

Storage Foundation 8.0 Configuration and Upgrade Guide - Linux

Last updated: 2021-12-28

Legal Notice

Copyright © 2021 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas 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. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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 Veritas Technologies LLC 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. VERITAS TECHNOLOGIES LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054
<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

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

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

infoscaledocs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1	Introduction and configuration of Storage Foundation	7
Chapter 1	Introducing Storage Foundation	8
	About Storage Foundation	8
	About Veritas Replicator Option	9
	About Veritas InfoScale Operations Manager	9
	About Veritas Services and Operations Readiness Tools (SORT)	10
Chapter 2	Configuring Storage Foundation	11
	Configuring Storage Foundation using the installer	11
	Configuring SF manually	12
	Configuring Veritas Volume Manager	12
	Configuring Veritas File System	12
	Configuring SFDB	14
Section 2	Upgrade of Storage Foundation	15
Chapter 3	Planning to upgrade Storage Foundation	16
	About the upgrade	16
	Supported upgrade paths	18
	Preparing to upgrade SF	22
	Getting ready for the upgrade	22
	Creating backups	24
	Determining if the root disk is encapsulated	25
	Pre-upgrade planning when VVR is configured	25
	Upgrading the array support	28
	Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches	28

Chapter 4	Upgrading Storage Foundation	31
	Upgrading Storage Foundation from previous versions to 8.0	31
	Upgrading Storage Foundation using the product installer	31
	Upgrading Volume Replicator	35
	Upgrading VVR without disrupting replication	35
	Upgrading SFDB	38
Chapter 5	Performing an automated SF upgrade using response files	39
	Upgrading SF using response files	39
	Response file variables to upgrade SF	40
	Sample response file for SF upgrade	43
Chapter 6	Performing post-upgrade tasks	44
	Optional configuration steps	44
	Re-joining the backup boot disk group into the current disk group	45
	Reverting to the backup boot disk group after an unsuccessful upgrade	45
	Recovering VVR if automatic upgrade fails	46
	Resetting DAS disk names to include host name in FSS environments	46
	Upgrading disk layout versions	47
	Upgrading VxVM disk group versions	48
	Updating variables	49
	Setting the default disk group	49
	Verifying the Storage Foundation upgrade	49
Section 3	Post configuration tasks	50
Chapter 7	Performing configuration tasks	51
	Switching on Quotas	51
	Enabling DMP support for native devices	51
	About configuring authentication for SFDB tools	52
	Configuring vxdbd for SFDB tools authentication	52

Section 4	Configuration and Upgrade reference	
	54
Appendix A	Installation scripts	55
	Installation script options	55
	About using the postcheck option	60
Appendix B	Configuring the secure shell or the remote shell for communications	63
	About configuring secure shell or remote shell communication modes before installing products	63
	Manually configuring passwordless ssh	64
	Setting up ssh and rsh connection using the installer -comsetup command	68
	Setting up ssh and rsh connection using the pwdutil.pl utility	69
	Restarting the ssh session	72
	Enabling rsh for Linux	72

Introduction and configuration of Storage Foundation

- [Chapter 1. Introducing Storage Foundation](#)
- [Chapter 2. Configuring Storage Foundation](#)

Introducing Storage Foundation

This chapter includes the following topics:

- [About Storage Foundation](#)
- [About Veritas InfoScale Operations Manager](#)
- [About Veritas Services and Operations Readiness Tools \(SORT\)](#)

About Storage Foundation

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

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

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

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

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

Note: The commands used for the Red Hat Enterprise Linux (RHEL) operating system in this document also apply to supported RHEL-compatible distributions.

About Veritas Replicator Option

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

File Replicator enables replication at the file level over IP networks. File Replicator leverages data duplication, provided by Veritas File System, to reduce the impact of replication on network resources.

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

Volume Replicator is available with Storage Foundation, Storage Foundation High Availability, Storage Foundation Cluster File System, Storage Foundation for Oracle RAC, and Storage Foundation for SybaseCE.

Before installing this option, read the Release Notes for the product.

To install the option, follow the instructions in the Installation Guide for the product.

About Veritas InfoScale Operations Manager

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

Veritas recommends using Veritas InfoScale Operations Manager to manage Storage Foundation and Cluster Server environments.

You can download Veritas InfoScale Operations Manager from:

https://www.veritas.com/content/support/en_US/downloads

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

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

<https://www.veritas.com/form/trialware/vcs-utilities>

Storage Foundation Management Server is deprecated.

About Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a Web site that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you manage your datacenter more efficiently and get the most out of your Veritas products.

SORT can help you do the following:

- | | |
|---|---|
| Prepare for your next installation or upgrade | <ul style="list-style-type: none">■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture.■ Analyze systems to determine if they are ready to install or upgrade Veritas products.■ Download the latest patches, documentation, and high availability agents from a central repository.■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems. |
| Manage risks | <ul style="list-style-type: none">■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDIs), and high availability agents from a central repository.■ Identify and mitigate system and environmental risks.■ Display descriptions and solutions for hundreds of Veritas error codes. |
| Improve efficiency | <ul style="list-style-type: none">■ Find and download patches based on product version and platform.■ List installed Veritas products and license keys.■ Tune and optimize your environment. |

Note: Certain features of SORT are not available for all products. Access to SORT is available at no extra cost.

To access SORT, go to:

<https://sort.veritas.com>

Configuring Storage Foundation

This chapter includes the following topics:

- [Configuring Storage Foundation using the installer](#)
- [Configuring SF manually](#)
- [Configuring SFDB](#)

Configuring Storage Foundation using the installer

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

To configure Storage Foundation

- 1 Go to the `/opt/VRTS/install/` installation directory.
- 2 Run the installer command with the configure option.

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

Or run the `/opt/VRTS/install/installer` command, then select the configure option:

```
Task Menu:
```

```
C) Configure a Product Component  
U) Uninstall a Product  
L) License a Product  
S) Start a Product  
D) View Product Descriptions  
X) Stop a Product  
O) Perform a Post-Installation Check  
?) Help
```

```
Enter a Task: [C,U,L,S,D,X,O,?] C
```

Configuring SF manually

You can manually configure different products within SF.

Configuring Veritas Volume Manager

Use the following procedures to configure Veritas Volume Manager. If you have installed and configured VxVM using the product installer, you do not need to complete the procedures in this section.

For information on setting up VxVM disk groups and volumes after installation, see "Configuring Veritas Volume Manager" in the *Storage Foundation Administrator's Guide*.

Configuring Veritas File System

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

```
/etc/fstab
```

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

See the *Storage Foundation Administrator's Guide*.

Loading and unloading the file system module

The `vxfs` file system module automatically loads on the first reference to a VxFS file system; this occurs when a user tries to mount a VxFS file system.

In some instances, you may find it efficient to load the file system module manually. For example, some larger class systems can have many dual interface I/O cards with multiple disk chains attached. The device interrogation process when such a system is rebooted can be very time consuming, so to avoid doing a reboot, use the `modprobe` command to load the `vxfs` module:

```
# modprobe vxfs ; modprobe vxportal ; modprobe fdd
```

Do not use the `insmod` command to load the `vxfs` module as `insmod` does not examine the module configuration file `/etc/modprobe.conf`.

To determine if the modules successfully loaded, use the `lsmod` command as shown here:

```
# lsmod | grep vxportal
vxportal                2952                0
vxfs                    3427960            0    fdd vxportal
# lsmod | grep fdd
fdd                     67212              0    (unused)
vxfs                    3427960            0    [fdd vxportal]
# lsmod | grep vxfs
vxfs                    3427960            0    [fdd vxportal]
```

The first field in the output is the module name. You can unload the modules by entering:

```
# rmmod fdd
# rmmod vxportal
# rmmod vxfs
```

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

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

Configuring SFDB

By default, SFDB tools are disabled that is the vxdbd daemon is not configured. You can check whether SFDB tools are enabled or disabled using the `/opt/VRTS/bin/sfae_config status` command.

To enable SFDB tools

- 1 Log in as root.
- 2 On SLES 15 systems, install the `insserv-compat` package manually. This package is required to enable the `vxdbd` daemon.
- 3 Run the following command to configure and start the `vxdbd` daemon. After you perform this step, entries are made in the system startup so that the daemon starts on a system restart.

```
#/opt/VRTS/bin/sfae_config enable
```

To disable SFDB tools

- 1 Log in as root.
- 2 Run the following command:

```
#/opt/VRTS/bin/sfae_config disable
```

Upgrade of Storage Foundation

- [Chapter 3. Planning to upgrade Storage Foundation](#)
- [Chapter 4. Upgrading Storage Foundation](#)
- [Chapter 5. Performing an automated SF upgrade using response files](#)
- [Chapter 6. Performing post-upgrade tasks](#)

Planning to upgrade Storage Foundation

This chapter includes the following topics:

- [About the upgrade](#)
- [Supported upgrade paths](#)
- [Preparing to upgrade SF](#)
- [Using Install Bundles to simultaneously install or upgrade full releases \(base, maintenance, rolling patch\), and individual patches](#)

About the upgrade

This release supports upgrades from 6.2.1 and later versions. If your existing installation is from a pre-6.2.1 version, you must first upgrade to version 6.2.1, then follow the procedures mentioned in this document to upgrade the product.

The installer supports the following types of upgrade:

- Full upgrade
- Automated upgrade using response files

[Table 3-1](#) describes the product mapping after an upgrade.

Table 3-1 Veritas InfoScale product mapping after upgrade

Product (6.2.x and earlier)	Product (7.0 and later)	Component (7.0 and later)
SF Basic	No upgrade supported	Not applicable

Table 3-1 Veritas InfoScale product mapping after upgrade (*continued*)

Product (6.2.x and earlier)	Product (7.0 and later)	Component (7.0 and later)
SF	Veritas InfoScale Storage	SF
SF	Veritas InfoScale Foundation	SF
SF	Veritas InfoScale Enterprise	SF

Note: From 7.0 onwards, the existing Veritas InfoScale product upgrades to the higher version of the same product. For example, Veritas InfoScale Enterprise 7.4.1 gets upgraded to Veritas InfoScale Enterprise 7.4.2.

During the upgrade, the installation program performs the following tasks:

1. Stops the product before starting the upgrade
2. Upgrades the installed packages and installs additional packages

Slf license key files are required while upgrading to version 7.4 and later. The text-based license keys that are used in previous product versions are not supported when upgrading to version 7.4 and later. If you plan to upgrade any of the InfoScale products from a version earlier than 7.4, first contact Customer Care for your region to procure an applicable slf license key file. Refer to the following link for contact information of the Customer Care center for your region: https://www.veritas.com/content/support/en_US/contact-us.html.

If your current installation uses a permanent license key, you will be prompted to update the license to 8.0. Ensure that the license key file is downloaded on the local host, where you want to upgrade the product. The license key file must not be saved in the root directory (/) or the default license directory on the local host (/etc/vx/licenses/lic). You can save the license key file inside any other directory on the local host.

If you choose not to update your license, you will be registered with a keyless license. Within 60 days of choosing this option, you must install a valid license key file corresponding to the entitled license level.
3. You must configure the Veritas Telemetry Collector while upgrading, if you have do not already have it configured. For more information, refer to the *About telemetry data collection in InfoScale* section in the *Veritas Installation guide*.
4. Restores the existing configuration.

For example, if your setup contains an SF installation, the installer upgrades and restores the configuration to SF. If your setup included multiple components, the installer upgrades and restores the configuration of the components.

5. Starts the configured components.

Supported upgrade paths

If you are on an unsupported operating system version, ensure that you first upgrade to a supported version of the operating system. Also, upgrades between major operating system versions are not supported, for example, from RHEL 7 to RHEL 8. If you plan to move between major operating system versions, you need to reinstall the product. For supported operating system versions, see the *Veritas InfoScale Release Notes*.

[Table 3-2](#) lists the supported upgrade paths for upgrades on RHEL and Oracle Linux.

Table 3-2 Supported upgrade paths on RHEL and Oracle Linux

From product version	From OS version	To OS version	To product version	To Component
6.2.1	RHEL 6 Update 4, 5, 6, 7, 8, 9, 10 RHEL 7 GA, Update 1, 2, 3, 4, 5, 6, 7, 8, 9 Oracle Linux 6 Update 4, 5, 6, 7 Oracle Linux 7 GA, Update 2, 3	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 9 Oracle Linux 8 Update 2, 4	Veritas InfoScale Storage 8.0	SF
7.1	RHEL 6 update 4, 5, 6, 7, 8, 9 RHEL 7 Update 1, 2, 3, 4 Oracle Linux 6 Update 6, 7, 8 Oracle Linux 7 Update 1, 2	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 9 Oracle Linux 8 Update 2, 4	Veritas InfoScale Storage 8.0	SF

Table 3-2 Supported upgrade paths on RHEL and Oracle Linux (*continued*)

From product version	From OS version	To OS version	To product version	To Component
7.2	RHEL 6 update 6, 7, 8, 9 RHEL 7 Update 1, 2, 3, 4 Oracle Linux 6 Update 4, 5, 6, 7, 8, 9 Oracle Linux 7 Update 1, 2, 3, 4	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 9 Oracle Linux 8 Update 2, 4	Veritas InfoScale Storage 8.0	SF
7.3	RHEL 6 update 7, 8, 9 RHEL 7 Update 1,2, 3, 4 Oracle Linux 6 Update 7, 8, 9 Oracle Linux 7 Update 1, 2, 3	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 9 Oracle Linux 8 Update 2, 4	Veritas InfoScale Storage 8.0	SF
7.3.1	RHEL 6 update 7, 8, 9 RHEL 7 Update 3, 4, 5 Oracle Linux 6 Update 7, 8, 9 Oracle Linux 7 Update 3, 4	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 9 Oracle Linux 8 Update 2, 4	Veritas InfoScale Storage 8.0	SF

Table 3-2 Supported upgrade paths on RHEL and Oracle Linux (*continued*)

From product version	From OS version	To OS version	To product version	To Component
7.4	RHEL 6 Update 7, 8, 9, 10 RHEL 7 Update 3, 4, 5 Oracle Linux 6 Update 7, 8, 9 Oracle Linux 7 Update 3, 4	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 9 Oracle Linux 8 Update 2, 4	Veritas InfoScale Storage 8.0	SF
7.4.1	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 4, 5, 6, 7, 8, 9 Oracle Linux 8 Update 1, 2	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 9 Oracle Linux 8 Update 2, 4	Veritas InfoScale Storage 8.0	SF
7.4.2	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 7 Oracle Linux 8 Update 1	RHEL 7 Update 9 RHEL 8 Update 2, 4 Oracle Linux 7 Update 9 Oracle Linux 8 Update 2, 4	Veritas InfoScale Storage 8.0	SF

Table 3-3 lists the supported upgrade paths for upgrades on SLES.

Table 3-3 Supported upgrade paths on SLES

From product version	From OS version	To OS version	To product version	To component
6.2.1	SLES11 SP2, SP3, SP4 SLES 12 GA, SP1, SP2, SP3, SP4	SLES 12 SP5 SLES 15 SP2	Veritas InfoScale Storage 8.0	SF
7.1, 7.2	SLES 11 SP3, SP4 SLES 12 GA, SP1, SP2	SLES 12 SP5 SLES 15 SP2	Veritas InfoScale Storage 8.0	SF
7.3	SLES 11 SP3, SP4 SLES 12 SP1, SP2	SLES 12 SP5 SLES 15 SP2	Veritas InfoScale Storage 8.0	SF
7.3.1	SLES 11 SP3, SP4 SLES 12 SP2, SP4, SP5	SLES 12 SP5 SLES 15 SP2	Veritas InfoScale Storage 8.0	SF
7.4	SLES 11 SP3, SP4 SLES 12 SP2, SP3	SLES 12 SP5 SLES 15 SP2	Veritas InfoScale Storage 8.0	SF
7.4.1	SLES 11 SP3, SP4 SLES 12 SP2, SP3, SP4, SP5 SLES 15 GA, SP1, SP2	SLES 12 SP5 SLES 15 SP2	Veritas InfoScale Storage 8.0	SF
7.4.2	SLES 12 SP 4, SP5 SLES 15 SP1, SP2	SLES 12 SP5 SLES 15 SP2	Veritas InfoScale Storage 8.0	SF

Preparing to upgrade SF

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

Getting ready for the upgrade

Complete the following tasks before you perform the upgrade:

- Review the *Veritas InfoScale 8.0 Release Notes* for any late-breaking information on upgrading your system.
- Review the Veritas Technical Support website for additional information: https://www.veritas.com/support/en_US.html
- You can configure the Veritas Telemetry Collector while upgrading, if you have do not already have it configured. For more information, refer to the *About telemetry data collection in InfoScale* section in the *Veritas Installation guide*.
- Make sure that the administrator who performs the upgrade has root access and a good knowledge of the operating system's administration.
- Make sure that all users are logged off and that all major user applications are properly shut down.
- Make sure that you have created a valid backup.
See “[Creating backups](#)” on page 24.
- Ensure that you have enough file system space to upgrade. Identify where you want to copy the RPMs, for example `/packages/Veritas` when the root file system has enough space or `/var/tmp/packages` if the `/var` file system has enough space.
Do not put the files under `/tmp`, which is erased during a system restart.
Do not put the files on a file system that is inaccessible before running the upgrade script.
You can use a Veritas-supplied disc for the upgrade as long as modifications to the upgrade script are not required.
If `/usr/local` was originally created as a slice, modifications are required.
- Comment out any application commands or processes that are known to hang if their file systems are not present in the startup scripts.
In case of RHEL 7 and SLES 12 systems, some startup scripts are located at `/etc/vx/`, and the startup scripts of the following services are located at:

Service name	Startup script location and file name
amf.service	<code>/opt/VRTSamf/bin/amf</code>

Service name	Startup script location and file name
<code>gab.service</code>	<code>/opt/VRTSgab/gab</code>
<code>llt.service</code>	<code>/opt/VRTSllt/llt</code>
<code>vcs.service</code>	<code>/opt/VRTSvcs/bin/vcs</code>
<code>vcsmm.service</code>	<code>/opt/VRTSvcs/rac/bin/vcsmm</code>
<code>vxfen.service</code>	<code>/opt/VRTSvcs/vxfen/bin/vxfen</code>

The remaining startup scripts for RHEL 7 and SLES 12 are located at `/etc/init.d/`, like all the other startup scripts for the other supported RHEL distributions.

- Make sure that the current operating system supports version 8.0 of the product. If the operating system does not support it, plan for a staged upgrade.

Note: Before you upgrade RHEL 7.7 OS on a virtual machine, you need to first upgrade Veritas InfoScale 8.0. Later upgrade RHEL 7.7 OS, else the virtual machine may go in an unstable state.

Use `-ignorechecks` CPI option on RHEL 7.0 to RHEL 7.6 version to successfully upgrade Veritas InfoScale product.

- Schedule sufficient outage time and downtime for the upgrade and any applications that use the Veritas InfoScale products. Depending on the configuration, the outage can take several hours.
- Any swap partitions not in `rootdg` must be commented out of `/etc/fstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/fstab` and not mounted during the upgrade. The active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.
- Make sure that the file systems are clean before upgrading.
- Upgrade arrays (if required).
See [“Upgrading the array support”](#) on page 28.
- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. Removing the disk in this manner offers a fallback point.
- Make sure that DMP support for native stack is disabled (`dmp_native_support=off`). If DMP support for native stack is enabled

(`dmp_native_support=on`), the installer may detect it and ask you to restart the system.

- If you want to upgrade the application clusters that use CP server based fencing to version 6.1 and later, make sure that you first upgrade VCS or SFHA on the CP server systems to version 6.1 and later. And then, from 7.0.1 onwards, CP server supports only HTTPS based communication with its clients and IPM-based communication is no longer supported. CP server needs to be reconfigured if you upgrade the CP server with IPM-based CP server configured. For instructions to upgrade VCS or SFHA on the CP server systems, refer to the relevant Configuration and Upgrade Guides.

Creating backups

Save relevant system information before the upgrade.

To create backups

- 1 Log in as superuser.
- 2 Before the upgrade, ensure that you have made backups of all data that you want to preserve.
- 3 Back up information in files such as `/boot/grub/menu.lst`, `/etc/grub.conf` or `/etc/lilo.conf`, and `/etc/fstab`.
- 4 Installer verifies that recent backups of configuration files in VxVM private region have been saved in `/etc/vx/cbr/bk`.

If not, a warning message is displayed.

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

- 5 Copy the `fstab` file to `fstab.orig`:


```
# cp /etc/fstab /etc/fstab.orig
```
- 6 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 7 If you install Veritas InfoScale Enterprise 8.0 software, follow the guidelines that are given in the *Cluster Server Configuration and Upgrade Guide* for information on preserving your VCS configuration across the installation procedure.

- 8 Back up the external `quotas` and `quotas.grp` files.

If you are upgrading from 6.2.1, you must also back up the `quotas.grp.64` and `quotas.64` files.

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

Determining if the root disk is encapsulated

Note: Root Disk Encapsulation (RDE) on Linux Distribution is not supported from 7.3.1 release onwards.

Check if the system's root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootdg/rootvol` is listed as being mounted as the root (`/`) file system.

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

Pre-upgrade planning when VVR is configured

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, Veritas 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 *Veritas InfoScale™ 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 *Veritas InfoScale™ Replication Administrator's Guide* for more information.

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

Considerations for upgrading SF to 7.4 or later on systems with an ongoing or a paused replication

Typically, you can upgrade SF in a setup where VVR is configured. However, InfoScale does not support upgrade from version 7.3.1 or earlier to version 7.4 or later with an ongoing or a paused replication. To upgrade InfoScale from these earlier versions to 7.4 or later, perform the following steps:

1. Stop replication to the Secondary using the `vradmin stoprep` command for all RVGs.
2. Upgrade InfoScale to version 7.4 or later at the primary and the secondary sites.
3. Upgrade the disk group version at the primary and the secondary sites.
4. Start replication using the `vradmin -a startrep` command.

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

For information regarding VVR support for replicating across Storage Foundation versions, refer to the *Veritas InfoScale Release Notes*.

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

8.0, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 8.0, 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 8.0 (DG version <= 140)	VVR 8.0 (DG version >= 310)	VVR calculates checksum TCP connections?
Primary	Secondary	Yes
Secondary	Primary	Yes
Primary and Secondary		Yes
	Primary and Secondary	No

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

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

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

Planning and upgrading VVR to use IPv6 as connection protocol

SF supports using IPv6 as the connection protocol.

This release supports the following configurations for VVR:

- VVR continues to support replication between IPv4-only nodes with IPv4 as the internet protocol
- VVR supports replication between IPv4-only nodes and IPv4/IPv6 dual-stack nodes with IPv4 as the internet protocol
- VVR supports replication between IPv6-only nodes and IPv4/IPv6 dual-stack nodes with IPv6 as the internet protocol
- VVR supports replication between IPv6 only nodes
- VVR supports replication to one or more IPv6 only nodes and one or more IPv4 only nodes from a IPv4/IPv6 dual-stack node

Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches

- VVR supports replication of a shared disk group only when all the nodes in the cluster that share the disk group are at IPv4 or IPv6

Upgrading the array support

The Veritas InfoScale 8.0 release includes all array support in a single RPM, `VRTSaslapm`. The array support RPM includes the array support previously included in the `VRTSvxvm` RPM. The array support RPM also includes support previously packaged as external Array Support Libraries (ASLs) and array policy modules (APMs).

See the 8.0 Hardware Compatibility List for information about supported arrays.

When you upgrade Storage Foundation products with the product installer, the installer automatically upgrades the array support. If you upgrade Storage Foundation products with manual steps, you should remove any external ASLs or APMs that were installed previously on your system. Installing the `VRTSvxvm` RPM exits with an error if external ASLs or APMs are detected.

After you have installed Veritas InfoScale 8.0, Veritas provides support for new disk arrays through updates to the `VRTSaslapm` RPM.

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

Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches

Beginning with version 6.2.1, you can easily install or upgrade your systems directly to a base, maintenance, patch level or a combination of multiple patches and packages together in one step using Install Bundles. With Install Bundles, the installer has the ability to merge so that customers can install or upgrade directly to maintenance or patch levels in one execution. The various scripts, RPMs, and patch components are merged, and multiple releases are installed together as if they are one combined release. You do not have to perform two or more install actions to install or upgrade systems to maintenance levels or patch levels.

Releases are divided into the following categories:

Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches

Table 3-5 Release Levels

Level	Content	Form factor	Applies to	Release types	Download location
Base	Features	RPMs	All products	Major, minor, Service Pack (SP), Platform Release (PR)	FileConnect
Maintenance	Fixes, new features	RPMs	All products	Maintenance Release (MR), Rolling Patch (RP)	Veritas Services and Operations Readiness Tools (SORT)
Patch	Fixes	RPMs	Single product	P-Patch, Private Patch, Public patch	SORT, Support site

When you install or upgrade using Install Bundles:

- InfoScale products are discovered and assigned as a single version to the maintenance level. Each system can also have one or more patches applied.
- Base releases are accessible from FileConnect that requires customer serial numbers. Maintenance and patch releases can be automatically downloaded from SORT.
- Patches can be installed using automated installers.
- Patches can now be detected to prevent upgrade conflict. Patch releases are not offered as a combined release. They are only available from Veritas Technical Support on a need basis.

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

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

In the example below:

- 8.0 is the base version

Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches

- 8.0.1 is the maintenance version
- 8.0.1.1000 is the patch version for 8.0.1
- 8.0.0.1000 is the patch version for 8.0

1. Base + maintenance:

This integration method can be used when you install or upgrade from a lower version to 8.0.1.

Enter the following command:

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

2. Base + patch:

This integration method can be used when you install or upgrade from a lower version to 8.0.0.100.

Enter the following command:

```
# installer -patch_path <path_to_patch>
```

3. Maintenance + patch:

This integration method can be used when you upgrade from version 8.0 to 8.0.1.100.

Enter the following command:

```
# installmr -patch_path <path_to_patch>
```

4. Base + maintenance + patch:

This integration method can be used when you install or upgrade from a lower version to 8.0.1.100.

Enter the following command:

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

Note: You can add a maximum of five patches using `-patch_path <path_to_patch> -patch2_path <path_to_patch> ... -patch5_path <path_to_patch>`

Upgrading Storage Foundation

This chapter includes the following topics:

- [Upgrading Storage Foundation from previous versions to 8.0](#)
- [Upgrading Volume Replicator](#)
- [Upgrading SFDB](#)

Upgrading Storage Foundation from previous versions to 8.0

Note: Root Disk Encapsulation (RDE) is not supported on Linux from 8.0 onwards.

If you are running an earlier release of Storage Foundation, you can upgrade to the latest version using the procedures described in this chapter.

See [“Upgrading Storage Foundation using the product installer”](#) on page 31.

If you need to upgrade your kernel with Storage Foundation 8.0 already installed, use the kernel upgrade procedure.

For information on upgrading the kernel, refer to the *Storage Foundation Administrator's Guide*.

Upgrading Storage Foundation using the product installer

Note: Root Disk Encapsulation (RDE) is not supported on Linux from 8.0 onwards.

For details on installing and upgrading Veritas InfoScale using the installer with the `-yum` option, refer to the *Veritas InfoScale Installation guide*.

Use this procedure to upgrade Storage Foundation (SF).

To upgrade SF from previous versions to 8.0

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

```
# df -k | 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 filesystem  
# mount -t vxfs filesystem mountpoint  
# umount mountpoint
```

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

There may be a pending large 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** If a cache area is online, you must take the cache area offline before you upgrade the VxVM RPM. Use the following command to take the cache area offline:

```
# sfcache offline cachename
```

- 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 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
```

- 8** 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.
- 9** Perform any necessary preinstallation checks.
- 10** To invoke the installer, run the `installer` command on the disc as shown in this example:

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

- 11** Enter `G` to upgrade and press Return.

- 12** 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 RHEL6.

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.

- 13** The installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.
- 14** The installer discovers if any of the systems that you are upgrading have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the system's boot disk group before the upgrade proceeds. If you want to split the boot disk group to create a backup, answer **y**.
- 15** The installer then prompts you to name the backup boot disk group. Enter the name for it or press **Enter** to accept the default.
- 16** You are prompted to start the split operation. Press **y** to continue.

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

- 17** Stop the product's processes.

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

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

- 18** The installer stops, uninstalls, reinstalls, and starts specified RPMs.
- 19** If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node that you recorded in step 8.
- 20** Restart all the volumes by entering the following command for each disk group:

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

21 Remount all VxFS file systems and Storage Checkpoints on all nodes:

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

22 You can perform the following optional configuration steps:

- If you want to use features of Storage Foundation 8.0 for which you do not currently have an appropriate license installed, obtain the license and run the `vxlicinst` command to add it to your system.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
See [“Upgrading VxVM disk group versions”](#) on page 48.

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

See [“Re-joining the backup boot disk group into the current disk group”](#) on page 45.

See [“Reverting to the backup boot disk group after an unsuccessful upgrade”](#) on page 45.

Upgrading Volume Replicator

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

You have the option to upgrade without disrupting replication.

Upgrading VVR without disrupting replication

This section describes the upgrade procedure from an earlier version of VVR to the current version of VVR when replication is in progress, assuming that you do not need to upgrade all the hosts in the RDS simultaneously.

You may also need to set up replication between versions.

See [“Planning an upgrade from the previous VVR version”](#) on page 26.

When both the Primary and the Secondary have the previous version of VVR installed, the upgrade should be performed on the secondary site first and then on the primary site.

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

Upgrading VVR sites for InfoScale 7.3.1 or earlier

Use the product installer to first upgrade VVR on the Secondaries and then on the Primary.

To upgrade a Secondary

- 1 Stop the replication to a Secondary by initiating `stoprep` on the Primary.

```
# /usr/sbin/vradmin -g <disk_group_name> stoprep <RVG_name>  
<secondary_hostname>
```

- 2 Verify that the replication has stopped.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

- 3 Upgrade VVR from any version between 6.2.1 and 7.3.1 to VVR 8.0 on the Secondary.

- 4 Start the replication to the Secondary host by initiating `startrep` on the Primary.

```
# /usr/sbin/vradmin -g <disk_group_name> startrep <RVG_name>  
<secondary_hostname>
```

- 5 Verify that the replication has started.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

To upgrade the Primary

- 1 Verify that the replication status is consistent and up-to-date.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

- 2 Take the applications and the mount points down.

- 3 Stop the replication to a Secondary by initiating `stoprep` on the Primary.

```
# /usr/sbin/vradmin -g <disk_group_name> stoprep <RVG_name>  
<secondary_hostname>
```

- 4 Verify that the replication has stopped.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

- 5 Upgrade VVR from any version between 6.2.1 and 7.3.1 to VVR 8.0 on the Primary.

- 6 Start the replication to the Secondary host by initiating `startrep` on the Primary.

```
# /usr/sbin/vradmin -g <disk_group_name> startrep <RVG_name>  
<secondary_hostname>
```

- 7 Verify that the replication has started.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

- 8 Mount all the file systems and start all the applications on the Primary.

Upgrading VVR sites with InfoScale 7.4 or later

Use the product installer to first upgrade VVR on the Secondaries and then on the Primary.

To upgrade a Secondary

- 1 Pause the replication to a Secondary by initiating `pauserep` on the Primary.

```
# /usr/sbin/vradmin -g <disk_group_name> pauserep <RVG_name>  
<secondary_hostname>
```

- 2 Verify that the replication has paused.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

- 3 Upgrade from VVR 7.4 or later to VVR 8.0 on the Secondary.

- 4 Resume the replication to the Secondary by initiating `resumerep` on the Primary.

```
# /usr/sbin/vradmin -g <disk_group_name> resumerep <RVG_name>  
<secondary_hostname>
```

- 5 Verify that the replication has resumed.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

To upgrade the Primary

- 1 Verify that the replication status is consistent and up-to-date.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

- 2 Take the applications and the mount points down.

- 3 Pause the replication to the Secondary by initiating `pauserep` on the Primary.

```
# /usr/sbin/vradmin -g <disk_group_name> pauserep <RVG_name>  
<secondary_hostname>
```

- 4 Verify that the replication has paused.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

- 5 Upgrade from VVR 7.4 or later to VVR 8.0 on the Primary.

- 6 Resume the replication to the Secondary by initiating `resumerep` on the Primary.

```
# /usr/sbin/vradmin -g <disk_group_name> resumerep <RVG_name>  
<secondary_hostname>
```

- 7 Verify that the replication has resumed.

```
# /usr/sbin/vradmin -g <disk_group_name> -l repstatus <RVG_name>
```

- 8 Mount all the file systems and start all the applications on the Primary.

Post-upgrade tasks for VVR sites

To upgrade disk group and disk layout versions on replication hosts

- 1 Upgrade the disk group version on all the Secondaries for all the disk groups.

```
# /usr/sbin/vxdg upgrade <disk_group_name>
```

- 2 Upgrade the disk group version on the Primary for all the disk groups.

```
# /usr/sbin/vxdg upgrade <disk_group_name>
```

- 3 Upgrade the disk layout version (DLV) on the Primary for all the VxFS file systems.

```
# /opt/VRTS/bin/vxupgrade -n 17 <vxfs_mount_point_name>
```

```
# /opt/VRTS/bin/fstyp -v <disk_path_for_mount_point_volume>
```

The DLV upgrade changes are automatically replicated to the secondaries.

Upgrading SFDB

While upgrading to 8.0, the SFDB tools are enabled by default, which implies that the `vxdbd` daemon is configured. You can enable the SFDB tools, if they are disabled.

To enable SFDB tools

- 1 Log in as root.
- 2 Run the following command to configure and start the `vxdbd` daemon.

```
# /opt/VRTS/bin/sfae_config enable
```

Note: If any SFDB installation with authentication setup is upgraded to 8.0, the commands fail with an error. To resolve the issue, setup the SFDB authentication again. For more information, see the *Veritas InfoScale™ Storage and Availability Management for Oracle Databases* or *Veritas InfoScale™ Storage and Availability Management for DB2 Databases*.

Performing an automated SF upgrade using response files

This chapter includes the following topics:

- [Upgrading SF using response files](#)
- [Response file variables to upgrade SF](#)
- [Sample response file for SF upgrade](#)

Upgrading SF using response files

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

To perform automated SF upgrade

- 1 Make sure the systems where you want to upgrade SF meet the upgrade requirements.
- 2 Make sure the pre-upgrade tasks are completed.
- 3 Copy the response file to the system where you want to upgrade SF.
- 4 Edit the values of the response file variables as necessary.

- 5 Mount the product disc and navigate to the folder that contains the installation program.
- 6 Start the upgrade from the system to which you copied the response file. For example:

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

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

Response file variables to upgrade SF

Table 5-1 lists the response file variables that you can define to configure SF.

Table 5-1 Response file variables for upgrading SF

Variable	Description
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{upgrade}	Upgrades all RPMs installed. List or scalar: list Optional or required: required
CFG{keys}{keyless} CFG{keys}{licensefile}	CFG{keys}{keyless} gives a list of keyless keys to be registered on the system. CFG{keys}{licensefile} gives the absolute file path to the permanent license key to be registered on the system. List or scalar: list Optional or required: required

Table 5-1 Response file variables for upgrading SF (*continued*)

Variable	Description
CFG{opt}{keyfile}	<p>Defines the location of an ssh keyfile that is used to communicate with all remote systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{tmppath}	<p>Defines the location where a working directory is created to store temporary files and the RPMs that are needed during the install. The default location is /opt/VRTStmp.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{logpath}	<p>Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{mirrordgname}{system}	<p>If the root dg is encapsulated and you select split mirror is selected:</p> <p>Splits the target disk group name for a system.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{splitmirror}{system}	<p>If the root dg is encapsulated and you select split mirror is selected:</p> <p>Indicates the system where you want a split mirror backup disk group created.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>

Table 5-1 Response file variables for upgrading SF *(continued)*

Variable	Description
CFG{opt}{disable_dmp_native_support}	<p>If it is set to 1, Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools is disabled after upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases RPM upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{patch_path}	<p>Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{patch2_path}	<p>Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{patch3_path}	<p>Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{patch4_path}	<p>Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>

Table 5-1 Response file variables for upgrading SF (*continued*)

Variable	Description
CFG{opt}{patch5_path}	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional

Sample response file for SF upgrade

The following example shows a response file for upgrading Storage Foundation with keyless license key.

```
our %CFG;
$CFG{accepteula}=1;
$CFG{keys}{keyless}=[ qw(STORAGE) ];
$CFG{prod}="STORAGE80";
$CFG{opt}{upgrade}=1;
$CFG{systems}=[ qw(sys1) ];
1;
```

The following example shows a response file for upgrading Storage Foundation with permanent license key.

```
our %CFG;
$CFG{accepteula}=1;
$CFG{keys}{licensefile}=["<path_to_license_key_file>"];
$CFG{opt}{noipc}=1;
$CFG{opt}{upgrade}=1;
$CFG{prod}="STORAGE80";
$CFG{systems}=[ qw(sys1) ];
1;
```

Performing post-upgrade tasks

This chapter includes the following topics:

- [Optional configuration steps](#)
- [Re-joining the backup boot disk group into the current disk group](#)
- [Reverting to the backup boot disk group after an unsuccessful upgrade](#)
- [Recovering VVR if automatic upgrade fails](#)
- [Resetting DAS disk names to include host name in FSS environments](#)
- [Upgrading disk layout versions](#)
- [Upgrading VxVM disk group versions](#)
- [Updating variables](#)
- [Setting the default disk group](#)
- [Verifying the Storage Foundation upgrade](#)

Optional configuration steps

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

You can perform the following optional configuration steps:

- If Volume Replicator (VVR) is configured, do the following steps in the order shown:
 - Reattach the RLINKs.
 - Associate the SRL.

- To encapsulate and mirror the boot disk, follow the procedures in the "Administering Disks" chapter of the *Storage Foundation Administrator's Guide*.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
See "[Upgrading VxVM disk group versions](#)" on page 48.

Re-joining the backup boot disk group into the current disk group

Note: Root Disk Encapsulation (RDE) is not supported on Linux from 7.3.1 onwards.

Perform this procedure to rejoin the backup boot disk if you split the mirrored boot disk during upgrade. After a successful upgrade and reboot, you no longer need to keep the boot disk group backup.

To re-join the backup boot disk group

- ◆ Re-join the *backup_bootdg* disk group to the boot disk group.

```
# /etc/vx/bin/vxrootadm -Y join backup_bootdg
```

where the `-Y` option indicates a silent operation, and *backup_bootdg* is the name of the backup boot disk group that you created during the upgrade.

Reverting to the backup boot disk group after an unsuccessful upgrade

Note: Root Disk Encapsulation (RDE) is not supported on Linux from 7.3.1 onwards.

Perform this procedure if your upgrade was unsuccessful and you split the mirrored boot disk to back it up during upgrade. You can revert to the backup that you created when you upgraded.

To revert the backup boot disk group after an unsuccessful upgrade

- 1 To determine the boot disk groups, look for the *rootvol* volume in the output of the `vxprint` command.

```
# vxprint
```

- 2 Use the `vx dg` command to find the boot disk group where you are currently booted.

```
# vx dg bootdg
```

- 3 Boot the operating system from the backup boot disk group.
- 4 Join the original boot disk group to the backup disk group.

```
# /etc/vx/bin/vxrootadm -Y join original_bootdg
```

where the `-Y` option indicates a silent operation, and *original_bootdg* is the boot disk group that you no longer need.

Recovering VVR if automatic upgrade fails

If the upgrade fails during the configuration phase, after displaying the VVR upgrade directory, the configuration needs to be restored before the next attempt. Run the scripts in the upgrade directory in the following order to restore the configuration:

```
# restoresrl  
# adddcn  
# srlprot  
# attrlink  
# start.rvg
```

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

Resetting DAS disk names to include host name in FSS environments

If you are on a version earlier than 7.1, the VxVM disk names in the case of DAS disks in FSS environments, must be regenerated to use the host name as a prefix. The host prefix helps to uniquely identify the origin of the disk. For example, the device name for the disk *disk1* on the host *sys1* is now displayed as *sys1_disk1*.

To regenerate the disk names, run the following command:

```
# vxddladm -c assign names
```

Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout version 13, 14, 15, 16, and 17. You can local mount disk layout version 7, 8, 9, 10, 11, and 12 to upgrade to a later disk layout version.

Note: If you plan to use 64-bit quotas, you must upgrade to the disk layout version 10 or later.

Disk layout version 7, 8, 9, 10, 11, and 12 are deprecated and you cannot cluster mount an existing file system that has any of these versions. To upgrade a cluster file system from any of these deprecated versions, you must local mount the file system and then upgrade it using the `vxupgrade` utility or the `vxfsconvert` utility.

The `vxupgrade` utility enables you to upgrade the disk layout while the file system is online. However, the `vxfsconvert` utility enables you to upgrade the disk layout while the file system is offline.

If you use the `vxupgrade` utility, you must incrementally upgrade the disk layout versions. However, you can directly upgrade to a desired version, using the `vxfsconvert` utility.

For example, to upgrade from disk layout version 7 to a disk layout version 17, using the `vxupgrade` utility:

```
# vxupgrade -n 8 /mnt
# vxupgrade -n 9 /mnt
# vxupgrade -n 10 /mnt
# vxupgrade -n 11 /mnt
# vxupgrade -n 12 /mnt
# vxupgrade -n 13 /mnt
# vxupgrade -n 14 /mnt
# vxupgrade -n 15 /mnt
# vxupgrade -n 16 /mnt
# vxupgrade -n 17 /mnt
```

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

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

Note: Veritas recommends that before you begin to upgrade the product version, you must upgrade the existing file system to the highest supported disk layout version. Once a disk layout version has been upgraded, it is not possible to downgrade to the previous version.

Use the following command to check your disk layout version:

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

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

Upgrading VxVM disk group versions

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions. VxVM can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks or use the features, upgrade the existing disk groups.

For 8.0, the Veritas Volume Manager disk group version is different than in previous VxVM releases. Veritas recommends that you upgrade the disk group version if you upgraded from a previous VxVM release.

After upgrading to SF 8.0, you must upgrade any existing disk groups that are organized by ISP. Without the version upgrade, configuration query operations continue to work fine. However, configuration change operations will not function correctly.

For more information about ISP disk groups, refer to the *Storage Foundation Administrator's Guide*.

Use the following command to find the version of a disk group:

```
# vxdg list diskgroup
```

To upgrade a disk group to the current disk group version, use the following command:

```
# vxdg upgrade diskgroup
```

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

Updating variables

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

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

Setting the default disk group

You may find it convenient to create a system-wide default disk group. The main benefit of creating a default disk group is that VxVM commands default to the default disk group. You do not need to use the `-g` option.

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

```
# vxctl defaultdg diskgroup
```

See the *Storage Foundation Administrator's Guide*.

Verifying the Storage Foundation upgrade

Refer to the *Verifying the Veritas InfoScale installation* chapter in the *Veritas InfoScale Installation Guide*.

Post configuration tasks

- [Chapter 7. Performing configuration tasks](#)

Performing configuration tasks

This chapter includes the following topics:

- [Switching on Quotas](#)
- [Enabling DMP support for native devices](#)
- [About configuring authentication for SFDB tools](#)

Switching on Quotas

This turns on the group and user quotas once all the nodes are upgraded to 8.0, if it was turned off earlier.

To turn on the group and user quotas

- ◆ Switch on quotas:

```
# vxquotaon -av
```

Enabling DMP support for native devices

Dynamic Multi-Pathing (DMP) is a component of SF. DMP supports Veritas Volume Manager (VxVM) volumes on DMP metadevices, and Veritas File System (VxFS) file systems on those volumes.

DMP can also provide multi-pathing functionality for the native operating system volumes and file systems on DMP devices.

For more information on using DMP with native devices, see the *Dynamic Multi-Pathing Administrator's Guide*.

After you install SF for the first time, use the following procedure to enable DMP support for native devices.

If DMP native support for native devices is enabled on a system before you upgrade SF, DMP native support is maintained when SF is upgraded.

To enable DMP support for native devices

- ◆ Turn on the tunable parameter to enable DMP support:

```
# vxddmpadm settune dmp_native_support=on
```

The `dmp_native_support` parameter is persistent.

About configuring authentication for SFDB tools

To configure authentication for Storage Foundation for Databases (SFDB) tools, perform the following tasks:

Configure the `vxdbd` daemon to require authentication

See [“Configuring vxdbd for SFDB tools authentication”](#) on page 52.

Add a node to a cluster that is using authentication for SFDB tools

Configuring vxdbd for SFDB tools authentication

To configure `vxdbd`, perform the following steps as the root user

- 1 Run the `sfae_auth_op` command to set up the authentication services.

```
# /opt/VRTS/bin/sfae_auth_op -o setup
Setting up AT
Starting SFAE AT broker
Creating SFAE private domain
Backing up AT configuration
Creating principal for vxdbd
```

- 2 Stop the `vxdbd` daemon.

```
# /opt/VRTS/bin/sfae_config disable
vxdbd has been disabled and the daemon has been stopped.
```

- 3 Enable authentication by setting the `AUTHENTICATION` key to `yes` in the `/etc/vx/vxdbed/admin.properties` configuration file.

If `/etc/vx/vxdbed/admin.properties` does not exist, then use `cp /opt/VRTSdbed/bin/admin.properties.example /etc/vx/vxdbed/admin.properties`.

- 4 Start the `vxdbd` daemon.

```
# /opt/VRTS/bin/sfae_config enable
vxdbd has been enabled and the daemon has been started.
It will start automatically on reboot.
```

The `vxdbd` daemon is now configured to require authentication.

Configuration and Upgrade reference

- [Appendix A. Installation scripts](#)
- [Appendix B. Configuring the secure shell or the remote shell for communications](#)

Installation scripts

This appendix includes the following topics:

- [Installation script options](#)
- [About using the postcheck option](#)

Installation script options

[Table A-1](#) shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas InfoScale product scripts, except where otherwise noted.

Table A-1 Available command line options

Command Line Option	Function
-allpkgs	Displays all RPMs required for the specified product. The RPMs are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-comcleanup	The <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.
-comsetup	The <code>-comsetup</code> option is used to set up the ssh or rsh communication between systems without requests for passwords or passphrases.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-configcps	The <code>-configcps</code> option is used to configure CP server on a running system or cluster.
-configure	Configures the product after installation.
-disable_dmp_native_support	Disables Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases RPM upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.
-fencing	Configures I/O fencing in a running cluster.
-fips	The <code>-fips</code> option is used to enable or disable security with fips mode on a running VCS cluster. It could only be used together with <code>-security</code> or <code>-securityonnode</code> option.
-hostfile <i>full_path_to_file</i>	Specifies the location of a file that contains a list of hostnames on which to install.
-install	Used to install products on system
-online_upgrade	Used to perform online upgrade. Using this option, the installer upgrades the whole cluster and also supports customer's application zero down time during the upgrade procedure. Now this option only supports VCS and ApplicationHA.
-patch_path	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-patch2_path	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-patch3_path	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch4_path	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch5_path	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-keyfile <i>ssh_key_file</i>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-I ssh_key_file</code> to every SSH invocation.
-kickstart <i>dir_path</i>	Produces a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of required RPMs in the correct order for installing, in a format that can be used for Kickstart installations. The <i>dir_path</i> indicates the path to the directory in which to create the file.
-license	Registers or updates product licenses on the specified systems.
-logpath <i>log_path</i>	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.
-noipc	Disables the installer from making outbound networking calls to Veritas Services and Operations Readiness Tool (SORT) in order to automatically obtain patch and release information updates.
-nolic	Allows installation of product RPMs without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-pkgtable	Displays product's RPMs in correct installation order by group.
-postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Veritas recommends doing a precheck before installing a product.
-prod	Specifies the product for operations.
-component	Specifies the component for operations.
-redirect	Displays progress details without showing the progress bar.
-require	Specifies an installer patch file.
-responsefile <i>response_file</i>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See "About configuring secure shell or remote shell communication modes before installing products" on page 63.
-security	The -security option is used to convert a running VCS cluster between secure and non-secure modes of operation.
-securityonenode	The -securityonenode option is used to configure a secure cluster node by node.
-securitytrust	The -securitytrust option is used to setup trust with another broker.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-settnables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the <code>-tunablesfile</code> option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-timeout	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 prevents the script from timing out. The <code>-timeout</code> option does not work with the <code>-serial</code> option.
-tmppath <i>tmp_path</i>	Specifies a directory other than <code>/opt/VRTStmp</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.
-tunables	Lists all supported tunables and create a tunables file template.
-tunables_file <i>tunables_file</i>	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-uninstall	This option is used to uninstall the products from systems.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.
-version	Checks and reports the installed products and their versions. Identifies the installed and missing RPMs and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing RPMs and patches where applicable. Lists the installed patches, patches, and available updates for the installed product if an Internet connection is available.
-yumgroupxml	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 and supported RHEL compatible distributions only.

About using the postcheck option

You can use the installer's post-check to determine installation-related problems and to aid in troubleshooting.

Note: This command option requires downtime for the node.

When you use the `postcheck` option, it can help you troubleshoot the following VCS-related issues:

- The heartbeat link does not exist.
- The heartbeat link cannot communicate.
- The heartbeat link is a part of a bonded or aggregated NIC.
- A duplicated cluster ID exists (if LLT is not running at the check time).
- The `VRTSllt` pkg version is not consistent on the nodes.
- The `llt-linkinstall` value is incorrect.

- The `/etc/llthosts` and `/etc/llttab` configuration is incorrect.
- the `/etc/gabtab` file is incorrect.
- The incorrect GAB `linkinstall` value exists.
- The VRTSgab pkg version is not consistent on the nodes.
- The `main.cf` file or the `types.cf` file is invalid.
- The `/etc/VRTSvcs/conf/sysname` file is not consistent with the hostname.
- The cluster UUID does not exist.
- The `uuidconfig.pl` file is missing.
- The VRTSvcs pkg version is not consistent on the nodes.
- The `/etc/vxfenmode` file is missing or incorrect.
- The `/etc/vxfendg` file is invalid.
- The `vxfen` link-install value is incorrect.
- The VRTSvxfen pkg version is not consistent.

The `postcheck` option can help you troubleshoot the following SFHA or SFCFSHA issues:

- Volume Manager cannot start because the `/etc/vx/reconfig.d/state.d/install-db` file has not been removed.
- Volume Manager cannot start because the `volboot` file is not loaded.
- Volume Manager cannot start because no license exists.
- Cluster Volume Manager cannot start because the CVM configuration is incorrect in the `main.cf` file. For example, the `Autostartlist` value is missing on the nodes.
- Cluster Volume Manager cannot come online because the node ID in the `/etc/llthosts` file is not consistent.
- Cluster Volume Manager cannot come online because Vxfen is not started.
- Cluster Volume Manager cannot start because `gab` is not configured.
- Cluster Volume Manager cannot come online because of a CVM protocol mismatch.
- Cluster Volume Manager group name has changed from "cvm", which causes CVM to go offline.

You can use the installer's post-check option to perform the following checks:

General checks for all products:

- All the required RPMs are installed.

- The versions of the required RPMs are correct.
- There are no verification issues for the required RPMs.

Checks for Volume Manager (VM):

- Lists the daemons which are not running (`vxattachd`, `vxconfigbackupd`, `vxesd`, `vxrelocd` ...).
- Lists the disks which are not in 'online' or 'online shared' state (`vxdisk list`).
- Lists the diskgroups which are not in 'enabled' state (`vx dg list`).
- Lists the volumes which are not in 'enabled' state (`vxprint -g <dgname>`).
- Lists the volumes which are in 'Unstartable' state (`vxinfo -g <dgname>`).
- Lists the volumes which are not configured in `/etc/fstab`.

Checks for File System (FS):

- Lists the VxFS kernel modules which are not loaded (`vxfs/fdd/vxportal`).
- Whether all VxFS file systems present in `/etc/fstab` file are mounted.
- Whether all VxFS file systems present in `/etc/fstab` are in disk layout 12 or higher.
- Whether all mounted VxFS file systems are in disk layout 12 or higher.

Checks for Cluster File System:

- Whether FS and ODM are running at the latest protocol level.
- Whether all mounted CFS file systems are managed by VCS.
- Whether `cvm` service group is online.

Configuring the secure shell or the remote shell for communications

This appendix includes the following topics:

- [About configuring secure shell or remote shell communication modes before installing products](#)
- [Manually configuring passwordless ssh](#)
- [Setting up ssh and rsh connection using the installer -comsetup command](#)
- [Setting up ssh and rsh connection using the pwdutil.pl utility](#)
- [Restarting the ssh session](#)
- [Enabling rsh for Linux](#)

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

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

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

Note: When installing on an RHEL5 / OEL5 system with SELinux enabled, only ssh is supported due to RedHat's SELinux policy restrictions.

You can set up ssh and rsh connections in many ways.

- You can manually set up the ssh and rsh connection with UNIX shell commands.
- You can run the `installer -comsetup` command to interactively set up ssh and rsh connection.
- You can run the password utility, `pwdutil.pl`.

This section contains an example of how to set up ssh password free communication. The example sets up ssh between a source system (sys1) that contains the installation directories, and a target system (sys2). This procedure also applies to multiple target systems.

Note: The product installer supports establishing passwordless communication.


Manually configuring passwordless ssh

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

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

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

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



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

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

To create the DSA key pair

- 1 On the source system (sys1), log in as root, and navigate to the root directory.

```
sys1 # cd /root
```

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

```
sys1 # ssh-keygen -t dsa
```

System output similar to the following is displayed:

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

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

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

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

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

- 5 Output similar to the following lines appears.

```
Your identification has been saved in /root/.ssh/id_dsa.  
Your public key has been saved in /root/.ssh/id_dsa.pub.  
The key fingerprint is:  
1f:00:e0:c2:9b:4e:29:b4:0b:6e:08:f8:50:de:48:d2 root@sys1
```

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

- 1 From the source system (sys1), move the public key to a temporary file on the target system (sys2).

Use the secure file transfer program.

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

Use the following command for secure file transfer:

```
sys1 # sftp sys2
```

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

```
Connecting to sys2 ...  
The authenticity of host 'sys2 (10.182.00.00)'  
can't be established. DSA key fingerprint is  
fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.  
Are you sure you want to continue connecting (yes/no)?
```

- 2 Enter `yes`.

Output similar to the following is displayed:

```
Warning: Permanently added 'sys2,10.182.00.00'  
(DSA) to the list of known hosts.  
root@sys2 password:
```

- 3 Enter the root password of sys2.
- 4 At the `sftp` prompt, type the following command:

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

The following output is displayed:

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

- 5 To quit the SFTP session, type the following command:

```
sftp> quit
```

- 6 Add the `id_dsa.pub` keys to the `authorized_keys` file on the target system. To begin the `ssh` session on the target system (sys2 in this example), type the following command on sys1:

```
sys1 # ssh sys2
```

Enter the root password of sys2 at the prompt:

```
password:
```

Type the following commands on sys2:

```
sys2 # cat /root/id_dsa.pub >> /root/.ssh/authorized_keys  
sys2 # rm /root/id_dsa.pub
```

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

```
sys1 # exec /usr/bin/ssh-agent $SHELL  
sys1 # ssh-add
```

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

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

To verify that you can connect to a target system

- 1 On the source system (sys1), enter the following command:

```
sys1 # ssh -l root sys2 uname -a
```

where sys2 is the name of the target system.

- 2 The command should execute from the source system (sys1) to the target system (sys2) without the system requesting a passphrase or password.
- 3 Repeat this procedure for each target system.

Setting up ssh and rsh connection using the installer -comsetup command

You can interactively set up the ssh and rsh connections using the `installer -comsetup` command.

Enter the following:

```
# ./installer -comsetup
```

```
Input the name of the systems to set up communication:
```

```
Enter the <platform> system names separated by spaces:
```

```
[q,?] sys2
```

```
Set up communication for the system sys2:
```

```
Checking communication on sys2 ..... Failed
```

```
CPI ERROR V-9-20-1303 ssh permission was denied on sys2. rsh permission was denied on sys2. Either ssh or rsh is required to be set up and ensure that it is working properly between the local node and sys2 for communication
```

```
Either ssh or rsh needs to be set up between the local system and sys2 for communication
```

```
Would you like the installer to setup ssh or rsh communication automatically between the systems?
```

```
Superuser passwords for the systems will be asked. [y,n,q,?] (y) y
```

```
Enter the superuser password for system sys2:
```

- 1) Setup ssh between the systems

- 2) Setup rsh between the systems
- b) Back to previous menu

Select the communication method [1-2,b,q,?] (1) 1

Setting up communication between systems. Please wait.
 Re-verifying systems.

Checking communication on sys2 Done

Successfully set up communication for the system sys2

Setting up ssh and rsh connection using the pwdutil.pl utility

The password utility, `pwdutil.pl`, is bundled under the `scripts` directory. The users can run the utility in their script to set up the ssh and rsh connection automatically.

```
# ./pwdutil.pl -h
```

Usage:

Command syntax with simple format:

```
pwdutil.pl check|configure|unconfigure ssh|rsh <hostname|IP addr>
[<user>] [<password>] [<port>]
```

Command syntax with advanced format:

```
pwdutil.pl [--action|-a 'check|configure|unconfigure']
            [--type|-t 'ssh|rsh']
            [--user|-u '<user>']
            [--password|-p '<password>']
            [--port|-P '<port>']
            [--hostfile|-f '<hostfile>']
            [--keyfile|-k '<keyfile>']
            [-debug|-d]
            <host_URI>
```

```
pwdutil.pl -h | -?
```

Table B-1 Options with pldutil.pl utility

Option	Usage
--action -a 'check configure unconfigure'	Specifies action type, default is 'check'.
--type -t 'ssh rsh'	Specifies connection type, default is 'ssh'.
--user -u '<user>'	Specifies user id, default is the local user id.
--password -p '<password>'	Specifies user password, default is the user id.
--port -P '<port>'	Specifies port number for ssh connection, default is 22
--keyfile -k '<keyfile>'	Specifies the private key file.
--hostfile -f '<hostfile>'	Specifies the file which list the hosts.
-debug	Prints debug information.
-h -?	Prints help messages.
<host_URI>	Can be in the following formats: <hostname> <user>:<password>@<hostname> <user>:<password>@<hostname>: <port>

You can check, configure, and unconfigure ssh or rsh using the `pldutil.pl` utility. For example:

- To check ssh connection for only one host:
`pldutil.pl check ssh hostname`
- To configure ssh for only one host:
`pldutil.pl configure ssh hostname user password`
- To unconfigure rsh for only one host:
`pldutil.pl unconfigure rsh hostname`
- To configure ssh for multiple hosts with same user ID and password:

```
pwdutil.pl -a configure -t ssh -u user -p password hostname1
hostname2 hostname3
```

- To configure ssh or rsh for different hosts with different user ID and password:

```
pwdutil.pl -a configure -t ssh user1:password1@hostname1
user2:password2@hostname2
```

- To check or configure ssh or rsh for multiple hosts with one configuration file:

```
pwdutil.pl -a configure -t ssh --hostfile /tmp/sshrsh_hostfile
```

- To keep the host configuration file secret, you can use the 3rd party utility to encrypt and decrypt the host file with password.

For example:

```
### run openssl to encrypt the host file in base64 format
# openssl aes-256-cbc -a -salt -in /hostfile -out /hostfile.enc
enter aes-256-cbc encryption password: <password>
Verifying - enter aes-256-cbc encryption password: <password>
```

```
### remove the original plain text file
# rm /hostfile
```

```
### run openssl to decrypt the encrypted host file
# pwdutil.pl -a configure -t ssh `openssl aes-256-cbc -d -a
-in /hostfile.enc`
enter aes-256-cbc decryption password: <password>
```

- To use the ssh authentication keys which are not under the default \$HOME/.ssh directory, you can use --keyfile option to specify the ssh keys. For example:

```
### create a directory to host the key pairs:
# mkdir /keystore
```

```
### generate private and public key pair under the directory:
# ssh-keygen -t rsa -f /keystore/id_rsa
```

```
### setup ssh connection with the new generated key pair under
the directory:
# pwdutil.pl -a configure -t ssh --keyfile /keystore/id_rsa
user:password@hostname
```

You can see the contents of the configuration file by using the following command:

```
# cat /tmp/sshrsh_hostfile
user1:password1@hostname1
user2:password2@hostname2
user3:password3@hostname3
user4:password4@hostname4

# all default: check ssh connection with local user
hostname5
The following exit values are returned:

0      Successful completion.
1      Command syntax error.
2      Ssh or rsh binaries do not exist.
3      Ssh or rsh service is down on the remote machine.
4      Ssh or rsh command execution is denied due to password is required.
5      Invalid password is provided.
255   Other unknown error.
```

Restarting the ssh session

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

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

To restart ssh

- 1 On the source installation system (sys1), bring the private key into the shell environment.

```
sys1 # exec /usr/bin/ssh-agent $SHELL
```

- 2 Make the key globally available for the user `root`

```
sys1 # ssh-add
```

Enabling rsh for Linux

The following section describes how to enable remote shell.

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

See [“Manually configuring passwordless ssh”](#) on page 64.

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

To enable rsh for RHEL

- ◆ Run the following commands to enable rsh passwordless connection:

```
# systemctl start rsh.socket
# systemctl start rlogin.socket
# systemctl enable rsh.socket
# systemctl enable rlogin.socket
# echo rsh >> /etc/securetty
# echo rlogin >> /etc/securetty
# echo "+ +" >> /root/.rhosts
```

To disable rsh for RHEL

- ◆ Run the following commands to disable rsh passwordless connection:

```
# systemctl stop rsh.socket
# systemctl stop rlogin.socket
# systemctl disable rsh.socket
# systemctl disable rlogin.socket
```