrade the second subcluster

- 1 Confirm that you are logged on as the superuser and you mounted the product disc.
- 2 Navigate to the folder that contains the SFHA Solutions 6.0.5 binary.

```
# cd /tmp/sfha6.0.5.
```

- 3 Confirm that SFHA is stopped on sys3 and sys4. Start the installmr program, specify the nodes in the second subcluster (sys3 and sys4).

```
# ./installmr -base_path /tmp/sfha6.0.1/ sys3  
sys4
```

The program starts with a copyright message and specifies the directory where it creates the logs.

- 4 From the opening Selection Menu, choose: G for **Upgrade a Product**.
- 5 The installer displays the list of RPMs that get removed, installed, and upgraded on the selected systems.
- 6 When you are prompted, reply **y** to stop appropriate processes.

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

The installer stops processes, uninstalls RPMs, and installs RPMs.

- 7 Enter **y** to agree to the End User License Agreement (EULA).
- 8 Monitor the installer program answering questions as appropriate until the upgrade completes.

After this step, for finishing the phased upgrade, see *Veritas Storage Foundation and High Availability 6.0.1 Installation Guide*.

Performing a phased SFCFSHA upgrade using Install Bundles

Performing a phased upgrade involves the following tasks:

- 1 Moving the service groups to the second subcluster.
See Veritas Storage Foundation Cluster File System High Availability 6.0.1 Installation Guide.
- 2 Upgrading the SFCFSHA stack on the first subcluster.
See “Upgrading the SFCFSHA stack on the first subcluster” on page 48.
- 3 Preparing the second subcluster.
See Veritas Storage Foundation Cluster File System High Availability 6.0.1 Installation Guide.
- 4 Activating the first subcluster.
See Veritas Storage Foundation Cluster File System High Availability 6.0.1 Installation Guide.
- 5 Upgrading the operating system on the second subcluster.
See Veritas Storage Foundation Cluster File System High Availability 6.0.1 Installation Guide.
- 6 Upgrading the second subcluster.
See “Upgrading the second subcluster” on page 49.
- 7 Completing the phased upgrade.

Before you start the upgrade on the first half of the cluster, back up the VCS configuration files `main.cf` and `types.cf` which are in the directory `/etc/VRTSvcs/conf/config/`.

Upgrading the SFCFSHA stack on the first subcluster

After step 1, you now navigate to the installer program and start it.

To upgrade the SFCFSHA stack on the first subcluster

- 1 **Note:** This procedure is based on an "in-place" upgrade path; that is, if the operating system is upgraded, the release will be the same, but only the path level will change. If you are moving from major operating system release to another, you must uninstall the SFCFSHA stack before you upgrade the operating system. After the operating system is upgraded, you must reinstall SFCFSHA.

On the first half of the cluster, upgrade SFCFSHA by using the `installmr` script. For example use the `installmr` script with Install Bundles as shown below:

```
# ./installmr -base_path /tmp/sfha6.0.1/ sys1
```

where `<sys1>` is the node on the first subcluster.

Note: Do not reboot the nodes in the first subcluster until you complete the Preparing the second subcluster procedure.

After the upgrade for first half of the cluster is complete, no GAB ports will be shown in `gabconfig -a` output.

- 2 From the opening Selection Menu, choose: G for **Upgrade a Product**.

Upgrading the second subcluster

After step 3 to step 5, upgrade the second subcluster.

To upgrade the second subcluster

- 1 Enter the following:

```
# ./installmr -base_path /tmp/sfha6.0.1/ node_name
```

- 2 From the opening Selection Menu, choose: G for **Upgrade a Product**.

Performing a phased upgrade of SF Oracle RAC using Install Bundles

The phased upgrade methodology involves upgrading half of the nodes in the cluster at a time. The procedure involves the following tasks:

- 1 Performing pre-upgrade tasks on the first half of the cluster.
See [“Step 1: Performing pre-upgrade tasks on the first half of the cluster”](#) on page 50.
- 2 Upgrading the first half of the cluster.
[Step 2: Upgrading the first half of the cluster](#)
- 3 Performing pre-upgrade tasks on the second half of the cluster.
See [“Step 3: Performing pre-upgrade tasks on the second half of the cluster”](#) on page 54.
- 4 Performing post-upgrade tasks on the first half of the cluster.
[Step 4: Performing post-upgrade tasks on the first half of the cluster](#)
- 5 Upgrading the second half of the cluster.
[Step 5: Upgrading the second half of the cluster](#)
- 6 Performing post-upgrade tasks on the second half of the cluster.
[Step 6: Performing post-upgrade tasks on the second half of the cluster](#)

Step 1: Performing pre-upgrade tasks on the first half of the cluster

Perform the following pre-upgrade steps on the first half of the cluster.

To perform the pre-upgrade tasks on the first half of the cluster

- 1 Back up the following configuration files: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/config/types.cf.save
# cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
/etc/VRTSvcs/conf/config/OracleTypes.cf.save
# cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
/etc/VRTSvcs/conf/config/PrivNIC.cf.save
# cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
/etc/VRTSvcs/conf/config/MultiPrivNIC.cf.save
```

The installer verifies that recent backups of configuration files in the VxVM private region are saved in `/etc/vx/cbr/bk`.

If not, the following warning message is displayed: `Warning: Backup /etc/vx/cbr/bk directory.`

- 2 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.
- 3 Stop the applications configured under VCS. Stop the Oracle database:

- If the Oracle RAC instance is managed by VCS:

```
# hagrpl -offline oracle_group -sys sys1
# hagrpl -offline oracle_group -sys sys2
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the first half of the cluster and shut down the instances:

For Oracle RAC 12c:

```
$ srvctl stop instance -db db_name \
-node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \  
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \  
-i instance_name
```

- 4 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw  
# hagrps -modify oracle_group AutoStart 0  
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

For Oracle 12c:

```
$ srvctl modify database -db db_name -policy MANUAL
```

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y MANUAL
```

- 5 Unmount the CFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the CFS mount point:

```
# mount | grep vxfs | grep cluster  
# fuser -cu /mount_point
```

- Unmount the CFS file system:

```
# umount /mount_point
```

- 6 Stop the parallel service groups and switch over failover service groups on each of the nodes in the first half of the cluster:

```
# hastop -local -evacuate
```

- 7 Unmount the VxFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS mount point:

```
# mount | grep vxfs  
  
# fuser -cu /mount_point
```

- Unmount the VxFS file system:

```
# umount /mount_point
```

- 8 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.

- 9 If you plan to upgrade the operating system, stop all ports.

For 5.1x and 6.0 versions:

```
# /opt/VRTS/install/installsfrac -stop sys1 sys2
```

Step 2: Upgrading the first half of the cluster

Perform the following steps to upgrade the first half of the cluster.

To upgrade the first half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 6 On the first half of the cluster, upgrade SF Oracle RAC by using the `installmr` script. When you invoke the installer, select the **Full Upgrade** option. The installer automatically detects the phased upgrade though you select the Full Upgrade option.

```
# cd /tmp/sfha6.0.5  
  
# ./installmr -base_path /tmp/sfha6.0.1 sys1 sys2
```

Note: Do not reboot the nodes in the first subcluster until you complete the instructions for preparing the second subcluster.

See [“Step 3: Performing pre-upgrade tasks on the second half of the cluster”](#) on page 54.

After you complete the upgrade of the first half of the cluster, no GAB ports will be shown in the output when you run the `gabconfig -a` command.

If you are upgrading from 5.0 releases that use regular license keys (not vxkeyless), then the installer shows the following warning. Select 'n' when prompted for additional licenses.

```
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not  
updated to 6.0 on sys1. It's recommended to upgrade to a 6.0 key.  
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not updated  
to 6.0 on sys2. It's recommended to upgrade to a 6.0 key.  
  
SFRAC is licensed on the systems  
Do you wish to enter additional licenses? [y,n,q,b] (n) n
```

Step 3: Performing pre-upgrade tasks on the second half of the cluster

Perform the following pre-upgrade steps on the second half of the cluster.

To perform the pre-upgrade tasks on the second half of the cluster

- 1 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.

Note: The downtime starts now.

- 2 Stop all applications that are configured under VCS. Stop the Oracle database:

- If the Oracle RAC instance is managed by VCS:

```
# hagrps -offline oracle_group -sys sys3
# hagrps -offline oracle_group -sys sys4
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the second half of the cluster and shut down the instances:

For Oracle RAC 12c:

```
$ srvctl stop instance -db db_name \
-node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \
-i instance_name
```

- 3 Unmount the CFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs | grep cluster
# fuser -cu /mount_point
```

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

- 4 Stop VCS on each of the nodes in the second half of the cluster:

```
# hastop -local
```

- 5 Unmount the VxFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs
```

```
# fuser -cu /mount_point
```

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

- 6 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.

- 7 Stop all ports.

For 5.1x and 6.0 versions:

```
# /opt/VRTS/install/installsfrac -stop sys1 sys2
```

Step 4: Performing post-upgrade tasks on the first half of the cluster

Perform the following post-upgrade steps on the first half of the cluster.

To perform the post-upgrade tasks on the first half of the cluster

- 1 On any one node on the first half of the cluster, force GAB to form a cluster.

```
# /etc/init.d/llt start
```

```
# /etc/init.d/gab start
```

```
# gabconfig -x
```

- 2 On the first half of the cluster, start SF Oracle RAC:

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

```
# ./installsfrac<version> -start node1 node2
```

Where *<version>* is the specific release version.

- 3 On the first half of the cluster, manually mount the VxFS or CFS file systems that are not managed by VCS.
- 4 Relink the SF Oracle RAC libraries with Oracle:
 Choose the option **Relink Oracle Database Binary** from the program menu.
- 5 Bring the Oracle database service group online.
 If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys node_name
```

If the Oracle database is not managed by VCS:

For Oracle RAC 12c:

```
$ srvctl start instance -db db_name \  
-node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl start instance -d db_name \  
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \  
-i instance_name
```

Note: The downtime ends here.

- 6 On the first half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

Step 5: Upgrading the second half of the cluster

Perform the following steps to upgrade the second half of the cluster.

To upgrade the second half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.
 For instructions, see the operating system documentation.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/l1ttab` file to its original name:

```
# mv /etc/l1ttab.save /etc/l1ttab
```

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
- 6 On the second half of the cluster, upgrade SF Oracle RAC with the `installmr` script.

When you invoke the installer, select the **Full Upgrade** option. The installer automatically detects the phased upgrade though you select the Full Upgrade option.

```
# cd /tmp/sfha6.0.5
```

```
# ./installmr -base_path /tmp/sfha6.0.1 sys3 sys4
```

If you are upgrading from 5.0 releases that use regular license keys (not vxkeyless), then the installer shows the following warning. Select 'n' when prompted for additional licenses.

```
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not
updated to 6.0 on sys1. It's recommended to upgrade to a 6.0 key.
CPI WARNING V-9-40-5323 SFRAC license version 5.0 is not updated
to 6.0 on sys2. It's recommended to upgrade to a 6.0 key.
```

```
SFRAC is licensed on the systems
```

```
Do you wish to enter additional licenses? [y,n,q,b] (n) n
```

Step 6: Performing post-upgrade tasks on the second half of the cluster

Perform the following post-upgrade steps on the second half of the cluster.

To perform the post-upgrade tasks on the second half of the cluster

- 1 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 2 On the second half of the cluster, start SF Oracle RAC:

```
# cd /opt/VRTS/install  
  
# ./installsfrac<version> -start node3 node4
```

Where *<version>* is the specific release version.

- 3 Relink the SF Oracle RAC libraries with Oracle:

If you want the installer to relink the Oracle Database Binary, you can choose the option **Relink Oracle Database Binary** from the menu.

Complete the remaining tasks to finish the upgrade.

- 4 Upgrade VxVM disk group version.

For instructions, see the chapter "Post-upgrade tasks" in the *Veritas Storage Foundation for Oracle RAC 6.0.1 Installation and Configuration Guide*.

- 5 Upgrade disk layout version.

For instructions, see the chapter "Post-upgrade tasks" in the *Veritas Storage Foundation for Oracle RAC 6.0.1 Installation and Configuration Guide*.

- 6 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrpl -online oracle_group -sys node3  
# hagrpl -online oracle_group -sys node4
```

If the Oracle database is not managed by VCS:

For Oracle RAC 12c:

```
$ srvctl start instance -db db_name \  
-node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl start instance -d db_name \  
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \  
-i instance_name
```

- 7 ■ If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy AUTOMATIC
```

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y AUTOMATIC
```

- 8 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 9 Set or change the product license level, if required.
- 10 Migrate the SFDB repository database.

As root, dump out the old Sybase Adaptive Server Anywhere (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
```

Performing an automated upgrade using response files with Install Bundles

- [Performing an automated upgrade of VCS, SFHA, or SFCFSHA using response files with Install Bundles](#)
- [Upgrading SF Oracle RAC using a response file](#)

Performing an automated upgrade of VCS, SFHA, or SFCFSHA using response files with Install Bundles

Typically, you can use the response file that the installer generates after you perform VCS, SFHA, or SFCFSHA upgrade with Install Bundles on one system to upgrade

VCS, SFHA, or SFCFSHA on other systems. You can also create a response file using the `makeresponsefile` option of the installer.

To perform automated VCS, SFHA, or SFCFSHA upgrade

- 1 Make sure the systems where you want to upgrade VCS, SFHA, or SFCFSHA 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 SFHA.
- 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 the `/tmp/sfha6.0.5` directory. For example:

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

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

Upgrading SF Oracle RAC using a response file

Perform the steps in the following procedure to upgrade to 6.0.5 using a response file.

To upgrade SF Oracle RAC using a response file

- 1 Upgrade the operating system, if required.
For instructions, see the operating system documentation.
- 2 Create a response file using one of the available options.

Note: Make sure that you replace the host names in the response file with the names of the systems that you plan to upgrade.

For more information, refer to *Veritas Storage Foundation for Oracle RAC 6.0.1 Installation and Configuration Guide*.

- 3 Navigate to the product directory on the installation media that contains the SF Oracle RAC installation program.

4 Start the installation:

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

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

5 Complete the post-upgrade steps.

See the chapter "Performing post-upgrade tasks" in SF Oracle RAC 6.0.1 Installation Guide.

Performing a rolling upgrade using Install Bundles

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

- [Supported rolling upgrade paths](#)
- [Performing a rolling upgrade of VCS, SFHA, and SFCFSHA with Install Bundles](#)
- [Performing a rolling upgrade of SF Oracle RAC with Install Bundles](#)

Supported rolling upgrade paths

You can perform a rolling upgrade using the `installmr` script with Install Bundles.

[Table 3-5](#) shows the versions for which you can perform a rolling upgrade to 6.0.5.

Table 3-5 Supported rolling upgrade paths for Install Bundles

Platform	version
Linux	5.1, 5.1RPs 5.1SP1, 5.1SP1RPs 5.1SP1PR2, 6.0, 6.0RP1 Note: SLES doesn't support Install Bundles rolling upgrade from 5.1 and 5.1RPx to 6.0.5.

Performing a rolling upgrade of VCS, SFHA, and SFCFSHA with Install Bundles

Before you start the rolling upgrade, make sure that Veritas Cluster Server (VCS) is running.

To perform a rolling upgrade

- 1 Complete the preparatory steps on the first sub-cluster.

2 Log in as superuser.

3 Change to the `/tmp/sfha6.0.5` directory.

4 Start the installer.

```
# ./installmr -base_path /tmp/sfha6.0.1
```

5 From the menu, select `Upgrade` and from the sub menu, select `Rolling Upgrade`.

6 Enter one system of the cluster on which you would like to perform rolling upgrade.

7 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type `y` to continue.

8 The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type `y` to continue. If you choose to specify the nodes, type `n` and enter the names of the nodes.

9 The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type `y` to continue or quit the installer and address the precheck's warnings.

10 Review the end-user license agreement, and type `y` if you agree to its terms.

11 After the installer shows the package list, it detects if there are online failover service groups on the nodes to be upgraded. If there are online failover service groups, the installer prompts you to do one of the following:

- Manually switch service groups
- Use the CPI to automatically switch service groups

The downtime is the time that it normally takes for the service group's failover.

12 The installer prompts you to stop the applicable processes. Type `y` to continue.

The installer evacuates all service groups to the node or nodes that are not upgraded at this time. The installer stops parallel service groups on the nodes that are to be upgraded.

13 The installer stops relevant processes, uninstalls old kernel RPMs, and installs the new RPMs.

The installer performs the upgrade configuration and re-starts processes.

If some processes fail to start, you may need to reboot the nodes and manually check the cluster's status.

14 Complete the preparatory steps on the nodes that you have not yet upgraded.

- 15 The installer begins phase 1 of the upgrade on the remaining node or nodes. Type **y** to continue the rolling upgrade.

If the installer prompts to reboot nodes, reboot the nodes. Then, restart the installer.

The installer repeats step 8 through step 13.

For clusters with larger number of nodes, this process may repeat several times. Service groups come down and are brought up to accommodate the upgrade.

- 16 When phase 1 of the rolling upgrade completes, begin phase 2 of the upgrade. Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.

- 17 The installer determines the remaining RPMs to upgrade. Press **Enter** to continue.

- 18 The installer stops Veritas Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.

The installer performs prechecks, uninstalls old RPMs, and installs the new RPMs. It performs post-installation tasks, and the configuration for the upgrade.

- 19 A prompt message appears to ask if you would like to send the information about this installation to Symantec to help improve installation in the future?

Type **y** or **n** to help Symantec improve the installation.

- 20 If you have network connection to the Internet, the installer checks for updates.

If updates are discovered, you can apply them now.

- 21 A prompt message appears to ask if the user would like to read the summary file. You can choose **y** if you want to read the install summary file.

Performing a rolling upgrade of SF Oracle RAC with Install Bundles

Use a rolling upgrade to upgrade Veritas Storage Foundation for Oracle RAC to the latest release with minimal application downtime.

Using the `installmr` script with Install Bundles, you can upgrade to 6.0.5 from releases earlier than 6.0.1.

- [Preparing to perform a rolling upgrade to SF Oracle RAC 6.0.5](#)
- [Performing a rolling upgrade of SF Oracle RAC using the installer with Install Bundles](#)

Preparing to perform a rolling upgrade to SF Oracle RAC 6.0.5

Perform the preparatory steps in this section if you are performing a rolling upgrade of the cluster. Before you upgrade, make sure that your systems meet the hardware and software requirements for this release.

To prepare to upgrade SF Oracle RAC

Perform the steps on the first subcluster.

- 1 Log in as superuser to one of the nodes in the subcluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `CRSResource.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```
# cp /etc/VRTSvcs/conf/config/main.cf \  
/etc/VRTSvcs/conf/config/main.cf.save
```

- 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 will be displayed after `installmr upgrade prechecks`.

Warning: Backup `/etc/vx/cbr/bk` directory.

- 4 Stop the Oracle RAC resources on each node.
 - If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

```
# hagrps -offline oracle_group -sys node_name
```

- If the database instances are not managed by VCS, then run the following on one node:

For Oracle RAC 12c:

```
# srvctl stop instance -db db_name -node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \  
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \  
-i instance_name
```

- 5 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts. Failing to perform this step results in the database attempting to come online after the upgrade; the attempt fails due to the presence of old libraries on the system.

```
# haconf -makerw  
# hagrps -modify oracle_group AutoStart 0  
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy MANUAL
```

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y MANUAL
```

- 6 Unmount all the CFS file system which is not under VCS control.

```
# mount |grep vxfs | grep cluster  
  
# fuser -m /mount_point  
  
# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

- 7 Take all the parallel VCS service groups offline on each of the nodes in the current subcluster:

```
# hagrps -offline grp_name -sys sys_name
```

- 8 Unmount all the VxFS file system which is not under VCS control.

```
# mount |grep vxfs  
  
# fuser -m /mount_point  
  
# umount /mount_point
```

- 9 If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to SF Oracle RAC 6.0.5.

See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 38.

Performing a rolling upgrade of SF Oracle RAC using the installer with Install Bundles

Before you start the rolling upgrade, make sure that Veritas Cluster Server (VCS) is running.

To perform a rolling upgrade

- 1 If you want to upgrade the operating system, perform the following steps:

Note: If you plan to upgrade the operating system, make sure that you upgrade all nodes before you start rolling upgrade of SF Oracle RAC.

- Change to the `/opt/VRTS/install` directory on the node where you want to upgrade the operating system:

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

- Stop SF Oracle RAC:

```
# ./installsfrac <version> -stop
```

Where `<version>` is the specific release version.

- Upgrade the operating system. For instructions, see the operating system documentation.
- Reboot the nodes:

```
# shutdown -r now
```

- 2 Complete the preparatory steps on the first sub-cluster.

See [“Preparing to perform a rolling upgrade to SF Oracle RAC 6.0.5”](#) on page 65.

- 3 Log in as superuser.
- 4 Change to the `sfha6.0.5` directory.
- 5 Start the `installmr` script.

```
# ./installmr -base_path /tmp/sfha6.0.1/
```

- 6 From the menu, select `Upgrade` and from the sub menu, select `Rolling Upgrade`.
- 7 The installer suggests system names for the upgrade. Enter `Yes` to upgrade the suggested systems, or enter `No`, and then enter the name of any one system in the cluster on which you want to perform a rolling upgrade.
- 8 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type `y` to continue.
- 9 The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type `y` to continue. If you choose to specify the nodes, type `n` and enter the names of the nodes.
- 10 The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type `y` to continue or quit the installer and address the precheck's warnings.
- 11 Review the end-user license agreement, and type `y` if you agree to its terms.
- 12 The installer lists the RPMs to upgrade on the selected node or nodes.
- 13 After the installer shows the package list, it detects if there are online failover service groups on the nodes to be upgraded. If there are online failover service groups, the installer prompts you to do one of the following:
 - Manually switch service groups
 - Use the CPI to automatically switch service groups

The downtime is the time that it normally takes for the service group's failover.

- 14 The installer prompts you to stop the applicable processes. Type **y** to continue.
The installer evacuates all service groups to the node or nodes that are not upgraded at this time. The installer stops parallel service groups on the nodes that are to be upgraded.
- 15 The installer stops relevant processes, uninstalls old kernel RPMs, and installs the new RPMs.
The installer performs the upgrade configuration and re-starts processes.
If some processes fail to start, you may need to reboot the nodes and manually check the cluster's status.

Note: The Oracle service group is offline as the AutoStart attribute is set to 0 to prevent the service group from starting automatically. The service group is started later in the process.

- 16 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 17 Relink the SF Oracle RAC libraries with Oracle:
Choose the option **Relink Oracle Database Binary** from the program menu.
- 18 Bring the Oracle database service group online.
 - If VCS manages the Oracle database:

```
# hagrps -online oracle_group -sys node_name
```
 - If VCS does not manage the Oracle database:
For Oracle RAC 11g:

```
$ srvctl start database -d db_name
```


For Oracle RAC 12c:

```
$ srvctl start database -db db_name
```
- 19 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 20 Complete the preparatory steps on the nodes that you have not yet upgraded.

- 21 If the nodes are rebooted, restart the installer and continue phase-1 for second sub-cluster. Type **y** to continue the rolling upgrade.

The installer repeats step 7 through step 17.

For clusters with larger number of nodes, this process may repeat several times. Service groups come down and are brought up to accommodate the upgrade.

This completes phase 1 of the upgrade.

- 22 ■ If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- If VCS does not manage the Oracle database, change the management policy for the database to automatic:

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y AUTOMATIC
```

For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy AUTOMATIC
```

- 23 If the Sybase database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the service group online when VCS starts:

```
# haconf -makerw
# hagrps -modify sybasece AutoStart 1
# haconf -dump -makero
```

- 24 Migrate the SFDB repository database.

As root, dump out the old Sybase Adaptive Server Anywhere (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
```

- 25 When phase 1 of the rolling upgrade completes, begin phase 2 of the upgrade. Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.

- 26 Phase 2 of the upgrade begins here. This phase includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.
- 27 The installer determines the remaining RPMs to upgrade. Press **Enter** to continue.
- 28 The installer stops Veritas Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.

The installer performs prechecks, uninstalls old RPMs, and installs the new RPMs. It performs post-installation tasks, and the configuration for the upgrade.
- 29 When the following message appears, type **y** or **n** to help Symantec improve the installation:


```
Would you like to send the information about this installation to  
Symantec to help improve installation in the future?
```
- 30 If you have network connection to the Internet, the installer checks for updates. If updates are discovered, you can apply them now.
- 31 A prompt message appears to ask if the user would like to read the summary file. You can choose **y** if you want to read the install summary file.

Upgrading to 6.0.5 from 6.0.1 or later

This chapter includes the following topics:

- [Supported upgrade paths from 6.0.1 or later](#)
- [Prerequisites for upgrading to 6.0.5](#)
- [Downloading required software to upgrade to 6.0.5](#)
- [Performing a full upgrade to 6.0.5 on a cluster](#)
- [Upgrading to 6.0.5 on a standalone system](#)
- [Performing a phased upgrade to SFCFSHA and SFRAC](#)
- [Upgrading to 6.0.5 on a system that has encapsulated boot disk](#)
- [Performing a rolling upgrade using the installmr script](#)
- [Upgrading the operating system](#)
- [Upgrading SFHA with the Veritas Web-based installer](#)
- [Verifying software versions](#)

Supported upgrade paths from 6.0.1 or later

You can upgrade your products from the following releases to 6.0.5 with the script-based installer, the web-based installer or manual steps:

- 6.0.1 to 6.0.5
- 6.0.2 to 6.0.5

- 6.0.3 to 6.0.5
- 6.0.4 to 6.0.5

Prerequisites for upgrading to 6.0.5

If you are upgrading from 6.0.1 or 6.0.3, see the following list for prerequisites for upgrading to the 6.0.5 release:

- For any product in the Veritas Storage Foundation stack, you must have the 6.0.5 release binaries.
- Each system must have sufficient free space to accommodate patches.
- The full list of prerequisites can be obtained by running `./installmr -precheck`.
- Make sure to download the latest patches for the installer.

Downloading required software to upgrade to 6.0.5

This section describes how to download the latest patches for the installer.

To download required software to upgrade to 6.0.5

- 1 Download SFHA 6.0.5 from <http://sort.symantec.com/patches>.
- 2 Extract it to a directory, say `/tmp/sfha605`.

Performing a full upgrade to 6.0.5 on a cluster

Performing a full upgrade on a cluster requires stopping cluster failover functionality during the entire procedure.

Depending on your cluster's configuration, select one of the following procedures to upgrade to 6.0.5:

- [Performing a full upgrade to 6.0.5 on a Veritas Cluster Server](#)
- [Performing a full upgrade to 6.0.5 on an SFHA cluster](#)
- [Performing a full upgrade to 6.0.5 on an SFCFSHA cluster](#)
- [Performing a full upgrade to 6.0.5 on an SF Oracle RAC cluster](#)

Performing a full upgrade to 6.0.5 on a Veritas Cluster Server

The following procedure describes performing a full upgrade on a Veritas Cluster Server (VCS) cluster.

Note: You need to make sure that IPv4RouteOptions attribute is configured for MultiNICA resources, otherwise network connection may be interrupted.

To upgrade VCS

- 1 Make sure you have downloaded the latest software required for the upgrade.
- 2 Log in as superuser.

Note: Upgrade the Operating System and reboot the systems if required.

- 3 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 6.0.5 rolling patch binaries, change to the directory that contains the `installmr` script. Start the pre-upgrade check:

```
# ./installmr -precheck node1 node2 ... nodeN
```

- 4 Resolve any issues that the precheck finds.
- 5 Start the upgrade:

```
# ./installmr node1 node2 ... nodeN
```

- 6 After the upgrade, review the log files for any issues.

Performing a full upgrade to 6.0.5 on an SFHA cluster

The following procedure describes performing a full upgrade on an SFHA and VCS cluster.

To perform a full upgrade to 6.0.5 on an SFHA cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
- 2 Log in as superuser.
- 3 From the directory that contains the extracted and untarred 6.0.5 rolling patch binaries, change to the directory that contains the `installmr` script. Check the readiness of the nodes where you plan to upgrade. Start the pre-upgrade check:

```
# ./installmr -precheck node1 node2... nodeN
```

where `node1`, `node2` and `nodeN` are nodes to be upgraded.

- 4 If service groups have VxFS file systems mounted, make the service groups offline.
- 5 Start the upgrade:

```
# ./installmr node1 node2... nodeN
```

where *node1*, *node2* and *nodeN* are nodes to be upgraded.

Performing a full upgrade to 6.0.5 on an SFCFSHA cluster

The following procedure describes performing a full upgrade on an SFCFSHA cluster.

To perform a full upgrade to 6.0.5 on an SFCFSHA cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
- 2 Log in as superuser.
- 3 If applications are not managed by VCS, make the applications offline.
- 4 On each node, enter the following command to check if any Storage Checkpoints or VxFS file systems out of VCS control are mounted:

```
# mount | grep vxfs
```

If any Storage Checkpoints or VxFS file systems are mounted, on each node in the cluster unmount all Storage Checkpoints.

```
# umount /checkpoint_name
```

- 5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:
 - Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
 - Use the vxrvrg stop command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

- On the Primary node, use the vxrlink status command to verify that all RLINKs are up-to-date:

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

Caution: To avoid data corruption, do not proceed until all RLINKs are up-to-date.

- 6 Stop activity to all VxVM volumes.

For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes

- 7 Stop all VxVM volumes by entering the following command for each disk group on master node:

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

Verify that no volumes remain open:

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

- 8 On each node, stop the VCS command server:

```
# ps -ef | grep CmdServer  
# kill -9 pid_of_CmdServer
```

pid_of_CmdServer is the process ID of CmdServer.

- 9 If required, apply the OS kernel patches.
- 10 From the directory that contains the extracted and untarred 6.0.5 rolling patch binaries, change to the directory that contains the `installmr` script. Start the upgrade.

```
# ./installmr node1 node2
```

where *node1* and *node2* are nodes which are to be upgraded.

- 11 After all nodes in the cluster are upgraded, the processes will be restarted automatically. Should there be any problem, the `installmr` script will ask you to reboot the system. Then the application failover capability will be available.
- 12 If necessary, reinstate any missing mount points in the `/etc/filesystems` file on each node.
- 13 Restart all the volumes by entering the following command on the master node, for each disk group:

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

- 14 If you stopped any RVGs in step 5, restart each RVG:

```
# vxrvrg -g diskgroup start rvg_name
```

- 15 Remount all VxFS file systems and Storage Checkpoints on all nodes:

```
# mount /filesystem
```

Performing a full upgrade to 6.0.5 on an SF Oracle RAC cluster

The following procedure describes performing a full upgrade on an SF for Oracle RAC cluster.

To upgrade to 6.0.5 on an SF Oracle RAC cluster

- 1 Make sure you have downloaded the latest software required for the upgrade.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTSvcs/bin` is in your `PATH` so that you can execute all product commands.
- 4 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 5 Enter the following command to freeze HA service group operations on each node:

```
# hasys -freeze -persistent nodename
```

- 6 Make the configuration read-only:

```
# haconf -dump -makero
```

- 7
 - If the Oracle database is managed by VCS, set the `AutoStart` value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 0
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:
 For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy MANUAL
```

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y MANUAL
```

- 8 If Oracle Clusterware is not controlled by VCS, enter the following command on each node of the cluster to stop Oracle Clusterware:

```
# GRID_HOME/bin/crsctl stop crs
```

- 9 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS. If the applications are under VCS control:

```
# hagrps -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

- 10 Unmount the VxFS file system, which is not under VCS control.

```
# mount |grep vxfs
```

```
# fuser -m /mount_point
```

```
# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

- 11 Stop VCS.

```
# hastop -all
```

- 12 If required, apply the OS kernel patches.

See *Oracle's* documentation for the procedures.

- 13 From the directory that contains the extracted and untarred 6.0.5 rolling patch binaries, change to the directory that contains the installmr script. Enter:

```
# ./installmr node1 node2
```

where *node1* and *node2* are nodes which are to be upgraded.

- 14 Follow the instructions from the installer. If there is some module load/unload issue, reboot all of the nodes of the cluster.

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

- 15 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node.

- 16 Relink the SF Oracle RAC libraries with Oracle:

Choose the option **Relink Oracle Database Binary** from the program menu.

- 17 From any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 18 Enter the following command on each node to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent nodename
```

- 19 Make the configuration read-only:

```
# haconf -dump -makero
```

- 20 Enter the following command on each node to take service groups online:

```
# hagrps -online service_group -sys nodename
```

- 21 Restart all the volumes by entering the following command for each disk group:

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

- 22 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

- 23 If Oracle Clusterware is not controlled by VCS, enter the following command on each node to start Oracle Clusterware.

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

24 Check if the VEA service was restarted:

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

If the VEA service is not running, restart it:

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

25 Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

```
# hagrps -online Oracle_group -any
```

- If the Oracle database is not managed by VCS:

For Oracle RAC 11g:

```
$ srvctl start database -d db_name
```

For Oracle RAC 12c:

```
$ srvctl start database -db db_name
```

- 26** ■ If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
```

```
# hagrps -modify oracle_groupname AutoStart 1
```

```
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y AUTOMATIC
```

For Oracle RAC 12c:


```
$ srvctl modify database -db db_name -policy AUTOMATIC
```

27 Upgrade Oracle RAC.

Note: Oracle RAC 11g Release 1 Clusterware is not supported. Make sure that you install Oracle RAC 11g Release 2 Grid Infrastructure in order to use the Oracle RAC 11g Release 1 database. All database versions starting from Oracle 11g Release 2 and later are supported.

For instructions, see the chapter *Upgrading Oracle RAC* in *Veritas Storage Foundation™ for Oracle® RAC Installation and Configuration Guide*.

Note: The procedure for Oracle RAC 12c is the same with that for Oracle RAC 11g Release 2.

Upgrading to 6.0.5 on a standalone system

You can use this procedure to upgrade on a standalone system that runs SF.

To upgrade to 6.0.5 on a standalone system

- 1 Make sure you have downloaded the latest software required for the upgrade.
- 2 Log in as superuser.
- 3 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.
- 4 If required, apply the OS kernel patches.
- 5 Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

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

- 6 Unmount all Storage Checkpoints and file systems:

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

- 7 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

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

Caution: To avoid data corruption, do not proceed until all RLINKs are up-to-date.

- 8 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

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

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

Verify that no volumes remain open:

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

- 10 Check if the VEA service is running:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is running, stop it:

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

- 11 Navigate to the folder that contains the installation program. Run the `installmr` script:

```
# ./installmr nodename
```

- 12 If necessary, reinstate any missing mount points in the `/etc/fstab` file.

- 13 Restart all the volumes by entering the following command for each disk group:

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

14 If you stopped any RVGs in step 7, restart each RVG:

```
# vxrvrg -g diskgroup start rvg_name
```

15 Remount all VxFS file systems and Storage Checkpoints:

```
# mount /filesystem
# mount /checkpoint_name
```

16 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctrl status
```

If the VEA service is not running, restart it:

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

Performing a phased upgrade to SFCFSHA and SFRAC

This section describes how to perform a phased upgrade to SFCFSHA and SFRAC.

Performing a phased upgrade of SFCFSHA

Performing a phased upgrade involves the following tasks:

- Preparing the first subcluster
- Upgrading the SFCFSHA stack on the first subcluster
- Preparing the second subcluster
- Activating the first subcluster
- Upgrading the operating system on the second subcluster
- Upgrading the second subcluster
- Finishing the phased upgrade

Before you start the upgrade on the first half of the cluster, back up the VCS configuration files `main.cf` and `types.cf` which are in the directory `/etc/VRTSvcs/conf/config/`.

Prerequisites for a phased upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade. Note that your applications have downtime during this procedure.

Planning for a phased upgrade

Plan out the movement of the service groups from node-to-node to minimize the downtime for any particular service group. Some basic guidelines follow:

- Split the cluster in half. If the cluster has an odd number of nodes, calculate $(n+1)/2$, and start the upgrade with the even number of nodes.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.

Phased upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the phased upgrade:

- While you perform the upgrades, do not start any modules. Also, do not add or remove service groups to any of the nodes.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.
- You can perform a phased upgrade when the root disk is encapsulated.

Preparing the first subcluster

To move the service groups to the second subcluster

- 1 Switch failover groups from the first half of the cluster to one of the nodes in the second half of the cluster. In this procedure, `sys1` is a node in the first half of the cluster and `sys4` is a node in the second half of the cluster. Enter the following:

```
# hagrps -switch failover_group -to sys4
```

- 2 On the first half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the applications.

- 3 On the first half of the cluster, unmount the VxFS or CFS file systems that are not managed by VCS.

```
# mount | grep vxfs
```

Verify that no processes use the VxFS or CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS or CFS mount point with the mechanism provided by the application.

Unmount the VxFS or CFS file system. Enter the following:

```
# umount /mount_point
```

- 4 On the first half of the cluster, bring all the VCS service groups offline including CVM group. Enter the following:

```
# hagrps -offline group_name -sys sys1 sys2
```

When the CVM group becomes OFFLINE, all the parallel service groups such as the CFS file system will also become OFFLINE on the first half of the cluster nodes.

- 5 Verify that the VCS service groups are offline on all the nodes in first half of the cluster. Enter the following:

```
# hagrps -state group_name
```

- 6 Freeze the nodes in the first half of the cluster. Enter the following:

```
# haconf -makerw
```

```
# hasys -freeze -persistent sys1
```

```
# haconf -dump -makero
```

- 7 If I/O fencing is enabled, then on each node of the first half of the cluster, change the contents of the `/etc/vxfenmode` file to configure I/O fencing in disabled mode.

If you are in the CPS fencing mode, Symantec recommends you making a backup of the `/etc/vxfenmode` file for future use:

```
# cp /etc/vxfenmode /etc/vxfenmode.bak
```

To change the contents of the `/etc/vxfenmode` file to configure I/O fencing in disabled mode, enter the following:

```
# cp /etc/vxfen.d/vxfenmode_disabled /etc/vxfenmode
# cat /etc/vxfenmode
[root@swlx08 ~]# cat /etc/vxfenmode
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=disabled
```

- 8 If the cluster-wide attribute `UseFence` is set to `SCSI3`, then reset the value to `NONE` in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.
- 9 Verify that only GAB ports a, b, d and h are open. Enter the following:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 6b5901 membership 01
Port b gen 6b5904 membership 01
Port d gen 6b5907 membership 01
Port h gen ada40f membership 01
```

Do not stop VCS. Port h should be up and running.

- 10 In the first half of the cluster, stop all VxVM and CVM volumes. Enter the following command for each disk group:

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

Verify that no volumes remain open. Enter the following:

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

- 11 On first half of the cluster, upgrade the operating system on all the nodes, if applicable. For instructions, see the upgrade paths for the operating system.

Upgrading the SFCFSHA stack on the first subcluster

To upgrade the SFCFSHA stack on the first subcluster

- ◆ **Note:** This procedure is based on an "in-place" upgrade path; that is, if the minor operating system is upgraded, the release will be the same, but only the path level will change. If you are moving from major operating system release to another, you must uninstall the SFCFSHA stack before you upgrade the operating system. After the operating system is upgraded, you must reinstall SFCFSHA.

On the first half of the cluster, upgrade SFCFSHA by using the `installmr` script. For example use the `installmr` script as shown below:

```
# ./installmr sys1 sys2
```

where `<sys1>` and `<sys2>` are the nodes on the first subcluster.

After the upgrade for first half of the cluster is complete, no GAB ports will be shown in `gabconfig -a` output.

To upgrade your operating system, follow the normal procedures for your platform.

Note: After the installation completes, you can safely ignore any instructions that the installer displays.

Preparing the second subcluster

To prepare the second subcluster

- 1 On the second half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application. [Downtime starts now.]
- 2 On the second half of the cluster, unmount the VxFS and CFS file systems that are not managed by VCS. Enter the following:

```
# mount | grep vxfs
```

Verify that no processes use the VxFS and CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS and CFS mount point with the mechanism provided by the application.

Unmount the VxFS and CFS file system. Enter the following:

```
# umount /mount_point
```

- 3 On the second half of the cluster, unfreeze all the VCS service groups on all the nodes using the following commands:

```
# haconf -makerw  
# hagrps -unfreeze group_name -persistent  
# haconf -dump -makero
```

- 4 On the second half of the cluster, bring all the VCS service groups offline, including CVM group. Enter the following:

```
# hagrps -offline group_name -sys sys3 sys4
```

- 5 On the second half of the cluster, verify that the VCS service groups are offline. Enter the following:

```
# hagrps -state group_name
```

- 6 Stop VCS on the second half of the cluster. Enter the following:

```
# hastop -local
```


- 7 On each node of the second half of the cluster, change the contents of the `/etc/vxfenmode` file to configure I/O fencing in disabled mode.

```
# cp /etc/vxfen.d/vxfenmode_disabled /etc/vxfenmode
# cat /etc/vxfenmode#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=disabled
```

- 8 If the cluster-wide attribute `UseFence` is set to `SCSI3`, reset the value to `NONE` in the `/etc/VRTSvcs/conf/config/main.cf` file, in second half of the cluster.
- 9 On the second half on cluster, stop the following SFCFSHA modules: `VCS`, `VxFEN`, `ODM`, `GAB`, and `LLT`. Enter the following:

```
# /etc/init.d/vxglm stop
# /etc/init.d/vxodm stop
# /etc/init.d/vxgms stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

- 10** On each node in the first half of the cluster, enable fencing. Replace the `/etc/vxfemode` file with your previous `vxfenmode` file. Take the `scsi3_dmp` file as an example:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode
# cat /etc/vxfenmode
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# sybase     - use scsi3 disks in kernel but coordinate membership
# with Sybase ASE
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=scsi3
#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
scsi3_disk_policy=dmp
```

- 11** If the cluster-wide attribute `UseFence` is set to `NONE`, reset the value to `SCSI3` in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.

Activating the first subcluster

To activate the first subcluster

- 1 Restart the upgraded nodes in the first half of the cluster:

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

When the first half of the cluster nodes come up, no GAB ports are OPEN. The following command does not show any GAB ports:

```
# /sbin/gabconfig -a
GAB Port Memberships
=====
```

- 2 If required, force gab to form a cluster after the upgraded nodes are rebooted in first half of the cluster.

```
# /sbin/gabconfig -x
```

GAB ports a, b, d and h appear in `gabconfig -a` command output.

Note: If port b and h are not up, you need to bring fencing and VCS manually online.

- 3 On the first half of the cluster, bring the VCS service groups online. Enter the following:

```
# hagrps -online group_name -sys node_name
```

After you bring the CVM service group ONLINE, all the GAB ports u, v, w and f come ONLINE and all the CFS mounts service groups also come ONLINE automatically. Only failover service groups need to be brought ONLINE manually.

- 4 Manually mount the VxFS and CFS file systems that are not managed by VCS. [Downtime ends now.]

Upgrading the operating system on the second subcluster

To upgrade the operating system on the second subcluster

- ◆ Enter the following.

```
# chkconfig vcs off
# chkconfig vxfen off
# chkconfig gab off
# chkconfig llt off
```

On the second half of the cluster, upgrade the operating system, if applicable. For instructions, see the upgrade paths for the operating system.

Upgrading the second subcluster

To upgrade the second subcluster

- ◆ Enter the following:

```
# ./installmr sys3 sys4
```

Completing the phased upgrade

To complete the phased upgrade

- 1 Verify that the cluster UUID on the nodes in the second subcluster is the same as the cluster UUID on the nodes in the first subcluster. Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -display nodename
```

If the cluster UUID differs, manually copy the cluster UUID from a node in the first subcluster to the nodes in the second subcluster. For example:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -copy -from_sys \  
node01 -to_sys node03 node04
```

- 2 On each node in the second half of the cluster, enable fencing. Enter the following:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode  
# cat /etc/vxfenmode  
#  
# vxfen_mode determines in what mode VCS I/O Fencing should work.  
#  
# available options:  
# scsi3      - use scsi3 persistent reservation disks  
# customized - use script based customized fencing  
# sybase     - use scsi3 disks in kernel but coordinate  
#             membership with Sybase ASE  
# disabled   - run the driver but don't do any actual fencing  
#  
vxfen_mode=scsi3  
#  
# scsi3_disk_policy determines the way in which I/O Fencing  
# communicates with the coordination disks.  
#  
# available options:  
# dmp - use dynamic multipathing  
# raw - connect to disks using the native interface  
#  
scsi3_disk_policy=dmp
```

3 Enter the following before reboot of nodes:

```
# chkconfig vcs on
# chkconfig vxfs on
# chkconfig gab on
# chkconfig llt on
```

4 Restart the upgraded nodes in the second half of the cluster:

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

When second half of the nodes come up, all the GAB ports a, b, d, h, u, v, w and f are ONLINE. Also all the CFS mounts service groups come online automatically.

5 Manually mount the VxFS and CFS file systems that are not managed by VCS in the second half of the cluster.

6 Find out which node is the CVM master. Enter the following:

```
# vxctl -c mode
```

Performing phased upgrade of SF Oracle RAC to version 6.0.5

Table 4-1 illustrates the phased upgrade process. Each column describes the steps to be performed on the corresponding subcluster and the status of the subcluster when operations are performed on the other subcluster.

Table 4-1 Summary of phased upgrade

First half of the cluster	Second half of the cluster
---------------------------	----------------------------

SF Oracle RAC cluster before the upgrade:




Table 4-1 Summary of phased upgrade (*continued*)





First half of the cluster	Second half of the cluster
<p>STEP 1: Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"> Switch failover applications. Stop all parallel applications. <p>See “Step 1: Performing pre-upgrade tasks on the first half of the cluster” on page 96.</p> <p>STEP 2: Upgrade SF Oracle RAC.</p> <p>See “Step 2: Upgrading the first half of the cluster” on page 98.</p>	<p>The second half of the cluster is up and running.</p> 
<p>The first half of the cluster is not running.</p> 	<p>STEP 3: Perform the following pre-upgrade steps:</p> <ul style="list-style-type: none"> Stop all parallel and failover applications. Stop SF Oracle RAC. <p>See “Step 3: Performing pre-upgrade tasks on the second half of the cluster” on page 99.</p> <p>The downtime starts now.</p>
<p>STEP 4: Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"> Start SF Oracle RAC. Start all applications. <p>See “Step 4: Performing post-upgrade tasks on the first half of the cluster” on page 101.</p> <p>The downtime ends here.</p>	<p>The second half of the cluster is not running.</p> 
<p>The first half of the cluster is up and running.</p> 	<p>STEP 5: Upgrade SF Oracle RAC.</p> <p>See “Step 5: Upgrading the second half of the cluster” on page 102.</p> <p>STEP 6: Perform the following post-upgrade steps:</p> <ul style="list-style-type: none"> Start SF Oracle RAC. Start all applications. <p>See “Step 6: Performing post-upgrade tasks on the second half of the cluster” on page 103.</p>

Table 4-1 Summary of phased upgrade (*continued*)

First half of the cluster	Second half of the cluster
---------------------------	----------------------------

The phased upgrade is complete and both the first and the second half of the cluster are running.



Step 1: Performing pre-upgrade tasks on the first half of the cluster

Perform the following pre-upgrade steps on the first half of the cluster.

To perform the pre-upgrade tasks on the first half of the cluster

- 1 Back up the following configuration files: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

```

# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/config/types.cf.save
# cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
/etc/VRTSvcs/conf/config/OracleTypes.cf.save
# cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
/etc/VRTSvcs/conf/config/PrivNIC.cf.save
# cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
/etc/VRTSvcs/conf/config/MultiPrivNIC.cf.save
    
```

- 2 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 will be displayed after `installmr` upgrade prechecks.

Warning: Backup `/etc/vx/cbr/bk` directory.

3 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.

4 Stop the applications configured under VCS. Stop the Oracle RAC database:

- If the Oracle RAC instance is managed by VCS:

```
# hagrpl -offline oracle_group -sys sys1
# hagrpl -offline oracle_group -sys sys2
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the first half of the cluster and shut down the instances:

For Oracle RAC 12c:

```
$ srvctl stop instance -db db_name \
-node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \
-i instance_name
```

- 5 ■ If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

```
# haconf -makerw
# hagrpl -modify oracle_group AutoStart 0
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy MANUAL
```

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y MANUAL
```

- 6 Unmount the CFS file systems that are not managed by VCS.
 - Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs | grep cluster
# fuser -cu /mount_point
```

- Unmount the non-system CFS file system:

```
# umount /mount_point
```

- 7 Stop the parallel service groups and switch over failover service groups on each of the nodes in the first half of the cluster:

```
# hastop -local -evacuate
```

- 8 Unmount the VxFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs
# fuser -cu /mount_point
```

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

- 9 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.

- 10 If you plan to upgrade the operating system, stop all ports.

```
# /opt/VRTS/install/installsfrac -stop sys1 sys2
```

Step 2: Upgrading the first half of the cluster

Perform the following steps to upgrade the first half of the cluster.

To upgrade the first half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.
For instructions, see the operating system documentation.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the second subcluster without requests for a password.

- 6 Upgrade SF Oracle RAC. On the first half of the cluster, upgrade SFRAC by using the `installmr` script. For example use the `installmr` script as shown below:

```
# ./installmr sys1 sys2
```

Note: After you complete the upgrade of the first half of the cluster, no GAB ports will be shown in the output when you run the `gabconfig -a` command.

Step 3: Performing pre-upgrade tasks on the second half of the cluster

Perform the following pre-upgrade steps on the second half of the cluster.

To perform the pre-upgrade tasks on the second half of the cluster

- 1 Stop all applications that are not configured under VCS but dependent on Oracle RAC or resources controlled by VCS. Use native application commands to stop the application.

Note: The downtime starts now.

- 2 Stop all applications that are configured under VCS. Stop the Oracle RAC database:

- If the Oracle RAC instance is managed by VCS:

```
# hagrps -offline oracle_group -sys sys3
# hagrps -offline oracle_group -sys sys4
```

- If the Oracle RAC instance is not managed by VCS, log in as the Oracle user on one of the nodes in the second half of the cluster and shut down the instances:

For Oracle RAC 12c:

```
# srvctl stop instance -db db_name \
-node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \
-i instance_name
```

- 3 Unmount the CFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs | grep cluster

# fuser -cu /mount_point
```

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

- 4 Stop VCS on each of the nodes in the second half of the cluster:

```
# hstop -local
```

- 5 Unmount the VxFS file systems that are not managed by VCS.

- Make sure that no processes are running which make use of mounted shared file system. To verify that no processes use the VxFS or CFS mount point:

```
# mount | grep vxfs
```

```
# fuser -cu /mount_point
```

- Unmount the non-system VxFS file system:

```
# umount /mount_point
```

- 6 Verify that no VxVM volumes (other than VxVM boot volumes) remain open. Stop any open volumes that are not managed by VCS.

- 7 Stop all ports.

```
# /opt/VRTS/install/installsfrac -stop sys3 sys4
```

Step 4: Performing post-upgrade tasks on the first half of the cluster

Perform the following post-upgrade steps on the first half of the cluster.

To perform the post-upgrade tasks on the first half of the cluster

- 1 On any one node on the first half of the cluster, force GAB to form a cluster.

```
# /etc/init.d/llt start
```

```
# /etc/init.d/gab start
```

```
# gabconfig -x
```

- 2 On the first half of the cluster, start SF Oracle RAC:

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

```
# ./installsfrac605 -start sys1 sys2
```

- 3 On the first half of the cluster, manually mount the VxFS or CFS file systems that are not managed by VCS.
- 4 Relink the SF Oracle RAC libraries with Oracle.
- 5 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys node_name
```

If the Oracle database is not managed by VCS:

For Oracle RAC 12c:

```
$ srvctl start instance -db db_name \  
-node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl start instance -d db_name \  
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \  
-i instance_name
```

Note: The downtime ends here.

- 6 On the first half of the cluster, start all applications that are not managed by VCS. Use native application commands to start the applications.

Step 5: Upgrading the second half of the cluster

Perform the following steps to upgrade the second half of the cluster.

To upgrade the second half of the cluster

- 1 If you plan to upgrade the operating system, rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system, if required.

For instructions, see the operating system documentation.

- 3 If you upgraded the operating system, restart the nodes:

```
# shutdown -r now
```

- 4 Rename the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

- 5 Make sure that you can run secure shell or remote shell from the node where you launched the installer to the nodes in the first subcluster without requests for a password.
- 6 On the second half of the cluster, upgrade SF Oracle RAC. Navigate to the product directory on the installation media.

Upgrade SFRAC by using the `installmr` script. For example use the `installmr` script as shown below:

```
# ./installmr sys3 sys4
```

Step 6: Performing post-upgrade tasks on the second half of the cluster

Perform the following post-upgrade steps on the second half of the cluster.

To perform the post-upgrade tasks on the second half of the cluster

- 1 Manually mount the VxFS and CFS file systems that are not managed by VCS.
- 2 On the second half of the cluster, start SF Oracle RAC:

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

```
# ./installsfrac605 -start sys3 sys4
```

- 3 Relink the SF Oracle RAC libraries with Oracle.

4 Upgrade VxVM disk group version.

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions and can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks, you need to upgrade existing disk group version to 180.

Check the existing disk group version:

```
# vxvg list dg_name | grep -i version
```

If the disk group version is not 180, run the following command on the master node to upgrade the version:

```
# vxvg -T 180 upgrade dg_name
```

5 Upgrade disk layout version.

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

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

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

Support for disk layout Version 4 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.

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

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

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

6 Bring the Oracle database service group online.

If the Oracle database is managed by VCS:

```
# hagrps -online oracle_group -sys sys3
# hagrps -online oracle_group -sys sys4
```

If the Oracle database is not managed by VCS:

For Oracle RAC 12c:

```
$ srvctl start instance -db db_name \
node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl start instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \
-i instance_name
```

- 7 ■ If the Oracle database is managed by VCS, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 1
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to automatic:

For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy AUTOMATIC
```

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y AUTOMATIC
```

- 8 Start all applications that are not managed by VCS. Use native application commands to start the applications.

9 Set or change the product license level, if required.

10 Migrate the SFDB repository database.

As root, dump out the old Sybase Adaptive Server Anywhere (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
```

Upgrading to 6.0.5 on a system that has encapsulated boot disk

You can use this procedure to upgrade to 6.0.5 on a system that has encapsulated boot disk.

Note: Upgrading with encapsulated boot disk from 6.0.1 to 6.0.5 requires reboot.

To upgrade to 6.0.5 on a system that has encapsulated boot disk

- 1 Manually unmount file systems and stop open volumes.
- 2 Upgrade to 6.0.5 using `installmr` command.
- 3 After upgrading, reboot the system to have the new VM drivers take effect.
- 4 If the mirror of boot disk is split during upgrading, re-join the mirrors manually when systems reboot after upgrading to 6.0.5.

Performing a rolling upgrade using the installmr script

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

- [About rolling upgrades](#)
- [Prerequisites for a rolling upgrade](#)
- [Performing a rolling upgrade on kernel packages for VCS, SFHA and SFCFSHA: phase 1](#)
- [Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSHA : phase 2](#)
- [Performing a rolling upgrade on kernel packages for SF Oracle RAC: phase 1](#)

- [Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2](#)

About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime.

Rolling upgrades take two discrete phases. In the first, you upgrade the kernel packages with exception of VCS packages and agent packages. In the second, you upgrade the non-kernel packages, which are VCS packages and agents packages.

You can perform a rolling upgrade for the following products:

- Veritas Cluster Server
- Storage Foundation and High Availability
- Storage Foundation Cluster File System High Availability
- Storage Foundation for Oracle RAC
- Symantec VirtualStore
- Storage Foundation for Sybase CE

Prerequisites for a rolling upgrade

Meet the following prerequisites before performing a rolling upgrade:

- Make sure that the product you want to upgrade supports rolling upgrades.
- Make a plan on splitting up your clusters into sub-clusters for the purpose of upgrade. A sub-cluster can include one or more nodes. This division helps to keep service groups running during the upgrade.
- Make sure you are logged in as superuser.
- VCS must be running before performing the rolling upgrade.
- Make sure you have downloaded the latest software required for the upgrade.

Limitation: During VCS and agents upgrade, you must bring down the application High Availability (HA) for several minutes. This does not affect the application running on the cluster. You can restore the application's high availability after VCS and the agent packages are upgraded.

Performing a rolling upgrade using the `installmr` script

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

Performing a rolling upgrade on kernel packages for VCS, SFHA and SFCFSA: phase 1

Note that in the following instructions a sub-cluster can represent one or more nodes in a full cluster, but is represented by `nodeA` as `subcluster1` and `nodeB` as `subcluster2`.

To perform the rolling upgrade on kernel packages: phase 1

- 1 On the first sub-cluster, start the installer for the rolling upgrade with the `-rollingupgrade_phase1` option.

```
# ./installmr -rollingupgrade_phase1 nodeA
```

You can also use the web-based installer to perform the upgrade. If you are using the web-based installer, start the web-based installer with the `./webinstaller start` command, select **Rolling Upgrade** from the task list, make sure the **Phase-1: Upgrade Kernel packages** radio button is checked, and then click **Next**.

- 2 Note that if the boot-disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
- 3 The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.
If you are using the web-based installer, input one node of the cluster. The web-based installer detects the whole cluster, and then recommend some nodes (NodeA) as the subcluster to run the rolling upgrade phase 1. The web-based installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel packages.
- 4 The installer loads new kernel modules and starts all the relevant processes and brings all the service groups online.
- 5 If the boot disk is encapsulated, reboot the first sub-cluster's system. Otherwise go to step 6.

- 6** After rolling upgrade phase 1 is completed on nodeA, the following message displays:

It is recommended to perform rolling upgrade phase 1 on the systems nodeB in the next step.

Would you like to perform rolling upgrade phase 1 on the systems?
[y,n,q] (y)

If you choose *y*, it continues to run rolling upgrade phase 1 by itself on nodeB.

If you choose *n* or *q*, you need to complete step 1 to step 4 on nodeB.

- 7** After rolling upgrade phase 1 of the cluster, the following message displays:

Would you like to perform rolling upgrade phase 2 on the cluster?
[y,n,q] (y)

- If you choose *y*, it continues to run rolling upgrade phase 2 of the cluster by itself. You don't need to run phase 2:

After phase 2 upgrade, verify the cluster's status:

```
# hastatus -sum
```

- If you choose *n* or *q*, you need to use the following steps to finish rolling upgrade phase 2 of the cluster:

Performing a rolling upgrade on non-kernel packages for VCS, SFHA and SFCFSA : phase 2

In this phase installer installs all non-kernel packages on all the nodes in cluster and restarts VCS cluster.

To perform the rolling upgrade on non-kernel packages: phase 2

- 1 Start the installer for the rolling upgrade with the `-rollingupgrade_phase2` option. Specify all the nodes in the cluster:

```
# ./installmr -rollingupgrade_phase2 nodeA nodeB nodeC nodeD
```

You can also use the web-based installer to perform the upgrade. If you are using the web-based installer, select **Rolling Upgrade** from the task list, make sure the **Phase-2: Upgrade non Kernel packages** radio button is checked, and then click **Next**.

- 2 The installer checks system communications, patch versions, product versions, and completes prechecks. It upgrades non-kernel patches. It also verifies completion of phase 1.

If you are using the web-based installer, input one node of the cluster, the web-based installer detects the whole cluster to run the rolling upgrade phase 2. The web-based installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1.

- 3 Installer will start HA daemon (had) on all nodes, HA will be available once HA daemon is up.
- 4 Verify the cluster's status:

```
# hastatus -sum
```

- 5 If you want to upgrade CP server systems that use VCS or SFHA to 6.0.1, make sure that you upgraded all application clusters to version 6.0.1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

Performing a rolling upgrade on kernel packages for SF Oracle RAC: phase 1

Note that in the following instructions that a subcluster can represent one or more nodes in a full cluster, but is represented by nodeA, nodeB as subcluster1 and nodeC, nodeD as subcluster2.

To perform the rolling upgrade on kernel: phase 1

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`.

For example:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
  /etc/VRTSvcs/conf/config/main.cf.save
# cp /etc/VRTSvcs/conf/config/types.cf \
  /etc/VRTSvcs/conf/config/types.cf.save
# cp /etc/VRTSvcs/conf/config/OracleTypes.cf \
  /etc/VRTSvcs/conf/config/OracleTypes.cf.save
# cp /etc/VRTSvcs/conf/config/PrivNIC.cf \
  /var/tmp/PrivNIC.cf.save
# cp /etc/VRTSvcs/conf/config/MultiPrivNIC.cf \
  /var/tmp/MultiPrivNIC.cf.save
```

- 3
 - If the Oracle database is managed by VCS, set the `AutoStart` value to 0 to prevent the database service group from starting automatically when VCS starts. Failing to perform this step results in the database attempting to come online after the upgrade; the attempt fails due to the presence of old libraries on the system.

```
# haconf -makerw
# hagrps -modify oracle_group AutoStart 0
# haconf -dump -makero
```

- If the Oracle database is not managed by VCS, change the management policy for the database to manual:

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y MANUAL
```

For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy MANUAL
```

- 4 Stop the applications that use VxFS or VxVM disk groups on each node, whether local or CFS.

If the applications are under VCS control:

```
# hagrps -offline grp_name -sys node_name
```

If the applications are not under VCS control:

Use native application commands to stop the application.

- 5 Stop the Oracle RAC resources on each node.

- If the database instances are managed by VCS, take the corresponding VCS service groups offline. As superuser, enter:

```
# hagrps -offline oracle_group -sys nodeA
# hagrps -offline oracle_group -sys nodeB
```

- If the database instances are not managed by VCS, then run the following on one node:

For Oracle RAC 12 c:

```
$ srvctl stop instance -db db_name -node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl stop instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl stop instance -d db_name \
-i instance_name
```

- 6 Switch over all failover service groups to the other nodes in the cluster:

```
# hagrps -switch grp_name -to node_name
```

- 7 Take all the VCS service groups offline:

```
# hagrps -offline grp_name -sys node_name
```


- 8 Unmount all the VxFS file system which is not under VCS control.

```
# mount |grep vxfs

# fuser -m /mount_point

# umount /mount_point
```

Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu /mount_point
```

Note: Installer will automatically stop all the applications, database instances, filesystems and volumes which are under VCS control on nodes, while using the `rollingupgrade_phase1` option.

- 9 On the sub-cluster, start the installer for the rolling upgrade with the `-rollingupgrade_phase1` option.

```
# ./installmr -rollingupgrade_phase1 nodeA nodeB
```

You can also use the web-based installer to perform the upgrade. If you are using the web-based installer, start the web-based installer with the `./webinstaller start` command, select **Rolling Upgrade** from the task list, make sure the **Phase-1: Upgrade Kernel packages** radio button is checked, and then click **Next**.

- 10 Note that if the boot-disk is encapsulated, you do not need to perform an unencapsulation for upgrades.
- 11 The installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel patches.

If you are using the web-based installer, input one node of the cluster. The web-based installer detects the whole cluster, and then recommend some nodes (NodeA) as the subcluster to run the rolling upgrade phase 1. The web-based installer checks system communications, package versions, product versions, and completes prechecks. It then upgrades applicable kernel.
- 12 Relink the SF Oracle RAC libraries with Oracle:

Choose the option **Relink Oracle Database Binary** from the program menu.
- 13 If the boot disk is encapsulated, the installer strongly recommends a reboot of the nodes. Reboot the nodes as prompted by the installer.

14 For minimal down time, bring the oracle groups and database up manually on subcluster on which phase 1 is completed:

- If the database instances are managed by VCS, take the corresponding VCS service groups online. As superuser, enter:

```
# hagrps -online oracle_group -sys nodeA
# hagrps -online oracle_group -sys nodeB
```

- If the database instances are not managed by VCS, then run the following on one node:

For Oracle RAC 12c:

```
# srvctl start instance -db db_name \
-node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
# $ srvctl start instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
# $ srvctl start instance -d db_name \
-i instance_name
```

15 Manually mount the VxFS and CFS file systems that are not managed by VCS.

16 Bring the Oracle database service group online.

- If VCS manages the Oracle database:

```
# hagrps -online oracle_group -sys node_name
```

- If VCS does not manage the Oracle database:

For Oracle RAC 12c:

```
$ srvctl start instance -db db_name \
-node node_name
```

For Oracle RAC 11.2.0.2 and later versions:

```
$ srvctl start instance -d db_name \
-n node_name
```

For Oracle RAC 11.2.0.1 and earlier versions:

```
$ srvctl start instance -d db_name \  
-i instance_name
```

- 17 Start all applications that are not managed by VCS. Use native application commands to start the applications.
- 18 After rolling upgrade phase 1 is completed on nodeA and nodeB, the following message displays:

It is recommended to perform rolling upgrade phase 1 on the systems nodeC and nodeD in the next step.

Would you like to perform rolling upgrade phase 1 on the systems?
[y,n,q] (y)

- If you choose *y*, first complete step 4 to step 8 on the remaining subcluster. Then it continues to run rolling upgrade phase 1 on nodeC and nodeD by itself.
- If you choose *n* or *q*, go to step 15.

- 19 After rolling upgrade phase 1 of the cluster, the following message displays:

Would you like to perform rolling upgrade phase 2 on the cluster?
[y,n,q] (y)

- If you choose *y*, it continues to run rolling upgrade phase 2 of the cluster by itself. You don't need to run phase 2: See [“Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2”](#) on page 116. After rolling upgrade phase 2, complete step 18 to step 21 (except step 20) verify the cluster's status:

```
# hastatus -sum
```
- If you choose *n* or *q*, you need to complete step 18 to step 21 and run rolling upgrade phase 2: See [“Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2”](#) on page 116.

- 20 Before you proceed to phase 2, complete step 4 to 17 on the remaining subcluster.

- 21 Perform one of the following steps:

- If VCS manages the Oracle database, reset the AutoStart value to 1 to enable VCS to bring the database service group online automatically when VCS starts:

```
# haconf -makerw  
# hagrps -modify oracle_group AutoStart 1  
# haconf -dump -makero
```

- If the VCS does not manage the Oracle database, change the management policy for the database to automatic:
 For Oracle RAC 12c:

```
$ srvctl modify database -db db_name -policy AUTOMATIC
```

For Oracle RAC 11g:

```
$ srvctl modify database -d db_name -y AUTOMATIC
```

Performing a rolling upgrade on non-kernel packages for SF Oracle RAC: phase 2

In this phase, the installer installs all non-kernel RPMs on all the nodes in cluster and restarts VCS cluster.

To perform the rolling upgrade on non-kernel packages: phase 2

- 1 Start the installer for the rolling upgrade with the `-rollingupgrade_phase2` option. Specify all the nodes in the cluster:

```
./installmr -rollingupgrade_phase2 nodeA nodeB nodeC nodeD
```

You can also use the web-based installer to perform the upgrade. If you are using the web-based installer, select **Rolling Upgrade** from the task list, make sure the **Phase-2: Upgrade non Kernel packages** radio button is checked, and then click **Next**.

- 2 The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1.
 If you are using the web-based installer, input one node of the cluster, the web-based installer detects the whole cluster to run the rolling upgrade phase 2. The web-based installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1.
- 3 Installer will start HA daemon (had) on all nodes, HA will be available once HA daemon is up.
- 4 Verify the cluster's status:

```
# hastatus -sum
```

Upgrading the operating system

This section describes how to upgrade the operating system on a Veritas Storage Foundation and High Availability Solutions (SFHA) node where you plan to upgrade to 6.0.5 for Red Hat Enterprise Linux (RHEL) 5, RHEL 6, SUSE Linux Enterprise (SLES) 10, SLES 11, Oracle Linux (OL) 5, and OL 6.

Preparing for OS upgrades

- 1 Rename the `/etc/llttab` file to prevent LLT from starting automatically when the node starts:

```
# mv /etc/llttab /etc/llttab.save
```

- 2 Upgrade the operating system on all nodes in the cluster.
For instructions, see the operating system documentation.

Note: If required, restart the nodes with the `shutdown -r now` command.

- 3 After the system restarts, restore the `/etc/llttab` file to its original name:

```
# mv /etc/llttab.save /etc/llttab
```

To upgrade the operating system to a later version

- 1 Stop Storage Foundation.
- 2 Upgrade to the latest operating system.
- 3 Upgrade to 6.0.5.
- 4 Start Storage Foundation.

Upgrading SFHA with the Veritas Web-based installer

This section describes upgrading SFHA with the Veritas Web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

To perform a full upgrade for SFHA

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.

- 3 On the **Select a task and a product** page, select **Install 6.0.5** from the **Task** drop-down list, and click **Next**.
- 4 Stop all applications accessing the file system. Unmount all mounted filesystems before installation.
- 5 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Next**.
The installer detects the product that is installed on the specified system.
- 6 The installer stops the processes. Choose to restore and reuse the previous configuration on all systems. Click **Next** to start the processes.
- 7 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.
- 8 Click **Finish**. The installer prompts you for another task.
To perform a rolling upgrade with Symantec Web-based installer, refer to 6.0.1 Installation Guides for your products.

Verifying software versions

To list the Veritas RPMs installed on your system, enter the following command:

```
# rpm -qa | grep VRTS
```

The output version for 6.0.5 is 6.0.500.0, and the VRTSperl version is 5.14.2.20.

Removing Veritas Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [About removing Veritas Storage Foundation and High Availability Solutions 6.0.5](#)
- [Uninstalling Veritas Storage Foundation Cluster File System High Availability 6.0.5](#)
- [Uninstalling Veritas Storage Foundation for Oracle RAC](#)

About removing Veritas Storage Foundation and High Availability Solutions 6.0.5

Symantec recommends that you follow the steps in the following sections to remove all the installed Veritas software, and then perform a complete reinstallation of the previous release.

For extensive uninstallation and reinstallation procedures, refer to the appropriate product's Installation Guide.

Uninstalling Veritas Storage Foundation Cluster File System High Availability 6.0.5

This section provides procedures for uninstalling Veritas Storage Foundation Cluster File System High Availability (SFCFSHA). You must complete the preparatory tasks before you uninstall SFCFS.

Preparing to uninstall Veritas Storage Foundation Cluster File System High Availability

The following procedure prepares your system for uninstalling Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

To prepare to uninstall Veritas Storage Foundation Cluster File System High Availability

- 1 Log in as the root user on any node in the cluster.
- 2 Verify that the following directories are set in your `PATH` environment variable:

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

- 3 Back up the following configuration files:

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

- 4 Determine if each node's root disk is under VxVM control and proceed as follows.

- Check if each node's root disk is under VxVM control:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- If the encapsulated root disk is mirrored, use the `vxrootadm` command to split the mirror.

For example, the following command removes a root disk mirror from the current root disk group:

```
# vxrootadm -Yv rmmirror rootdg_01-s0
```

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

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


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

- 5 If you have created any Veritas Volume Replicator (VVR) replicated volume groups (RVGs) on your system, perform the following steps:

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.
- Use the `vxrvrg stop` command to stop each RVG individually:

```
# vxrvrg -g diskgroup stop rvg_name
```

- On the Primary node, use the `vxrlink status` command to verify that all RLINKs are up-to-date:

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

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

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

- 7 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint1
# umount /filesystem1
```

If file system is mounted in a cluster, then use `cfsumount` command.

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

```
# vxvol -g dg1 stopall
```

To verify that no volumes are open:

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

- 9 Stop VCS:

```
# hastop -all
```

Uninstalling Veritas Storage Foundation Cluster File System High Availability

The following procedure uninstalls Veritas Storage Foundation Cluster File System High Availability (SFCFSHA).

To uninstall Veritas Storage Foundation Cluster File System High Availability

- 1 Log in as the root user on any node in the cluster.
- 2 Navigate to the directory that contains the uninstallation program:

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

- 3 Run the uninstallation program while specifying each node in the cluster:

```
# ./uninstallsfcfsha605 node1 node2
```

- 4 Confirm the uninstallation:

```
Are you sure you want to uninstall SFCFSHA [y,n,q] (y)
```

The installer stops the Veritas Storage Foundation Cluster File System High Availability processes and uninstalls the packages.

After uninstalling the Veritas Storage Foundation Cluster File System High Availability, refer to the *Veritas Storage Foundation Cluster File System High Availability 6.0.1 Installation Guide* to reinstall the 6.0.1 software.

Uninstalling Veritas Storage Foundation for Oracle RAC

The following procedure uninstalls Veritas Storage Foundation for Oracle RAC (SFRAC).

Note: This procedure will remove the complete SFRAC stack from all nodes.

To uninstall Veritas Storage Foundation for Oracle RAC

- 1 On each node, take the Oracle resources in the VCS configuration file (`main.cf`) offline.

```
# hagrpl -offline oracle_group -sys node_name
```

If the database is not managed by VCS, stop the Oracle database as follows:

```
$ srvctl stop database -d db_name
```

- 2 If Oracle Clusterware is not under VCS Control, then enter the following command on each node of the cluster to stop Oracle Clusterware.
 - For 11gR2 and 12c:

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

- 3 Stop the applications that use CVM or CFS that are not under VCS control
 - Using native application commands, stop the applications that use CVM or CFS on all nodes.
 - Verify that no processes use the CFS mount point:

```
# fuser -c mount_point
```

- 4 Unmount CFS file systems that are not under VCS control
 - Determine the file systems that need to be unmounted by checking the output of mount command.

```
# mount -v | grep vxfs | grep cluster
```

- Unmount each file system that is not controlled by VCS on each node:

```
# umount mount_point
```

- 5 Stop VCS to take the service groups on all nodes offline
 On any one node execute following command to stop VCS:

```
# hastop -all
```

- 6 Stopping the applications that use VxVM or VxFS that are not under VCS control
 - Using native application commands, stop the applications that use VxVM or VxFS.
 - Verify that no processes use the VxFS mount point:

```
# fuser -c mount_point
```

- 7 Unmounting VxFS file systems that are not under VCS control.
 - Determine the file systems that need to be unmounted by checking the output of mount command.

```
# mount -v | grep vxfs
```

Unmount each file system that is not controlled by VCS on each node:

```
# umount mount_point
```

- 8 Remove SF for Oracle RAC.
 - On any one node, navigate to the directory that contains the `uninstallsfrac` program:

```
# cd /opt/VRTS/install
```
 - Start the `uninstallsfrac` program:

```
# ./uninstallsfrac605
```
- 9 After uninstalling the SFRAC, refer to the *Veritas Storage Foundation for Oracle RAC 6.0.1 Installation and Configuration Guide* document to reinstall the SFRAC 6.0.1 software.

Using Oracle RAC High Availability IP (HAIP) with Oracle RAC 11.2.0.2 and later versions

This appendix includes the following topics:

- [Using Oracle High Availability IP \(HAIP\)](#)
- [Configuring static IP addresses for HAIP using the SFHA script-based installer](#)
- [Configuring static IP addresses for HAIP manually](#)

Using Oracle High Availability IP (HAIP)

Starting with Oracle RAC 11.2.0.2, Oracle introduced the High Availability IP (HAIP) feature for supporting IP address failover. The purpose of HAIP is to perform load balancing across all active interconnect interfaces and fail over existing non-responsive interfaces to available interfaces. HAIP has the ability to activate a maximum of four private interconnect connections. These private network adapters can be configured during the installation of Oracle Grid Infrastructure or after the installation using the `oifcfg` utility.

Oracle Grid Infrastructure currently creates an alias IP (known as virtual private IP) on a private network interface using the 169.254.*.* subnet for HAIP. However, if that subnet range is already in use, Oracle Grid Infrastructure will not attempt to use it. Therefore, to enable HAIP using the Oracle Grid Infrastructure installer, configure all private interfaces with static private IP addresses before starting the

Grid installation. During the installation of Oracle Grid Infrastructure, select these interfaces as 'private' interfaces. It is recommended to use all the LLT links as cluster interconnects.

Note: If you are installing Oracle RAC 11.2.0.2 or later versions, make sure that you configure HAIP before the installation of Oracle RAC. If you are upgrading to Oracle RAC 11.2.0.2 and later, ensure that you perform the HAIP configuration after the upgrade.

You can configure HAIP in one of the following ways:

SF Oracle RAC script-based installer See [“Configuring static IP addresses for HAIP using the SFHA script-based installer”](#) on page 126.

Manual See [“Configuring static IP addresses for HAIP manually”](#) on page 129.

Configuring static IP addresses for HAIP using the SFHA script-based installer

The high availability IP address is managed by Oracle Grid Infrastructure. SF Oracle RAC supports the configuration of multiple network interfaces for HAIP. One static IP address is required for each network interface. All nodes must use a common network interface for the private interconnects in the cluster.

Review the following information before you configure HAIP:

- Do not use the public network interface for Oracle HAIP configuration.
- All LLT links are used for configuring HAIP.
- The same network interfaces must be used for private interconnect on each node of the cluster.
- The private IP addresses must be added to `/etc/hosts` on all nodes before installing Oracle Clusterware. If you choose the installer to add IP addresses in `/etc/hosts` then only non-existent IP addresses in the file on any nodes are added.

To configure static IP addresses for HAIP using the SFHA script-based installer

- 1** Refer to "Preparing to install Oracle RAC using the SF Oracle RAC installer or manually" in the 6.0.1 *Veritas Storage Foundation™ for Oracle® RAC Installation and Configuration Guide* for steps of launching the SF Oracle RAC installer.

The installmr script does not support the following steps.

- 2** From the SF Oracle RAC menu, enter **3** to select the option **Oracle Network Configuration**.

```
1) Create Oracle Users and Groups
2) Create Storage for OCR and Voting disk
3) Oracle Network Configuration
4) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: [1-4,b,q] (1) 3
```

- 3** Enter **1** to select the option **Configure private IP addresses (For HAIP Configuration) - For Oracle RAC 11.2.0.2 and later**.

```
1) Configure private IP addresses (For HAIP Configuration)
- For Oracle RAC 11.2.0.2 and later
2) Configure private IP addresses (PrivNIC Configuration)
- For Oracle RAC prior to 11.2.0.2
3) Configure private IP addresses (MultiPrivNIC Configuration)
- For Oracle RAC prior to 11.2.0.2
4) Exit SF Oracle RAC Configuration
b) Back to previous menu
Choose option: 1-3,b,q (1) 1
```

The installer discovers available LLT links.

4 Enter **y** to add the IP addresses to the `/etc/hosts` file.

```
Do you want the Installer to add IP addresses in /etc/hosts
file? [y,n,q] (y)
```

Enter **n** if you choose to add the IP addresses to the file manually and provide the following information:

```
Enter the IP for net1 on sys1: [b] 192.168.12.1
Enter the IP for net2 on sys1: [b] 192.168.13.1
Enter the IP for net1 on sys2: [b] 192.168.12.2
Enter the IP for net2 on sys2: [b] 192.168.13.2
```

Go to step 6.

5 Perform this step only if you enabled the installer to add the IP address to the `/etc/hosts` file in the previous step. Provide the private IP address and the private node name for the IP address that must be added to the file.

Note: IP addresses used for a NIC on all nodes of your cluster must be in the same subnet, which must be different from the subnets for the IP addresses on other NICs. Otherwise Oracle Clusterware/Grid Infrastructure cannot communicate properly across the nodes. If the private IP address entries are already present in the `/etc/hosts` file on one of nodes in the cluster, the installer does not update the file with the specified IP addresses on any of the nodes in the cluster. Only those IP addresses that are not already present in the file on any node are added.

```
Enter the IP for net1 on sys1: [b] 192.168.12.1
Enter Hostname alias for the above IP address: [b] sys1-haip1
Enter the private IP for net2 on sys1: [b] 192.168.13.1
Enter Hostname alias for the above IP address: [b] sys1-haip2
Enter the IP for net1 on sys2: [b] 192.168.12.2
Enter Hostname alias for the above IP address: [b] sys2-haip1
Enter the private IP for net2 on sys2: [b] 192.168.13.2
Enter Hostname alias for the above IP address: [b] sys2-haip2
```

6 Enter the netmask information for the private network:

```
Enter the Netmask for private network: [b] (255.255.255.0)
```

The SF Oracle RAC installer now displays the configuration information.

7 Enter **y** to review and confirm the configuration information.

If you choose **y** in step 4, the installer updates the `/etc/hosts` file.

- 8 If you choose **n** in step 4, proceed to update the `/etc/hosts` file manually.
- 9 If you upgraded to Oracle RAC 11.2.0.2 or later from previous versions, and have PrivNIC or MultiPrivNIC configured on your systems, perform the following steps to remove the PrivNIC or MultiPrivNIC resources:

```
# haconf -makerw
# hares -unlink cssd res_name
# hares -delete res_name
# haconf -dump -makero
```

Configuring static IP addresses for HAIP manually

You can configure a static IP addresses for HAIP manually using native operating system commands. See the operating system documentation for the appropriate commands.

For example, to configure a static IP address:

Edit the file `/etc/sysconfig/network-scripts/ifcfg-interface_name`

```
DEVICE=interface_name
ONBOOT=yes
BOOTPROTO=none
IPADDR=ip-address
NETMASK=netmask
```

If you upgraded to Oracle RAC 11.2.0.2 or later from previous versions, and have PrivNIC or MultiPrivNIC configured on your systems, perform the following steps to remove the PrivNIC or MultiPrivNIC resources:

```
# haconf -makerw
# hares -unlink cssd res_name
# hares -delete res_name
# haconf -dump -makero
```