Veritas InfoScale™ 7.0.1 Installation Guide - Solaris 10 Sparc, Solaris 11 Sparc and Solaris 11 x64

VERITAS™

Veritas InfoScale™ Installation Guide

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

Product version: 7.0.1

Document version: 7.0.1 Rev 2

Legal Notice

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

Technical Support

Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within the company to answer your questions in a timely fashion.

Our support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization

- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information

- Upgrade assurance that delivers software upgrades

- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis

- Premium service offerings that include Account Management Services

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

www.veritas.com/support

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

Contacting Technical Support

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

www.veritas.com/support

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

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

- Product release level

- Hardware information

- Available memory, disk space, and NIC information

- Operating system

- Version and patch level

- Network topology

- Router, gateway, and IP address information

- Problem description:

  - Error messages and log files

  - Troubleshooting that was performed before contacting Technical Support

  - Recent software configuration changes and network changes

Licensing and registration

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

www.veritas.com/support

Customer service

Customer service information is available at the following URL:

www.veritas.com/support

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

- Questions regarding product licensing or serialization

- Product registration updates, such as address or name changes

- General product information (features, language availability, local dealers)

- Latest information about product updates and upgrades

- Information about upgrade assurance and support contracts

- Advice about technical support options

- Nontechnical presales questions

- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

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

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Contents

Installing the products for the first time

This chapter includes the following topics:

- Supported types of Installation

- Creating a new boot environment on Solaris 11

- Installing the Veritas InfoScale software using the Install Bundles feature

Supported types of Installation

You can use the script-based installer with install bundle method to install Veritas InfoScale 7.0.1.

Creating a new boot environment on Solaris 11

Before installing Veritas InfoScale 7.0.1 on a Solaris 11 host, you can optionally create a backup of the existing active boot environment (BE) and install Veritas InfoScale 7.0.1 on the present boot environment. This will help to rollback to the previous state of the operating system in the future if required.

To create a new boot environment as a backup

1   Identify the active boot environment (BE) by looking at the NR tag:

   # **beadm list**

2   Create the BE:

   # **beadm create *bename***

   For example,

   # **beadm create *pre_sfha_7.0.1***

Installing the Veritas InfoScale software using the Install Bundles feature

This section describes how to install a Veritas InfoScale product of 7.0 and 7.0.1 using the Install Bundles feature in one step.

Figure 1-1   Install flow of Veritas InfoScale
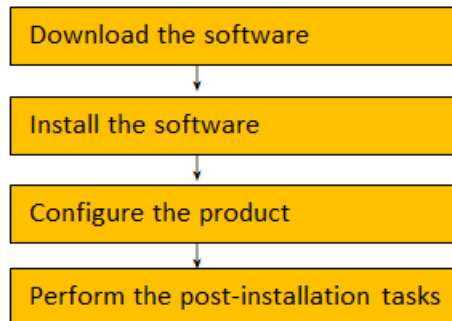
To install the Veritas InfoScale software 7.0.1 using Install Bundles:

1   Download 7.0 version of InfoScale product from https://myveritas.com.

2   Extract the tar ball into the /infoscale7.0/ directory.

3   Download Veritas InfoScale 7.0.1 from https://sort.veritas.com/patches.

4   Extract it to the /infoscale7.0.1 directory.

5 Change to the `/infoscale7.0.1` directory by entering:

```
# cd /infoscale7.0.1
```

6 Invoke the `installmr` script with `-base_path` option to install 7.0 and 7.0.1. Enter:

```
./installmr -base_path /infoscale7.0/dvd1-sol_sparc/sol10_sparc/
```

Note: The actual path varies depending on your operating system.

---

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

To configure the product, please see the 7.0 Configuration and Upgrade Guide.

Installing Veritas InfoScale using operating system-specific methods

This chapter includes the following topics:

- Installing Veritas InfoScale using JumpStart

- Installing Veritas InfoScale on Solaris 11 using Automated Installer

Installing Veritas InfoScale using JumpStart

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

Summary of tasks

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

2. Read the JumpStart installation instructions.

3. Generate the finish scripts.

   See "Generating the finish scripts" on page 14.

4. Modify the rules file for JumpStart.

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

5. Prepare installation resources.

   See "Preparing installation resources" on page 16.

6 Click **Upload finish.sh**, then continue with the operating systsem installation.

Note: The system is restarted after the packages are installed. If you choose to encapsulate the root disk on your systems, the systems start with an encapsulated root disk.

---

7 On each client node, run the following command to install the Veritas InfoScale product packages and patches:

For Solaris SPARC:

```
Ok> boot net:dhcp -v - install
```

8 After the reboot, run the `installer` command from the `/opt/VRTS/install` directory to configure the Veritas InfoScale software.

```
# /opt/VRTS/install/installer
```

Generating the finish scripts

Perform these steps to generate the finish script to install Veritas InfoScale.

To generate the finish script

1 Download the Veritas InfoScale 7.0.1 patch from SORT and run the `installmr` program to generate the scripts.

```
# ./installmr -jumpstart directory_to_generate_scripts
```

where the *directory_to_generate_scripts* is the location where you want to put the scripts.

For example:

```
# ./installmr -jumpstart /js_scripts
```

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

3 Specify a disk group name for the root disk.

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

4 Specify private region length.

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

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

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

6 JumpStart finish scripts of Veritas InfoScale products, and encapsulation scripts are generated in the directory you specified in step 1. Output resembles:

```
The finish scripts for InfoScale Availability is generated at
/export/jumpstart_availability.fin

The finish scripts for InfoScale Enterprise is generated at
/export/jumpstart_enterprise.fin

The finish scripts for InfoScale Foundation is generated at
/export/jumpstart_foundation.fin

The finish scripts for InfoScale Storage is generated at
/export/jumpstart_storage.fin
```

List the js_scripts directory.

```
# ls /js_scripts
```

You could select scripts according to the products you want to install and copy them to the BUILDSRC NFS shared location. For example, /export/config where you mounted the BUILDSRC.

For InfoScale Availability:

```
jumpstart_availability.fin
```

For InfoScale Enterprise:

```
jumpstart_enterprise.fin
```

For InfoScale Foundation:

```
jumpstart_foundation.fin
```

For InfoScale Storage:

```
jumpstart_storage.fin
```

7 Modify the JumpStart script according to your requirements. You must modify the BUILDSRC and ENCAPSRC values. Keep the values aligned with the resource location values.

```
BUILDSRC="hostname_or_ip:/path_to_pkgs_patches_scripts"
```

Example: **BUILDSRC=10.209.100.100:/export/config**

```
// If you don't want to encapsulate the root disk automatically
// comment out the following line.
ENCAPSRC="hostname_or_ip:/path_to_encap_script"
```

8 ■ If you want to install other products packages in the Veritas InfoScale, then use the installmr of 7.0.1 to get the list:

■ `minpkgs`

■ `recpkgs`

■ `allpkgs`
You can take the following command as an example:

```
# ./installmr -pkgtable
```

Use the list of packages that is generated to replace the package list in the finish scripts.

■ If you want to install other products patches in the Veritas InfoScale, then use the product-specific install command with one of the following options to get a list of patches in the order to be installed:

```
# ./installmr -listpatches
```

9 Once the installation is complete, refer to installation and configuration guide for the respective product from 7.0 to proceed with the configuration.

Preparing installation resources

Prepare resources for the JumpStart installation.

To prepare the resources

1   Copy the contents of Veritas InfoScale7.0 to buildsrc:

```
# cd /infoscale7.0
# cp -r pkgs BUILDSRC
```

Note: The VRTSaslapm package of 7.0 should be replaced with the 7.0.1 VRTSaslspm package.

2   Copy the contents of 7.0.1 patch to buildsrc:

```
# cd /infoscale7.0.1
# cp -r patches BUILDSRC
```

Note: After you copied the patches, you must uncompress them using the gunzip and tar commands.

3   Generate the response file for the package list that you found when you generated the finish script.

See

To view the patches, packages and operating systems for your Veritas InfoScale product use the `installmr -listpatches` command, type:

```
# ./installmr -listpatches

# cd BUILDSRC/dvd1-sol_sparc/sol10_sparc/pkgs/
# pkgask -r package_name.response -d /BUILDSRC/pkgs/packages_name.pkg
```

Note: The actual path varies depending on your operating system.

4 Create the adminfile file named admin under *BUILDSRC*/pkgs/ directory. The adminfile file's contents follow:

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

5 If you want to encapsulate the root disk automatically when perform the JumpStart installation, copy the scripts `encap_bootdisk_vm.fin` created when you generated the finish script to *ENCAPSRC*.

See "Generating the finish scripts" on page 14.

Adding language pack information to the finish file

For the language pack, copy the language packages from the Veritas InfoScale 7.0 to the shared storage.

```
# cd /infoscale7.0/pkgs
# cp -r * BUILDSRC/pkgs
```

Add lines for the language packages in the finish script. If the finish file resembles:

```
. . .
for PKG in VRTSperl VRTSvlic VRTSspt . . .

do
...
done
```

Installing Veritas InfoScale on Solaris 11 using Automated Installer

You can use the Oracle Solaris Automated Installer (AI) to install the Solaris 11 operating system and Storage Foundation product on multiple client systems in a

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

Using Automated Installer

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

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

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

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

- If two systems have different architectures or need to be installed with different versions of the Oracle Solaris OS, create two AI services. Then, associate each AI service with a different AI image

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

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

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

Using AI to install the Solaris 11 operating system and Veritas InfoScale products

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

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

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

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

2 Set up the Veritas InfoScale package repository by running the following commands to startup necessary SMF services and create directories:

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

3 Run the following commands to set up IPS repository for Veritas InfoScale SPARC and x86 packages:

```
# mkdir -p /ai/repo_symc_sparc
# pkgrepo create /ai/repo_symc_sparc
# pkgrepo add-publisher -s /ai/repo_symc_sparc Veritas
```

- For 7.0:

  ```
  # pkgrecv -s base_release_media/pkgs/VRTSpkgs.p5p -d
  /ai/repo_symc_sparc '*'
  ```

- For 7.0.1:

  ```
  # pkgrecv -s base_release_media/pkgs/VRTSpkgs.p5p -d
  /ai/repo_symc_sparc '*'

  # pkgrecv -s patch_release_media/patches/VRTSpatches.p5p -d
  /ai/repo_symc_sparc '*'
  ```

Set service property and enable it:

```
# svccfg -s pkg/server list
# svcs -a | grep pkg/server
# svccfg -s pkg/server add symcsparc
# svccfg -s pkg/server:symcsparc addpg pkg application
# svccfg -s pkg/server:symcsparc setprop pkg/port=10003
# svccfg -s pkg/server:symcsparc setprop pkg/inst_root=
/ai/repo_symc_sparc
# svccfg -s pkg/server:symcsparc addpg general framework
# svccfg -s pkg/server:symcsparc addpropvalue general/complete
astring: symcsparc
# svccfg -s pkg/server:symcsparc addpropvalue general/enable
boolean: true
# svcs -a | grep pkg/server
# svcadm refresh application/pkg/server:symcsparc
# svcadm enable application/pkg/server:symcsparc
```

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

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

Check the following URL on IE or Firefox browser:

http://*host*:10003

4   Set up the install service on the AI server.

Run the following command:

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

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

Create an install service.

For example:

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

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

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

```
# mkdir /ai/manifests
# media/installmr -ai /ai/manifests
```

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

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

or

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

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

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

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

Preparing to upgrade to Veritas InfoScale 7.0.1

This chapter includes the following topics:

Downloading required software to upgrade to 7.0.1

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

To download required software to upgrade to 7.0.1

1   Download InfoScale 7.0.1 from https://sort.Veritas.com/patches.

2   Extract it to a directory, say /infoscale.

3   On Solaris 11, Veritas recommends you creating a backup boot environment.

Note: If you are upgrading from versions earlier than 7.0 to 7.0.1 using the Install Bundles, you must download both 7.0 and 7.0.1.

Prerequisites for upgrading to 7.0.1

If you are upgrading from 7.0, see the following list for prerequisites for upgrading to the 7.0.1 release:

- For any product in the Veritas Storage Foundation stack, you must have the 7.0.1 release binaries.

- Each system must have sufficient free space to accommodate patches.

- The full list of prerequisites can be obtained by running `./installmr -precheck`.

- Make sure to download the latest patches for the installer.
  See "Downloading required software to upgrade to 7.0.1 " on page 23.

Supported upgrade types for Veritas InfoScale 7.0.1

Veritas InfoScale supports various ways of upgrading your cluster to the latest version. Choose a method that best suits your environment and supports your planned upgrade path.

Table 3-1 lists the supported types of upgrade.

Table 3-1   Supported types of upgrade

Type of upgrade	Abstract
Full upgrade	A full upgrade involves upgrading all the nodes in the cluster at the same time. All components are upgraded during the process. The cluster remains unavailable for the duration of the upgrade.
Online upgrade	The online upgrade involves upgrading the whole cluster and supporting customer's application zero down time during the upgrade procedure. Now it only support VCS component.
Solaris Live Upgrade or Boot Environment upgrade (On Solaris 11 x86 system)	Solaris Live Upgrade or Boot Environment upgrade provides a method of upgrading a system while the system continues to operate.

Supported upgrade paths for Veritas InfoScale 7.0.1

You can run the `installmr` script with Install Bundles to upgrade Veritas InfoScale to 7.0.1 .

For information on operating systems that are supported for 7.0.1, see *System requirements* in *Veritas InfoScale 7.0.1 Release Notes*.

The following figures list the supported upgrade paths to 7.0.1 on Solaris x64 and Sparc.

Figure 3-1   Supported upgrade paths to 7.0.1 on Solaris SPARC

Supported upgrade paths for
Solaris 10 SPARC and Solaris 11 Sparc
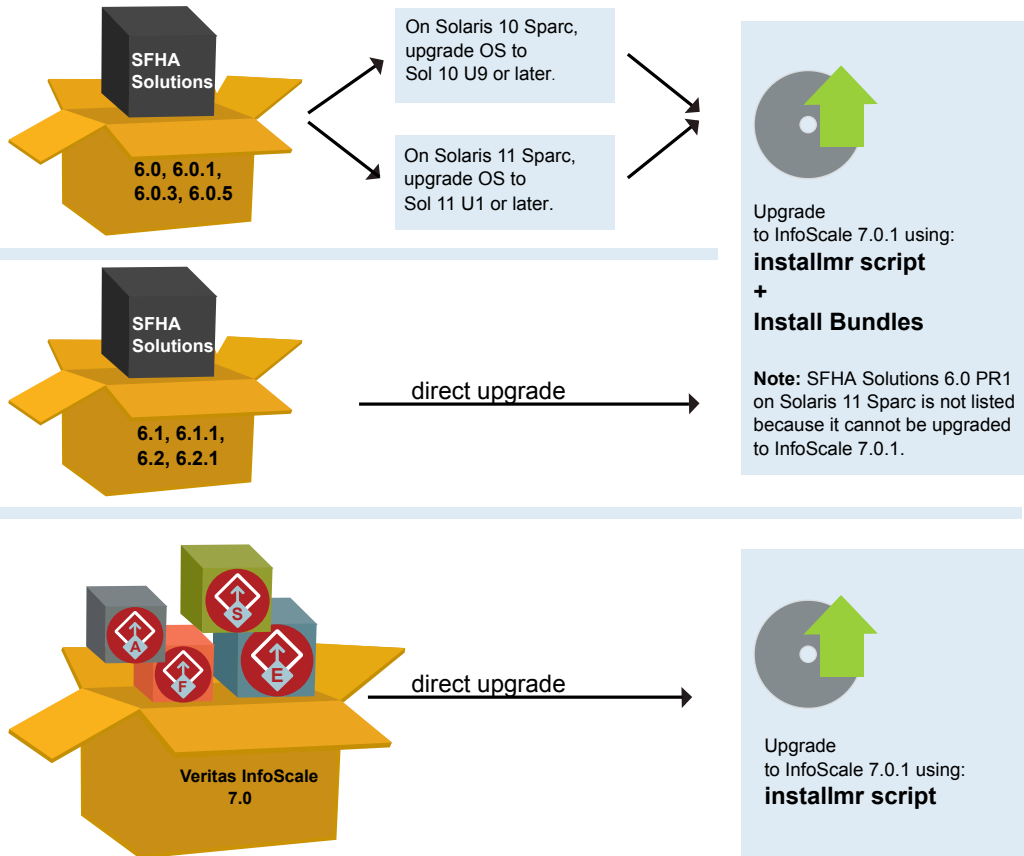
Figure 3-2   Supported upgrade paths to 7.0.1 on Solaris x64

Supported upgrade paths for Solaris 10 x64 and Solaris 11 x64

SFHA Solutions

6.0.5

On Solaris 10 x64, upgrade OS to Sol 10 U9 or later.

On Solaris 11 x64, upgrade OS to Sol 11 U1 or later.

Upgrade to InfoScale 7.0.1 using: **installmr script** **+** **Install Bundles**

Veritas InfoScale Availability 7.0

direct upgrade

Upgrade to InfoScale 7.0.1 using: **installmr script**
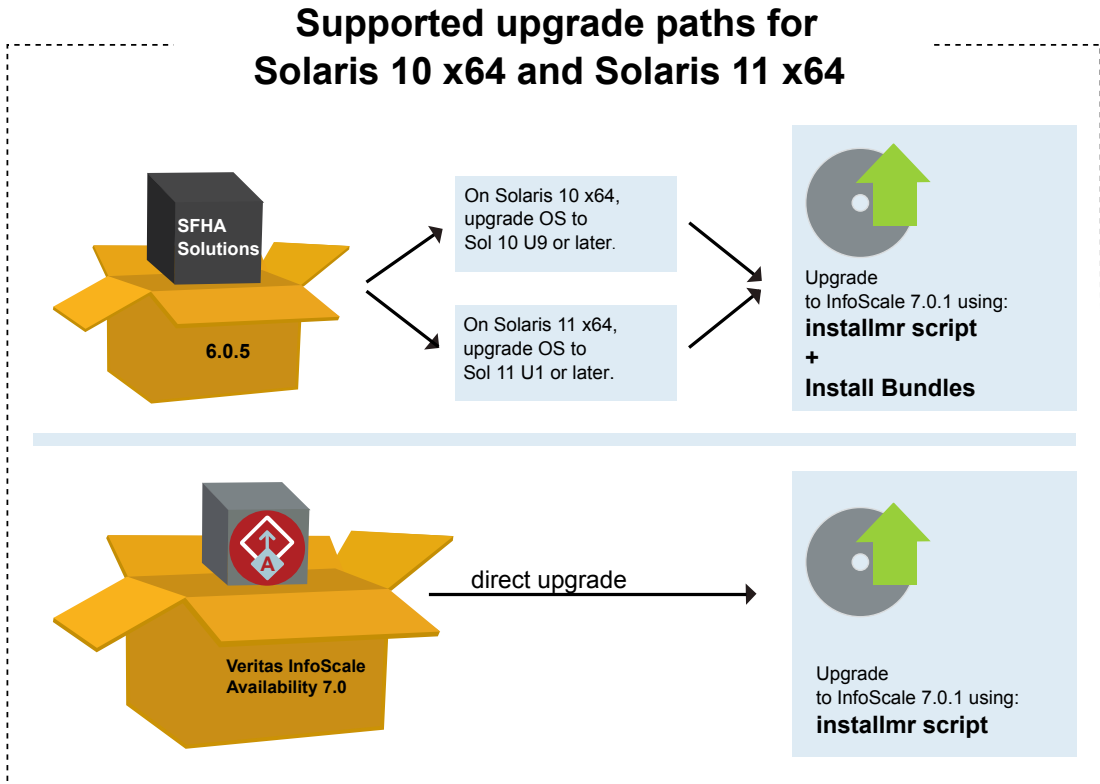
The following notes apply to the Solaris x64 platform:

- InfoScale Availability is the only supported product on the Solaris x64 platform.

- For InfoScale Availability 7.0.1, you can upgrade VCS 6.0.5 to InfoScale Availability 7.0.1. For VCS 6.0PR1, VCS 6.0.1 and VCS 6.0.3, you can upgrade to VCS 6.0.5 first and then upgrade VCS 6.0.5 to Availability 7.0.1.

- Coexistence of SF 6.0.5 with Availability 7.0.1 is supported.

- Upgrading the whole stack of SFHA, SFCFSHA, SFRAC is not supported.

Preparing to upgrade Volume Replicator

Before installing or upgrading Volume Replicator (VVR):

- Confirm that your system has enough free disk space to install VVR.

- Make sure you have root permissions. You must have root permissions to perform the install and upgrade procedures.

- If replication using VVR is configured, 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.

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

See the *Veritas InfoScale Replication Administrator's Guide* for more information.

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

Planning an upgrade from the previous VVR version

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

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

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

Also, if you specify TCP as the network protocol, the VVR versions on the Primary and Secondary determine whether the checksum is calculated. As shown in Table 3-2, if either the Primary or Secondary are running a version of VVR prior to 7.0.1, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 7.0.1, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

Table 3-2   VVR versions and checksum calculations

VVR prior to 7.0.1 (DG version <= 140)	VVR 7.0.1 (DG version >= 150)	VVR calculates checksum TCP connections?
Primary	Secondary	Yes
Secondary	Primary	Yes
Primary and Secondary		Yes
	Primary and Secondary	No

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

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

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

Additional settings for using VVR in a localized environment

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

Japanese version of VVR is installed, then the messages are displayed in the Japanese locale, if the client locale is Japanese.

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

- Install the required client locale from the Operating System disc.

- Install the required Volume Manager and VVR localized packages.

- Set the client locale, before using any of the VVR interfaces:

  - For the VVR command line, set the locale using the appropriate method for your operating system.

  - For VRW, select the locale from the VRW login page.

Downloading Veritas InfoScale 7.0.1

1 Download Veritas InfoScale 7.0 from https://myveritas.com.

2 Extract the tar ball into a directory called /infoscale7.0.

3 Download Veritas InfoScale 7.0.1 from https://sort.veritas.com/patches.

4 Extract it to a directory called /infoscale7.0.1.

Note: If you are upgrading from versions earlier than 7.0 to 7.0.1 using the Install Bundles, you must download both 7.0 and 7.0.1.

Upgrading to 7.0.1 from releases earlier than 7.0

This chapter includes the following topics:

Performing a full upgrade with Install Bundles

The following procedure describes how to upgrade to 7.0.1 with Install Bundles from releases earlier than 7.0.

Note: If you are upgrading from releases earlier than 7.0, Veritas suggests you upgrade with Install Bundles.

Performing a full upgrade of VCS using Install Bundles

You can use the installer to upgrade VCS.

To upgrade VCS using the product installer

1   Log in as superuser.

2   Change to the `/infoscale7.0.1` directory.

3   On Solaris 10, if zones are configured, you need to set AutoStart attribute to 0 for zone groups.

Set AutoStart attribute to 0 for zone groups:

`# haconf -makerw`

`# hagrp -modify <zonegroup> AutoStart 0`

`# haconf -dump -makero`

4   Invoke the `installmr` script with `-base_path` option to upgrade to 7.0.1:

`# ./installmr -base_path /infoscale7.0/dvd1-sol_sparc/sol10_sparc/`

Note: The actual path varies depending on your operating system version.

---

5   From the opening Selection Menu, choose: **G** for "Upgrade a Product."

6   Choose **1** for Full Upgrade.

7   Enter the names of the nodes that you want to upgrade. Use spaces to separate node names. Press the Enter key to proceed.

The installer runs some verification checks on the nodes.

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

The installer lists the packages to upgrade.

9 The installer displays the following question before it stops the product processes, if the cluster is configured in secure mode and the previous product version is less than 6.2:

```
Do you want to grant read access to everyone? [y,n,q,?]
```

To grant read access to all authenticated users, type **y**.

To grant permissions to specific user group, type **n**.

```
Do you want to provide any usergroups that you would like to\
grant read access?[y,n,q,?]
```

To specify user groups and grant them read access, type **y**.

To grant read access only to root users, type **n**. Then the installer grants read access read access to the root users.

Enter the user group names that you want to grant read access and separate them by spaces. If you want to grant read access to a user group on a specific node, enter *usergroup@node*. If you want to grant read access to user groups on any cluster node, enter *usergroup*. If some user groups are not created yet, create the user groups after configuration if needed.

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

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

The installer lists the nodes that Veritas recommends you to restart, if needed.

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

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

Note: 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.

---

12 On Solaris 10, perform the following steps to set the zone group appropriately
after upgrade.

Sync non-global zone with the global:

```
# zoneadm -z <zonename> attach -u
```

Reset the AutoStart attribute to 1:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 1
```

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

Make zone group online:

```
# hagrp -online <zonegroup> -sys <sysname>
```

Performing a full upgrade of SFHA using Install Bundles

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

Upgrading SFHA with Install Bundles

This section describes upgrading to the current Veritas InfoScale, and you do not
intend to upgrade your Solaris version. Only use this procedure if you are already
running a version of Solaris that is supported with 7.0.1.

To upgrade Veritas InfoScale

1 Log in as superuser.

2 Unmount any mounted VxFS file systems that are not managed by VCS.

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

If any VxFS file systems are mounted with the QuickLog feature, QuickLog must be disabled before upgrading.

3 If you are upgrading a high availability (HA) product, take all service groups offline.

List all service groups:

```
# /opt/VRTSvcs/bin/hagrp -list
```

For each service group listed, take it offline:

```
# /opt/VRTSvcs/bin/hagrp -offline service_group \
  -sys system_name
```

4 Enter the following commands on each node to freeze HA service group operations:

```
# haconf -makerw
# hasys -freeze -persistent nodename
# haconf -dump -makero
```

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

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

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

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

---

7   On Solaris 10, if zones are configured, you need to set AutoStart attribute to 0 for zone groups.

Set AutoStart attribute to 0 for zone groups:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 0
```

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

8   Change to the `/infoscale7.0.1` directory.

9   Run the `installmr` command as shown in this example:

```
# ./installmr -base_path /infoscale7.0/dvd1-sol_sparc/sol10_sparc/
```

Note: The actual path varies depending on your operating system version.

---

10   Enter `G` to upgrade and press Return.

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

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

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

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

13   You can perform this step if you upgrading from SFHA 5.1 SP1 for Solaris.

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 book disk group before the upgrade proceeds. If you want to split the boot disk group to create a backup, answer **y**.

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

---

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

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

---

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

16 The installer displays the following question before it stops the product processes, if the cluster is configured in secure mode and the previous product version is less than 6.2:

```
Do you want to grant read access to everyone? [y,n,q,?]
```

To grant read access to all authenticated users, type **y**.

To grant permissions to specific user group, type **n**.

```
Do you want to provide any usergroups that you would like to/
grant read access?[y,n,q,?]
```

To specify user groups and grant them read access, type **y**.

To grant read access only to root users, type **n**. Then the installer grants read access read access to the root users.

Enter the user group names that you want to grant read access and separate them by spaces. If you want to grant read access to a user group on a specific node, enter *usergroup@node*. If you want to grant read access to user groups on any cluster node, enter *usergroup*. If some user groups are not created yet, create the user groups after configuration if needed.

17 Stop the product's processes.

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

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

19 If the product is licensed with a stale (old) key, the installer would prompt users to update the key.

20 The installer verifies, configures, and starts the Veritas InfoScale software.

> **Note:** 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.

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

22 Unfreeze the service groups.

23 Take the service groups online.

24 On Solaris 10, perform the following steps to set the zone group appropriately after upgrade.

Sync non-global zone with the global:

```
# zoneadm -z <zonename> attach -u
```

Reset the AutoStart attribute to 1:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 1
```

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

Make zone group online:

```
# hagrp -online <zonegroup> -sys <sysname>
```

Performing a full upgrade of SFCFSHA using Install Bundles

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

Performing a full SFCFSHA upgrade with Install Bundles

Performing a full upgrade involves the following tasks:

- Ensuring that the file systems are clean

■   Performing the upgrade

Ensuring the file systems are clean

Before upgrading to SFCFSHA 7.0.1, ensure that the file systems are clean. To
ensure that the logs have been replayed and the file systems are marked clean:

To ensure the file systems are clean

1   Log in as superuser onto any node in the cluster.

2   Take the service group offline on each node of the cluster, which contains
VxFS and CFS resources:

```
# hagrp -offline group -any
```

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

Repeat this step for each SFCFSHA service group.

Note: This unmounts the CFS file systems.

---

3   Unmount all VxFS file systems not under VCS control:

```
# umount /mount_point
```

4   Check and repair each VxFS file system:

```
# fsck -F vxfs /dev/vx/rdsk/diskgroup/volume
```

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

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

Repeat this step for each file system.

Performing the upgrade

To perform the upgrade

1 Log in as superuser.

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

mount -p | grep vxfs

If any VxFS file systems are mounted, offline the group on each node of the cluster:

hagrp -offline group -any

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

If VxFS file systems are not managed by VCS then unmount them manually:

umount */mount_point*

Repeat this step for each SFCFSHA service group.

3 On Solaris 10, if zones are configured, you need to set AutoStart attribute to 0 for zone groups.

Set AutoStart attribute to 0 for zone groups:

 # haconf -makerw

 # hagrp -modify *<zonegroup>* **AutoStart 0**

 # haconf -dump -makero

4 Change to the /infoscale7.0.1 directory. Invoke the installmr script with -base_path option to upgrade to 7.0.1:

#./installmr -base_path /infoscale7.0/*dvd1-sol_sparc/sol10_sparc/*

Note: The actual path varies depending on your operating system version.

---

5 From the opening Selection Menu, choose: **G** for **Upgrade a Product**. Choose **1** for **Full Upgrade**.

6 You are prompted to enter the system names (in the following example, "sys1" and "sys2") on which the software is to be upgraded. Enter the system name or names and then press Return.

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

7 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA). Press **y** to agree and continue.

8 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**.

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

If the installer hangs or asks for a login password, sertup passwordless ssh or setup rsh from the system that run `installmr` to the system that need to be upgraded to 7.0.1. Then run the installer again.

10 After you accept EULA and the system checks complete, the installer displays a list of the packages that will be upgraded. Press Enter to continue with the upgrade.

11 The installer displays the following question before it stops the product processes, if the cluster is configured in secure mode and the previous product version is less than 6.2:

```
Do you want to grant read access to everyone? [y,n,q,?]
```

To grant read access to all authenticated users, type **y**.

To grant permissions to specific user group, type **n**.

```
Do you want to provide any usergroups that you would like to\
grant read access?[y,n,q,?]
```

To specify user groups and grant them read access, type **y**.

To grant read access only to root users, type **n**. Then the installer grants read access read access to the root users.

Enter the user group names that you want to grant read access and separate them by spaces. If you want to grant read access to a user group on a specific node, enter *usergroup@node*. If you want to grant read access to user groups on any cluster node, enter *usergroup*. If some user groups are not created yet, create the user groups after configuration if needed.

12 Output shows information that SFCFSHA must be stopped on a running system. Enter **y** to continue.

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

14 Press **Enter** again for summary information about logs and reboots.

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

Note: 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.

---

15 On Solaris 10, perform the following steps to set the zone group appropriately after upgrade.

Sync non-global zone with the global:

```
# zoneadm -z <zonename> attach -u
```

Reset the AutoStart attribute to 1:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 1
```

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

Make zone group online:

```
# hagrp -online <zonegroup> -sys <sysname>
```

16 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.

Performing a full upgrade of SF Oracle RAC using Install Bundles

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

- Preparing to perform a full upgrade to 7.0.1 on an SF Oracle RAC cluster
- Upgrading to SF Oracle RAC 7.0.1

Preparing to perform a full upgrade to 7.0.1 on an SF Oracle RAC cluster

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

To prepare to upgrade SF Oracle RAC

1 Log in as superuser to one of the nodes in the cluster.

2 Back up the following configuration files on your system: `main.cf`, `types.cf`, `CVMTypes.cf`, `CFSTypes.cf`, `OracleTypes.cf`, `OracleASMTypes.cf`, `PrivNIC.cf`, `MultiPrivNIC.cf`, `CRSResource.cf`, `/etc/llttab`, `/etc/llthosts`, `/etc/gabtab`, `/etc/vxfentab`, `/etc/vxfendg`, `/etc/vxfenmode`

For example:

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

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

If not, a warning message will be displayed after `installmr` upgrade prechecks.

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

---

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

If the applications are under VCS control:

```
# hagrp -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

5 Stop all Oracle RAC resources.

For Oracle RAC 11g, and Oracle RAC 12c:

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

  ```
  # hagrp -offline oracle_group -any
  ```

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

  - For Oracle RAC 11g:

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

  - For Oracle RAC 12c:

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

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

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

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

■ For Oracle RAC 11g:

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

■ For Oracle RAC 12c:

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

7 Stop VCS on all nodes:

```
# hastop -all
```

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

```
# mount -v |grep vxfs
```

```
# fuser -c /mount_point
```

```
# umount /mount_point
```

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

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

9 If you plan to continue using the Storage Foundation for Databases (SFDB) tools, you must prepare to migrate the SFDB repository database before upgrading to 7.0.1.

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

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

If you are upgrading to Solaris 10 Update 10, apply the following Oracle patches: `144524-02` (SPARC). See the Oracle documentation for instructions.

11 On Solaris 10, if zones are configured, you need to set AutoStart attribute to 0 for zone groups.

Set AutoStart attribute to 0 for zone groups:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 0
```

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

Pre-upgrade tasks for migrating the SFDB repository database

Perform the following before upgrading SF Oracle RAC.

Note: The Sfua_Base repository resource group will be removed from the main.cf file. It is not required as a separate service group for SF Oracle RAC 7.0.1.

To prepare to migrate the repository database

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

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

Warning: The Database Flashsnap clone database will not be able to be carried over after upgrading. You must create a new Database Flashsnap clone database after upgrading to 7.0.1.

Upgrading to SF Oracle RAC 7.0.1

This section provides instructions for upgrading to SF Oracle RAC 7.0.1.

- If required, upgrade the operating system.

- Upgrade to SF Oracle RAC 7.0.1.

- Bring the SF Oracle RAC online.

Upgrading the operating system

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

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

   # **mv /etc/llttab /etc/llttab.save**

2   If you are upgrading to Solaris 10 Update 10, apply the following Oracle (Solaris) patches. For instructions, see Oracle documentation.

   ■   For SPARC: `144524-02`

3   Upgrade the operating system on all nodes in the cluster.

   For instructions, see the operating system documentation.

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

   # **mv /etc/llttab.save /etc/llttab**

Upgrading SF Oracle RAC using Install Bundles

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

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

■   Verifies the compatibility of the systems before the upgrade.

■   Stops the SF Oracle RAC processes before the upgrade.

■   Uninstalls SF Oracle RAC.

■   Installs the SF Oracle RAC 7.0 packages on the nodes.

■   Installs the SF Oracle RAC 7.0.1 patches on the nodes.

■   Starts the SF Oracle RAC processes after the upgrade.

■   Displays the location of the log files, summary file, and response file.

To upgrade to SF Oracle RAC 7.0.1 using the installmr program

1   Log in as superuser.

2   Change to the `/infoscale7.0.1` directory.

3   Invoke the `installmr` script with `-base_path` option to upgrade to 7.0.1:

   #./installmr -base_path /infoscale7.0/*dvd1-sol_sparc/sol10_sparc/*

   Note: The actual path varies depending on your operating system.

   ---

4   From the opening Selection Menu, choose: **G** for "Upgrade a Product."

5   Select 1 for **Full upgrade**.

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

The installer verifies the systems for compatibility.

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

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

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

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

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

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

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

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

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 book disk group before the upgrade proceeds. If you want to split the boot disk group to create a backup, answer y.

7   Enter the name of the backup boot disk group when prompted. Press **Enter** to accept the default.

You are prompted to start the split operation.

8   Enter **y** to continue with the split operation.

The split operation can take some time to complete.

Note: Verify the boot device from which the system is set to boot. Make sure that the system is set to start from the boot device with the required version of SF Oracle RAC.

---

9 The installer displays the following question before it stops the product processes, if the cluster is configured in secure mode and the previous product version is lower than 6.2:

```
Do you want to grant read access to everyone? [y,n,q,?]
```

To grant read access to all authenticated users, type **y**.

To grant permissions to specific user group, type **n**.

```
Do you want to provide any usergroups that you would like to\
grant read access?[y,n,q,?]
```

To specify user groups and grant them read access, type **y**.

To grant read access only to root users, type **n**. Then the installer grants read access read access to the root users.

Enter the user group names that you want to grant read access and separate them by spaces. If you want to grant read access to a user group on a specific node, enter ***usergroup@node***. If you want to grant read access to user groups on any cluster node, enter ***usergroup***. If some user groups are not created yet, create the user groups after configuration if needed.

10 Enter **y** to stop the SF Oracle RAC processes.

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

The installer stops the processes and uninstalls SF Oracle RAC. After the uninstallation, the installer installs SF Oracle RAC 7.0.1 and starts 7.0.1 on all the nodes.

If the product is licensed with stale (old) key, the installer prompts users to update the key.

11 Install the language packages and patches if you would like to run SF Oracle RAC in a language other than English.

12 Relink the SF Oracle RAC libraries with Oracle:

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

Complete the remaining tasks to finish the upgrade.

13 On Solaris 10, perform the following steps to set the zone group appropriately after upgrade.

Sync non-global zone with the global:

```
# zoneadm -z <zonename> attach -u
```

Reset the AutoStart attribute to 1:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 1
```

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

Make zone group online:

```
# hagrp -online <zonegroup> -sys <sysname>
```

Bringing the application database online

1 Start all applications that are not managed by VCS. Use native application commands to start the applications.

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

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

■ For Oracle RAC 11g:

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

■ For Oracle RAC 12c:

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

3 Complete other post-upgrade steps.

For instructions, see the chapter *Performing post-upgrade tasks* in *Storage Foundation for Oracle RAC 7.0 Installation and Configuration Guide*.

4 Upgrade application, if required.

For information on Oracle RAC support, see:

http://www.Veritas.com/docs/DOC5081

For instructions, see the chapter *Upgrading application* in 7.0 SF Oracle RAC Installation Guide.

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

Note: If you want to upgrade all application clusters to version 7.0.1, make sure that you upgraded CP server systems that use VCS or SFHA to 7.0.1. Then, upgrade all application clusters to version 7.0.1.

Performing an automated upgrade using response files with Install Bundles

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

To perform automated upgrade using response files

1 Make sure the systems where you want to upgrade meet the upgrade requirements.

2 Copy the response file to one of the systems where you want to upgrade Veritas InfoScale.

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

4 Navigate to the folder that contains the installation program.

5 Start the upgrade from the system to the `/infoscale7.0.1` directory. For example:

```
# ./installmr -responsefile /infoscale7.0.1/response_file
```

Where /infoscale7.0.1/*response_file* is the response file's full path name.

6 Complete the post upgrade task as mentioned in the upgrade method for specific components.

Performing Boot Environment upgrade with Install Bundles on Solaris 11 x86 systems

Table 4-1 Upgrading VCS, SFHA, SFCFSHA, and SF Oracle RAC using BE upgrade

Step	Description
Step 1	Create a new BE on the primary boot disk.
	See "Creating a new Solaris 11 BE on the primary boot disk" on page 53.
Step 2	Upgrade VCS, SFHA, SFCFSHA, and SF Oracle RAC using the installer.
	See "Upgrading VCS, SFHA, SFCFSHA, and SF Oracle RAC using the installer for upgrading BE on Solaris 11" on page 53.
	To upgrade only Solaris
	See the Oracle documentation on Oracle Solaris 11 operating system.
Step 3	Switch the alternate BE to be the new primary.
	See "Completing the upgrade for VCS on BE on Solaris 11" on page 54.
	See "Completing the upgrade for SFHA, SFCFSHA, and SF Oracle RAC on BE on Solaris 11" on page 55.
Step 4	Verify BE upgrade of VCS, SFHA, SFCFSHA, and SF Oracle RAC.
	See "Verifying Solaris 11 BE upgrade" on page 56.

Creating a new Solaris 11 BE on the primary boot disk

At the end of the process, a new BE is created on the primary boot disk by cloning the primary BE.

To create a new BE on the primary boot disk

1   View the list of BE in the primary disk.

```
# beadm list
```

2   Create a new BE in the primary boot disk.

```
# beadm create beName
```

```
# beadm mount beName mountpoint
```

If VVR is configured, it is recommended that *<beName>* should have the value *altroot.5.11* and *<mountpoint>* should have the value /altroot.5.11.

Upgrading VCS, SFHA, SFCFSHA, and SF Oracle RAC using the installer for upgrading BE on Solaris 11

You can use the Veritas product installer to upgrade VCS, SFHA, SFCFSHA, and SF Oracle RAC on a BE.

At the end of the process, the VCS, SFHA, SFCFSHA, or SF Oracle RAC is installed on the alternate BE.

To perform BE upgrade of VCS, SFHA, SFCFSHA, and SF Oracle RAC using the installer

1   Upgrade OS before upgrading the product, if necessary.

2   Access your copy of the software on the network.

3   Run the installer script specifying the root path as the alternate BE:

```
# ./installmr -rootpath /altroot.5.11 -base_path /infoscale7.0/dvd1-sol_x
```

Note: The actual path varies depending on your operating system version.

4   Enter the names of the nodes that you want to upgrade to 7.0.1.

> **Note:** If you are upgrading SFHA, SFCFSHA, or SF Oracle RAC, make sure
> that the installed version of VxFS uses the disk layout version 6 or later. If you
> are on a previous disk layout version, upgrade the version before you proceed
> with the product installation.

The installer displays the list of packages to be installed or upgraded on the
nodes.

5   Press **Return** to continue with the installation.

6   Verify that the version of the Veritas packages on the alternate BE is 7.0.1.

    # **pkg -R /altroot.5.11 list *VRTS***

For example:

- For SFCFSHA:

        # **pkg -R /altroot.5.1list VRTSvxvm**

- For VCS:

- For SF Oracle RAC:

Review the installation logs at `/altroot.5.11/opt/VRTS/install/logs`.

Completing the upgrade for VCS on BE on Solaris 11

At the end of the process:

- The alternate BE is activated.

- The system is booted from the alternate BE.

To complete the BE upgrade

1   Activate the alternate BE.

```
# beadm activate altroot.5.11
```

2   _____

Note: Do not use the `reboot`, `halt`, or `uadmin` commands to restart the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate BE.

_____

```
# shutdown -g0 -y -i6
```

3   After the alternate BE is activated, you can switch BEs. If the root disk is encapsulated, refer to the procedure to switch the BEs manually.

Completing the upgrade for SFHA, SFCFSHA, and SF Oracle RAC on BE on Solaris 11

At the end of the process:

■   The alternate BE is activated.

■   The system is booted from the alternate BE.

To complete the BE upgrade

1   Activate the alternate BE.

```
# beadm activate altroot.5.11
```

2   _____

Note: Do not use the `reboot`, `halt`, or `uadmin` commands to restart the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate BE.

_____

```
# shutdown -g0 -y -i6
```

3   After the alternate BE is activated, you can switch BEs. If the root disk is encapsulated, refer to the procedure to switch the BEs manually.

4   After the upgrade, perform any required post-upgrade tasks such as upgrading the disk group.

5   After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.

Verifying Solaris 11 BE upgrade

To verify that BE upgrade completed successfully

1   Verify that the alternate BE is active.

     # **beadm list**

     If the alternate BE fails to be active, you can revert to the primary BE.

2   Make sure that GAB ports a and h are up.

3   Perform other verification as required to ensure that the new BE is configured
     correctly.

4   In a zone environment, verify the zone configuration.

If you have solaris10 brand zone on your system, you must manually upgrade the
packages inside the solaris10 brand zone with packages from Solaris 10 install
media.

If you have installed VRTSvxfs or VRTSodm packages inside the zones, you need
to manually upgrade these packages inside the zone.

Reverting to the primary BE on a Solaris 11 system

Boot the system to ok prompt.

View the available BEs.

To view the BEs, enter the following:

ok> **boot -L**

Select the option of the original BE to which you need to boot.

To boot to the BE

boot -Z *<path to boot env>*

For example:

```
{0} ok boot -L
Boot device: /virtual-devices@100/channel-devices@200/disk@0:a
File and args: -L
1 Oracle Solaris 11 11/11 SPARC
2 solaris-backup-1
Select environment to boot: [ 1 - 2 ]: 1
```

To boot the selected entry, enter the following:

```
boot [<root-device>] -Z rpool/ROOT/solaris

Program terminated
{0} ok boot -Z rpool/ROOT/solaris
```

Upgrading to 7.0.1 from 7.0

This chapter includes the following topics:

- Performing a full upgrade to 7.0.1 on a cluster

- Upgrading to 7.0.1 on a standalone system

- Upgrading Veritas products using Live Upgrade

- Manually installing packages on Solaris brand non-global zones

- Verifying software versions

Performing a full upgrade to 7.0.1 on a cluster

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

Depending on your cluster's configuration, select one of the following procedures to upgrade to 7.0.1:

- Performing a full upgrade to 7.0.1 on a Veritas Cluster Server

- Performing a full upgrade to 7.0.1 on an SFHA cluster

- Performing a full upgrade to 7.0.1 on an SFCFSHA cluster

- Performing a full upgrade to 7.0.1 on an SF Oracle RAC cluster
  See "Downloading required software to upgrade to 7.0.1 " on page 23.

Performing a full upgrade to 7.0.1 on a Veritas Cluster Server

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

To upgrade VCS

1. Make sure you have downloaded the latest software required for the upgrade.

2. Log in as superuser.

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

   ---

3. If you install VCS on Solaris 10 systems that run non-global zones, make sure that all non-global zones are booted and in the running state on each node before you upgrade the VCS stack in the global zone.

4. Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 7.0.1 rolling patch binaries, change to the directory that contains the installmr script. Start the pre-upgrade check:

   ```
   # ./installmr -precheck sys1 sys2 ... nodeN
   ```

5. Resolve any issues that the precheck finds.

6. On Solaris 10, if zones are configured, you need to set AutoStart attribute to 0 for zone groups.

   Set AutoStart attribute to 0 for zone groups:

   ```
   # haconf -makerw

   # hagrp -modify <zonegroup> AutoStart 0

   # haconf -dump -makero
   ```

7. Start the upgrade:

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

8 After the upgrade, review the log files for any issues.

9 On Solaris 10, perform the following steps to set the zone group appropriately after upgrade.

Sync non-global zone with the global:

```
# zoneadm -z <zonename> attach -u
```

Reset the AutoStart attribute to 1:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 1
```

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

Make zone group online:

```
# hagrp -online <zonegroup> -sys <sysname>
```

Performing a full upgrade to 7.0.1 on an SFHA cluster

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

To perform a full upgrade to 7.0.1 on an SFHA cluster

1 Make sure you have downloaded the latest software required for the upgrade.

2 Log in as superuser.

3 Verify that /opt/VRTS/bin and /opt/VRTSvcs/bin are in your PATH so that you can execute all product commands.

4 If you install SFHA on Solaris 10 systems that run non-global zones, make sure that all non-global zones are booted and in the running state on each node before you upgrade the SFHA stack in the global zone.

5 On each node in the cluster, make the VCS configuration read only:

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

6 Stop VCS.

To stop applications, unmount VxFS file systems and stop VxVM volumes managed by VCS.

```
# hastop -all
```

7 Stop all the applications that are using VxFS files systems and VxVM volumes which are not managed by VCS.

Use application's native commands to stop applications.

8 On each node, enter the following command to check if any Storage Checkpoints are mounted:

```
# df -F vxfs
```

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

```
# umount /checkpoint_name
```

9 On each node, enter the following command to check if any VxFS file systems are mounted.

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

```
# df -F vxfs
```

If any VxFS file systems are present, on each node in the cluster, stop IOs on the file systems, unmount all of the VxFS file systems:

```
# umount /filesystem
```

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the vxrvg stop command to stop each RVG individually:

  ```
  # vxrvg -g diskgroup stop rvg_name
  ```

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

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

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

11 Stop activity to all VxVM volumes that are not managed by VCS.

For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.Use application specific commands to stop the applications.

12 On each node, stop all VxVM volumes by entering the following command for each disk group, which are not managed by VCS:

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

Verify that no volumes remain open:

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

13 Deport all the disk groups which are not managed under VCS.

```
# vxdg deport diskgroup
```

14 If required, apply the OS kernel patches.

See Oracle's documentation for the procedures.

15 On each node, stop the VCS command server:

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

pid_of_CmdServer is the process ID of CmdServer.

16 Check the readiness of the nodes where you plan to upgrade. From the directory that contains the extracted and untarred 7.0.1 rolling patch binaries, change to the directory that contains the `installmr` script. Start the pre-upgrade check:

```
# ./installmr -precheck sys1 sys2 ... nodeN
```

where *sys1* and *sys2* are nodes which are to be upgraded.

Resolve any issue that the precheck finds.

17 On Solaris 10, if zones are configured, you need to set AutoStart attribute to 0 for zone groups.

Set AutoStart attribute to 0 for zone groups:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 0
```

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

18 Start the upgrade.

```
# ./installmr [-rsh]  sys1 sys2 ... nodeN
```

Review the output and follow the instructions to finish the upgrade.

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

```
# hagrp -online service_group -sys nodename
```

20 If necessary, reinstate any missing mount points in the /etc/vfstab file on each node.

21 Import all the diskgroups that are not managed by VCS:

```
# vxdg import diskgroup
```

22 Restart all the volumes by entering the following command for each disk group that are not managed by VCS:

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

23 If you stopped any RVGs in step 10, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

24 Remount all VxFS file systems on all nodes, which are not managed by VCS:

```
# mount -F vxfs blockdevice
        mountpoint
```

25 Remount all Storage Checkpoints on all nodes:

```
# mount -F vxfs -o ckpt=name
        blockdevice
        checkpoint_name
```

26 Start all applications which are using VxFS files systems that are not managed by VCS.

Use application native commands to start the applications.

27 On Solaris 10, perform the following steps to set the zone group appropriately after upgrade.

Sync non-global zone with the global:

```
# zoneadm -z <zonename> attach -u
```

Reset the AutoStart attribute to 1:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 1
```

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

Make zone group online:

```
# hagrp -online <zonegroup> -sys <sysname>
```

Performing a full upgrade to 7.0.1 on an SFCFSHA cluster

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

To perform a full upgrade to 7.0.1 on an SFCFSHA cluster

1 Make sure you have downloaded the latest software required for the upgrade.

2 Log in as superuser.

3 Verify that /opt/VRTS/bin and /opt/VRTSvcs/bin is in your PATH so that you can execute all product commands.

4 If you install SFCFSHA on Solaris 10 systems that run non-global zones, make sure that all the installed non-global zones are bootable.

Run the following command to get the state of non-global zones:

```
# zoneadm list -cv
```

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

```
# haconf -makerw
```

6 Enter the following command to freeze HA service group operations on any node:

```
# hagrp -freeze groupname -persistent
```

7 Make the configuration read-only:

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

8 Stop VCS. To stop applications, unmount VxFS/CFS file systems and stop VxVM or CVM volumes managed under VCS.

```
# hastop -all
```

9 Stop all applications which are using CFS file systems and VxVM volumes not managed by VCS. Use application native commands to stop applications.

10 On each node, enter the following command to check if any Storage Checkpoints are mounted:

```
# df -F vxfs
```

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

```
# cfsumount /checkpoint_name
```

11 On each node, enter the following command to check if any VxFS file systems are mounted:

```
# df -F vxfs
```

- If any VxFS/CFS file systems are present, on each node in the cluster, stop IOs on the file systems, unmount all of the VxFS/CFS file systems:

```
# umount /filesystem
```

Or

```
# cfsumount /filesystem
```

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the vxrvg stop command to stop each RVG individually:

```
# vxrvg -g diskgroup stop rvg_name
```

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

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

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

---

13 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.

14 On each node, stop all VxVM volumes by entering the following command for each disk group:

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

Verify that no volumes remain open:

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

15 Deport all the disk groups which are not managed under VCS:

```
# vxdg deport diskgroup
```

16 On each node, stop the VCS command server:

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

where *pid_of_CmdServer* is the process ID of CmdServer.

17 If required, apply the OS kernel patches.

See Oracle's documentation for the procedures.

18 On Solaris 10, if zones are configured, you need to set AutoStart attribute to 0 for zone groups.

Set AutoStart attribute to 0 for zone groups:

```
# haconf -makerw

# hagrp -modify <zonegroup> AutoStart 0

# haconf -dump -makero
```

19 From the directory that contains the extracted and untarred 7.0.1 rolling patch binaries, change to the directory that contains the installmr script. Start the upgrade.

```
# ./installmr [-rsh] sys1 sys2 ... nodeN
```

Review the output and follow the instructions to finish the upgrade..

20 If necessary, reinstate any missing mount points in the /etc/vfstab file on each node.

21 Make the VCS configuration writable again from any node:

```
# haconf -makerw
```

22 Enter the following command on any node to unfreeze HA service group operations:

```
# hagrp -unfreeze groupname -persistent
```

23 Make the configuration read-only:

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

24 Bring the CVM service group online on each node:

```
# hagrp -online cvm -sys nodename
```

25 Import all the VxVM or CVM diskgroups that are not managed by VCS:

```
# vxdg import diskgroup
```

or

```
$vxdg import -s disgroup
```

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

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

27 If you stopped any RVGs in step 12, restart each RVG:

```
# vxrvg -g diskgroup start rvg_name
```

28 Remount all VxFS/CFS file systems on all nodes:

```
# mount -F fstype
        blockdevice
        mountpoint
```

Or

```
# cfsmount /mountpoint
```

29 Remount all Storage Checkpoints on all nodes:

```
# cfsmount /checkpoint_name
```

Or

```
# mount vxfs -o ckpt=name
        blockdevicemountpoint
```

30 Start all applications which are using VxFS/CFS file systems that are not managed by VCS. Use the application native commands to start the applications.

31 On Solaris 10, perform the following steps to set the zone group appropriately after upgrade.

Sync non-global zone with the global:

```
# zoneadm -z <zonename> attach -u
```

Reset the AutoStart attribute to 1:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 1
```

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

Make zone group online:

```
# hagrp -online <zonegroup> -sys <sysname>
```

Performing a full upgrade to 7.0.1 on an SF Oracle RAC cluster

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

To upgrade to 7.0.1 on an SF Oracle RAC cluster

1 Make sure you have downloaded the latest software required for the upgrade.

2 Log in as superuser.

3 Verify that /opt/VRTSvcs/bin is in your PATH so that you can execute all product commands.

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

```
# haconf -makerw
```

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

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

6 Make the configuration read-only:

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

7   For Oracle RAC 11g and Oracle RAC 12c:

Stop all application resources.

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

  ```
  # hagrp -offline group_name -any
  ```

- If the database instances are not managed by VCS, then run the following on one node:
  For Oracle RAC 12c:

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

  For Oracle RAC 11g:

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

8
- If the Oracle database is managed by VCS, set the AutoStart value to 0 to prevent the database service group from starting automatically when VCS starts:

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

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

  For Oracle RAC 11g:

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

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

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

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

If the applications are under VCS control:

```
# hagrp -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

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

```
# mount -v |grep vxfs
```

```
# fuser -c /mount_point
```

```
# umount /mount_point
```

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

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

12 Stop VCS.

```
# hastop -all
```

13 If required, apply the OS kernel patches.

See *Oracle's* documentation for the procedures.

14 On Solaris 10, if zones are configured, you need to set AutoStart attribute to 0 for zone groups.

Set AutoStart attribute to 0 for zone groups:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 0
```

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

15 From the directory that contains the extracted and untarred 7.0.1 rolling patch binaries, change to the directory that contains the installmr script. Enter:

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

where sys1 and sys2 are nodes which are to be upgraded.

16 After the entire cluster is upgraded, follow the installer instructions to proceed further.

17 If necessary, reinstate any missing mount points in the /etc/vfstab file on each node.

18 Relink the SF Oracle RAC libraries with Oracle:

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

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

```
# haconf -makerw
```

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

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

21 Make the configuration read-only:

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

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

```
# hagrp -online service_group -sys nodename
```

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

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

24 Remount all VxFS file systems on all nodes:

```
# mount /filesystem
```

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

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

26 Bring the Oracle database service group online.

- If the Oracle database is managed by VCS:

    ```
    #  hagrp -online Oracle_group -any
    ```

- If the Oracle database is not managed by VCS:
  For Oracle RAC 11g:

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

    For Oracle RAC 12c:

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

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

```
# haconf -makerw
# hagrp -modify oracle_groupname AutoStart 1
# haconf -dump -makero
```

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

For Oracle RAC 11g:

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

For Oracle RAC 12c:

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

28 On Solaris 10, perform the following steps to set the zone group appropriately after upgrade.

Sync non-global zone with the global:

```
# zoneadm -z <zonename> attach -u
```

Reset the AutoStart attribute to 1:

```
# haconf -makerw
```

```
# hagrp -modify <zonegroup> AutoStart 1
```

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

Make zone group online:

```
# hagrp -online <zonegroup> -sys <sysname>
```

29 Upgrade Oracle RAC.

For information on Oracle RAC support, see:

http://www.Veritas.com/support/en_US/article.DOC5081

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

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

Upgrading to 7.0.1 on a standalone system

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

To upgrade to 7.0.1 on a standalone system

1   Make sure you have downloaded the latest software required for the upgrade.

2   Log in as superuser.

3   Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.

4   If required, apply kernel patches as mentioned in the System requirements.

     For more information, see System requirements in *Veritas™ Veritas InfoScale 7.0.1 Release Notes*.

     See IBM's documentation for the procedures.

5   Enter the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

     ```
     # df -F vxfs
     ```

6   Unmount all Storage Checkpoints and file systems:

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

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

     ■ Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

     ■ Use the `vxrvg stop` command to stop each RVG individually:

        ```
        # vxrvg -g diskgroup stop rvg_name
        ```

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

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

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

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

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

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

Verify that no volumes remain open:

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

10 Copy the patch archive downloaded from the patch central to temporary location, untar the archive and browse to the directory containing the installmr installer script. Enter the installmr script:

```
# ./installmr nodename
```

11 If necessary, reinstate any missing mount points in the /etc/vfstab file.

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

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

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

```
# vxrvg -g diskgroup start rvg_name
```

14 Remount all VxFS file systems and Storage Checkpoints:

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

Upgrading Veritas products using Live Upgrade

This section describes how to upgrade 7.0 to 7.0.1 using Live Upgrade.

Supported live upgrade paths:

- Upgrading Veritas Products without Solaris OS upgrade:
  - Upgrading Solaris 10 Update x 7.0 to Solaris 10 Update x 7.0.1
- Upgrading Veritas Products with Solaris OS upgrade
  - Upgrading Solaris 10 Update x 7.0 to Solaris 10 Update y 7.0.1
  - Upgrading Solaris 11 Update x 7.0 to Solaris 11 Update y 7.0.1

Prerequisites to upgrade to 7.0.1 using Live Upgrade:

- The node should have an alternate boot disk that is identical to the primary boot disk.

- Installation disc for 7.0 and 7.0.1 to be installed on the ABE.

- Installation disc for target OS to be installed on ABE.

- The latest list of required patches is available in the Oracle Solaris Live Upgrade Software:

  Patch Requirements (Doc ID 1004881.1) document in My Oracle Support (https://support.oracle.com/).

- If OS upgrade is involved, then remove the currently installed SUNWluu, SUNWlur and SUNWlucfg packages and install SUNWluu, SUNWlur, SUNWlucfg packages from target OS. Also replace SUNWluzone if zones are involved.

- The vxlustart script takes around 2-3 hours to complete uninterrupted. Veritas recommends to have a network connection that does not time out in the interim.

Upgrading Veritas products using Live Upgrade from 7.0 to 7.0.1 without OS upgrade

This section describes how to upgrade SF, SFHA, SFCFS, SFCFSHA, Sybase ASE CE, or SF for Oracle RAC from 7.0 to 7.0.1 using Live Upgrade where OS upgrade is not involved.

To upgrade your Veritas product using Live Upgrade

1 Ensure that 7.0 is installed and configured on PBE.

   See your Veritas product 7.0 Installation Guide for more information.

2 Run the vxlustart -V command to ensure there are no problems before beginning the Live Upgrade process.

   If the vxlustart -V command reports success, proceed with running the vxlustart command.

   If the vxlustart -V command reports errors, correct the problem, and run the vxlustart -V command again.

Note: This vxlustart -V command does not catch failures that are reported by Solaris Live Upgrade commands.

---

3   Run the `vxlustart` command to start the Live Upgrade for your Veritas product:

   # **cd /infoscale7.0.1/*dvd1-sol_sparc/sol10_sparc*/scripts**

   # **./vxlustart -v -u *target_os_version* -U -d *disk_name***

   Note: The actual path varies depending on your operating system.

   ---

4   Run the `installmr` command to upgrade your Veritas product:

   # **cd /infoscale7.0.1/*dvd1-sol_sparc/sol10_sparc***

   # **./installmr -rootpath /*altroot_path***

   Note: The actual path varies depending on your operating system.

   ---

5   Run the `vxlufinish` command to complete the Live Upgrade:

   ■   If the primary root disk is not encapsulated, run the following command:

      # **cd /infoscale7.0.1/*dvd1-sol_sparc/sol10_sparc*/scripts**

      # **./vxlufinish -u *target_os_version***

      Note: The actual path varies depending on your operating system.

   ■   If the primary root disk is encapsulated by VxVM, run the following command:

      # **cd /infoscale7.0.1/*dvd1-sol_sparc/sol10_sparc*/scripts**

      # **./vxlufinish -u *target_os_version* -g *diskgroup***

      Note: The actual path varies depending on your operating system.

      ---

6 If the Oracle database is managed by VCS, modify the VCS configuration file on the alternate root disk (`/altroot.5.10/etc/VRTSvcs/conf/config/main.cf`) to set the AutoStart value to 0. This prevents the database service group from starting automatically when VCS starts:

For SFRAC:

If the database is not managed by VCS, change the management policy for the database to manual on the primary boot disk:

For Oracle RAC 12c:

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

For Oracle RAC 11g:

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

7 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

Note: DO NOT use the reboot, halt, or uadmin commands to reboot the system. Use either the init or the shutdown commands to enable the system to boot using the alternate boot environment.

```
# shutdown -g0 -y -i6
```

8 In case of SFRAC, refer to the "Performing post-upgrade Tasks" section to relink Oracle RAC libraries with SF Oracle RAC from 7.0 Installation and Configuration guide.

9 Start the database group on all nodes:

For SFRAC:

10 If the Oracle database is managed by VCS, modify the VCS configuration file
(`/etc/VRTSvcs/conf/config/main.cf`) to set the AutoStart value to 1.

For SFRAC:

If the database is not managed by VCS, change the management policy for
the database to automatic on the primary boot disk:

For Oracle RAC 12c:

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

For Oracle RAC 10g and Oracle RAC 11g:

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

11 If you are on an unsupported version of Oracle RAC, upgrade Oracle RAC.

For instructions, see the chapter *Upgrading Oracle RAC* in this document.

12 Verify that the alternate boot environment is active.

lustatus

13 In a cluster environment, make sure that all the GAB ports are up. Note different
ports appear for different products.

gabconfig -a

Upgrading Veritas products using Live Upgrade from 7.0 to 7.0.1 with OS upgrade

This section describes how to upgrade SF, SFHA, SFCFS, SFCFSHA, Sybase
ASE CE, or SF for Oracle RAC from 7.0 to 7.0.1 using Live Upgrade where OS
upgrade is involved..

To upgrade your Veritas product using Live Upgrade

1   Ensure that 7.0 is installed and configured on PBE.

    See your Veritas product 7.0 Installation Guide for more information.

2   Run the `vxlustart -V` command to ensure there are no problems before
    beginning the Live Upgrade process.

    If the `vxlustart -V` command reports success, proceed with running the
    `vxlustart` command.

    If the `vxlustart -V` command reports errors, correct the problem, and run the
    `vxlustart -V` command again.

    Note: This `vxlustart -V` command does not catch failures that are reported
    by Solaris Live Upgrade commands.
    ---

3   Run the `vxlustart` command to start the Live Upgrade for your Veritas product:

    # **cd /infoscale7.0.1/*dvd1-sol_sparc/sol10_sparc*/scripts**

    # **./vxlustart -v -u *target_os_version* -s *osimage_path* -d *disk_name***

    Note: The actual path varies depending on your operating system.
    ---

4   Run the `installmr` command to upgrade your Veritas product:

    # **cd /infoscale7.0.1/*dvd1-sol_sparc/sol10_sparc***

    # **./installmr -rootpath /*altroot_path***

    Note: The actual path varies depending on your operating system.
    ---

5   Run the `vxlufinish` command to complete the Live Upgrade:

    ■ If the primary root disk is not encapsulated, run the following command:

       # **cd /infoscale7.0.1/*dvd1-sol_sparc/sol10_sparc*/scripts**

       # **./vxlufinish -u *target_os_version***

       Note: The actual path varies depending on your operating system.

- If the primary root disk is encapsulated by VxVM, run the following command:

  # **cd /infoscale7.0.1/*dvd1-sol_sparc/sol10_sparc*/scripts**

  # **./vxlufinish -u *target_os_version* -g *diskgroup***

  Note: The actual path varies depending on your operating system.

  ---

6 If the Oracle database is managed by VCS, modify the VCS configuration file on the alternate root disk (`/altroot.5.10/etc/VRTSvcs/conf/config/main.cf`) to set the AutoStart value to 0. This prevents the database service group from starting automatically when VCS starts:

For SFRAC:

```
group oradb_grp (
        SystemList = { galaxy = 0, nebula = 1 }
        AutoStart = 0
        AutoFailOver = 0
        Parallel = 1
        AutoStartList = { galaxy, nebula }
        )
.
.
```

If the database is not managed by VCS, change the management policy for the database to manual on the primary boot disk:

For Oracle RAC 12c:

$ **srvctl modify database -db db_name -policy manual**

For Oracle RAC 10g and Oracle RAC 11g:

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

7 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

Note: DO NOT use the reboot, halt, or uadmin commands to reboot the system. Use either the init or the shutdown commands to enable the system to boot using the alternate boot environment.

```
# shutdown -g0 -y -i6
```

8 In case of SFRAC, refer to the "Performing post-upgrade Tasks" section to relink Oracle RAC libraries with SF Oracle RAC from 7.0 Installation and Configuration guide.

9 Start the database group on all nodes:

For SFRAC:

10 If the Oracle database is managed by VCS, modify the VCS configuration file (`/etc/VRTSvcs/conf/config/main.cf`) to set the AutoStart value to 1.

For SFRAC:

If the database is not managed by VCS, change the management policy for the database to automatic on the primary boot disk:

For Oracle RAC 12c:

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

For Oracle RAC 10g and Oracle RAC 11g:

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

11 Verify that the alternate boot environment is active.

```
# lustatus
```

12 In a cluster environment, make sure that all the GAB ports are up. Note that different ports appear for different products.

```
# gabconfig -a
```

Upgrading SF Oracle RAC using Live Upgrade on Solaris 11

- Upgrading SF Oracle RAC only using Live Upgrade

- Creating a new boot environment on the alternate boot disk

- Upgrading SF Oracle RAC using the installer for a Live Upgrade

- Completing the Live Upgrade

- Verifying Live Upgrade of SF Oracle RAC 7.0.1

- Upgrading Solaris only using Live Upgrade

Before you upgrade SF Oracle RAC using Solaris Live Upgrade

Before you upgrade, perform the following procedure.

To prepare for the Live Upgrade

1 Make sure that the SF Oracle RAC installation media and the operating system installation images are available and on hand.

2 On the nodes to be upgraded, select an alternate boot disk that is at least the same size as the root partition of the primary boot disk.

   If the primary boot disk is mirrored, you need to break off the mirror for the alternate boot disk.

3 If the nodes to be upgraded are on Solaris 9, on the primary boot disk, patch the operating system for Live Upgrade. Patch 137477-01 is required. Verify that this patch is installed.

4 The version of the Live Upgrade packages must match the version of the operating system to which you want to upgrade on the alternate boot disk. If you are upgrading the Solaris operating system, do the following steps:

   - Remove the installed Live Upgrade packages for the current operating system version:
     All Solaris versions: SUNWluu, SUNWlur packages.
     Solaris 10 update 7 or later also requires: SUNWlucfg package.

   - From the new Solaris installation image, install the new versions of the following Live Upgrade packages:
     All Solaris versions: SUNWluu, SUNWlur, and SUNWlucfg packages.

   Solaris installation media comes with a script for this purpose named liveupgrade20. Find the script at
   /*cdrom*/*solaris_release*/Tools/Installers/liveupgrade20. If scripting, you can use:

   ```
   # /cdrom/solaris_release/Tools/Installers/liveupgrade20 \
   -nodisplay -noconsole
   ```

After you install the packages, install the latest Live Upgrade patch. For more information on required packages and patches, see the Oracle Metalink document: 1004881.1 or visit the following site:

https://support.oracle.com/

5 Veritas provides the `vxlustart` script that runs a series of commands to create the alternate boot disk for the upgrade.

To preview the commands, specify the `vxlustart` script with the `-V` option.

Veritas recommends that you preview the commands to ensure there are no problems before beginning the Live Upgrade process.

Note: The `vxlustart` script and the `vxlufinish` script are located in the `/infoscale7.0.1/`*dvd1-sol_sparc/sol10_sparc*`/scripts` directory.

The actual path varies depending on your operating system.

```
# cd /infoscale7.0.1/dvd1-sol_sparc/sol10_sparc/scripts

# ./vxlustart -V -u targetos_version -s osimage_path -d diskname
```

Note: The actual path varies depending on your operating system.

For usages of the vxlustart option,

For example, to preview the commands to upgrade only the Veritas product:

```
# ./vxlustart -V -u 5.10 -U -d disk_name
```

For example, to preview the commands for an upgrade to Solaris 10 update 6:

```
# ./vxlustart -V -u 5.10 -s /mnt/Solaris_10u6 -d c0t1d0s0
```

Note: This command prompts you to compare the patches that are installed on the image with the patches installed on the primary boot disk. If any patches are missing from the new operating system's image, note the patch numbers. To ensure the alternate boot disk is the same as the primary boot disk, you will need to install these patches on the alternate boot disk.

---

6 If the specified image is missing patches that are installed on the primary boot disk, note the patch numbers. To ensure that the alternate boot disk is the same as the primary boot disk, you need to install any missing patches on the alternate boot disk.

Upgrading the operating system and SF Oracle RAC using Live Upgrade

Perform the following steps to upgrade both the operating system and SF Oracle RAC using Live Upgrade.

To upgrade the operating system and SF Oracle RAC using Live Upgrade

1   Prepare to upgrade using Solaris Live Upgrade.

     See "Before you upgrade SF Oracle RAC using Solaris Live Upgrade" on page 83.

2   Create a new boot environment on the alternate boot disk.

     See "Creating a new boot environment on the alternate boot disk" on page 87.

3   Upgrade SF Oracle RAC using the installer or manually.

     See "Upgrading SF Oracle RAC using the installer for a Live Upgrade" on page 91.

4   Complete the Live Upgrade.

     See "Completing the Live Upgrade" on page 92.

5   Verify Live Upgrade of SF Oracle RAC.

     See "Verifying Live Upgrade of SF Oracle RAC 7.0.1" on page 96.

Upgrading SF Oracle RAC only using Live Upgrade

Perform the following steps to upgrade only SF Oracle RAC using Live Upgrade.

To upgrade only SF Oracle RAC using Live Upgrade

1   Prepare to upgrade using Solaris Live Upgrade.

     See "Before you upgrade SF Oracle RAC using Solaris Live Upgrade" on page 83.

2   Create a new boot environment on the alternate boot disk.

     See "Creating a new boot environment on the alternate boot disk" on page 87.

3   Upgrade SF Oracle RAC using the installer or manually.

     See "Upgrading SF Oracle RAC using the installer for a Live Upgrade" on page 91.

4   Complete the Live Upgrade.

     See "Completing the Live Upgrade" on page 92.

5   Verify Live Upgrade of SF Oracle RAC.

     See "Verifying Live Upgrade of SF Oracle RAC 7.0.1" on page 96.

Creating a new boot environment on the alternate boot disk

Run the `vxlustart` command on each node in the cluster to create a new boot environment on the alternate boot disk.

For usages of the vxlustart option,

If the `-U` option is specified, you can omit the `-s` option. The operating system is cloned from the primary boot disk.

Veritas recommends that you preview the commands with `-V` option to ensure there are no problems before beginning the Live Upgrade process. The `vxlustart` script is located on the distribution media, in the scripts directory.

For example, to preview the commands to upgrade only the Veritas product:

```
# ./vxlustart -V -u 5.10 -U -d disk_name
```

For example, to preview the commands for an upgrade to Solaris 10 update 6:

```
# ./vxlustart -V -u 5.10 -s /mnt/Solaris_10u6 -d c0t1d0s2
```

In the procedure examples, the primary or current boot environment resides on *Disk0 (c0t0d0s2)* and the alternate or inactive boot environment resides on *Disk1 (c0t1d0s2)*.

Note: This step can take several hours to complete. Do not interrupt the session as it may leave the boot environment unstable.

At the end of the process:

■ The Solaris operating system on the alternate boot disk is upgraded, if you have chosen to upgrade the operating system.

■ A new boot environment is created on the alternate boot disk by cloning the primary boot environment.

To create a new boot environment on the alternate boot disk

Perform the steps in this procedure on each node in the cluster.

1   Navigate to the install media for the Veritas products:

```
# cd /tmp/sfha7.0.1/scripts
```

2   View the list of VxVM disks on which you want to create the new boot environment.

```
# vxdisk list
```

3   Before you upgrade, make sure that you exclude the file system mount points on a shared storage that applications use from getting copied to the new boot environment. To prevent these shared mount points from being copied to the new boot environment, create a temporary file containing the file system mountpoints that need to be excluded.

```
# cat /var/tmp/file_list
- /ora_mnt
- /sap_mnt
```

where `/var/tmp/file_list` is a temporary file that contains the list of mount points to be excluded from the new boot environment. The items in the file list are preceded either by a '+' or '-' symbol.

The '+' symbol indicates that the mount point is included in the new boot environment.

The '-' symbol indicates that the mount point is excluded from the new boot.

Apart from file system mount points, you may choose to include or exclude other files.

If you have non-global zone in running state in the current boot environment and zone root path is on a VxVM, create another volume of same or more size for each zone root in alternate boot environment path using vxvm commands.

4 Run one of the following commands to create the alternate boot environment:

For example:

To upgrade the operating system:

```
# ./vxlustart -v -u 5.10 -s /mnt/sol10u9 -d
c0t1d0s2 -z /var/tmp/file_list
```

where `/mnt/sol10u9` is the path to the operating system image that contains the `.cdtoc` file.

To clone the operating system of current boot environment:

```
# ./vxlustart -v -u 5.10 -U -d c0t1d0s2 -z /var/tmp/file_list
```

If you have non-global zone with zone root path on VxVM, then to upgrade the OS:

```
# ./vxlustart -v -u 5.10 -U -d c0t1d0s2 -z
/var/tmp/file_list -w /zone1-rootpath:/dev/vx/dsk/rootpathdg_alt/
rootpathvol_alt:vxfs
```

Where `zone1-rootpath` is root path of zone in present boot environment.

5 Update the permissions, user name, and group name of the mount points (created on the ABE) to match that of the existing directories on the primary boot environment.

6 If zone root path is on VxVM, update the `/altroot.5.10/etc/VRTSvcs/conf/config/main.cf` file with new block device created in step 3 for all zones to reflect the ABE zone root paths.

7 Run one of the following commands to perform the upgrade:

To upgrade the operating system, by itself or together with upgrading the Veritas products:

```
# ./vxlustart -v -u targetos_version \
-s osimage_path -d disk_name
```

where *targetos_version* is the version of the operating system

osimage_path is the full path to the operating system image

disk_name is the name of the disk as displayed in the output of step 2.

To upgrade the Veritas product only:

```
# ./vxlustart -v -u 5.10 -U -d disk_name
```

The options to the vxlustart command are listed in the preupgrade section.

For example, to upgrade to Solaris 10 update 6:

```
# ./vxlustart -v -u 5.10 -s /mnt/Solaris_10u6
```

8 Create the mount points manually on the alternate boot environment as follows:

```
# for i in `cat /var/tmp/file_list` ; \
do mkdir -p /altroot.5.10/$i; done
```

9 Update the permissions, user name, and group name of the mount points (created on the ABE) to match that of the existing directories on the primary boot environment.

10 Review the output of df commands and note the new mount points. If the system is rebooted before completion of the upgrade or if the mounts become unmounted, you may need to remount the disks.

If you need to remount, run the command:

```
# vxlustart -r -u targetos_version -d disk_name
```

11 After the alternate boot disk is created and mounted on */altroot.5.10*, install any operating system patches or packages on the alternate boot disk that are required for the Veritas product installation:

```
# pkgadd -R /altroot.5.10 -d pkg_dir
```

Upgrading SF Oracle RAC using the installer for a Live Upgrade

You can use the Veritas product installer to upgrade SF Oracle RAC as part of the Live Upgrade.

On a node in the cluster, run the installer on the alternate boot disk to upgrade SF Oracle RAC on all the nodes in the cluster. The program uninstalls the existing version of SF Oracle RAC on the alternate boot disk during the process.

At the end of the process the following occurs:

■ SF Oracle RAC 7.0.1 is installed on the alternate boot disk.

To perform Live Upgrade of SF Oracle RAC 7.0.1 using the installer

1 Insert the product disc with SF Oracle RAC 7.0.1 or access your copy of the software on the network.

2 Run the installer script specifying the root path as the alternate boot disk:

```
# ./installmr -upgrade -rootpath /altroot.5.10 -base_path /infoscale7.0/d
```

Note: The actual path varies depending on your operating system version.

3 Enter the names of the nodes that you want to upgrade to SF Oracle RAC 7.0.1.

Note: Make sure that the installed version of VxFS uses the disk layout version 6 or later. If you are on a previous disk layout version, upgrade the version before you proceed with the SF Oracle RAC 7.0.1 installation.

The installer displays the list of packages to be installed or upgraded on the nodes.

4 Press **Return** to continue with the installation.

During Live Upgrade, if the OS of the alternate boot disk is upgraded, the installer will not update the VCS configurations for Oracle, Netlsnr, and Sybase resources. If cluster configurations include these resources, you will be prompted to run a list of commands to manually update the configurations after the cluster restarts from the alternate boot disks.

5 Verify that the version of the Veritas packages on the alternate boot disk is 7.0.1.

```
# pkginfo -R /altroot.5.10 -l VRTSpkgname
```

For example:

```
# pkginfo -R /altroot.5.10 -l VRTSdbac
```

Review the installation logs at `/altroot.5.10/opt/VRTS/install/logs`.

Completing the Live Upgrade

At the end of the process:

- If the original primary boot disk was encapsulated, the alternate boot disk is encapsulated.

- The alternate boot environment is activated.

- The system is booted from the alternate boot disk.

When completing the Live Upgrade process, take the following limitations into consideration for Solaris 10 Update 10:

- In a shared disk group environment, extra CFS mount entries are ignored when the `vxlustart` command is run, as they are included in `/etc/vfstab`. The entries must be manually removed before booting from the alternate boot environment.

- On Sparc, Live Upgrade from Solaris 9 to Solaris 10 Update 10 may fail using the `lucreate` command.

See the *Veritas™ Storage Foundation for Oracle RAC 7.0 Release notes* for more details.

To complete the Live Upgrade

1   Complete the Live Upgrade process using one of the following commands:

    If the primary root disk is not encapsulated, run the following command:

    ```
    # ./vxlufinish -u target_os_version
    Live Upgrade finish on the Solaris release <5.10>
    ```

    If the primary root disk is encapsulated by VxVM, run the following command:

    ```
    # ./vxlufinish -u target_os_version -g diskgroup
    Live Upgrade finish on the Solaris release <5.10>
    ```

    The Live Upgrade process encapsulates the alternate root disk if the primary root disk was encapsulated.

2   If the system crashes or reboots before Live Upgrade completes successfully, you may remount the alternate disk using the following command:

    ```
    # ./vxlustart -r -u target_os_version
    ```

    Then, rerun the `vxlufinish` command:

    ```
    # ./vxlufinish -u target_os_version
    ```

3   If the Oracle database is managed by VCS, modify the VCS configuration file on the alternate root disk (`/altroot.5.10/etc/VRTSvcs/conf/config/main.cf`) to set the AutoStart value to 0. This prevents the database service group from starting automatically when VCS starts:

    ```
    group oradb_grp (
            SystemList = { sys1 = 0, sys2 = 1 }
            AutoStart = 0
            AutoFailOver = 0
            Parallel = 1
            AutoStartList = { sys1, sys2 }
            )
    .
    .
    ```

    If the database is not managed by VCS, change the management policy for the database to manual on the primary boot disk:

    ```
    $ srvctl modify database -d db-name -y manual
    ```

4   Perform the following steps on the primary boot environment:

- Stop Oracle Clusterware on each node in the cluster:

  # *clus_home*/bin/crsctl crs stop

  where *clus_home* is the path of the Oracle Clusterware/Grid Infrastructure home directory

- Stop the applications using native application commands.

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

  # **fuser -cu mount-point**

- Take offline all VCS groups that contain CFSMount and CVMVolDg:

  # **hagrp -offline group -sys *sys1***
  # **hagrp -offline group -sys *sys2***

- Unmount the VxFS file systems:

  # **mount -v |grep vxfs**
  # **fuser -c */mount_point***
  # **umount */mount_point***

- Deport CVM disk groups:

  # **vxdg deport *diskgroup_name***

- Make sure that no disk groups are imported:

  # **vxdg list**
  NAME STATE ID

5 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

Note: DO NOT use the reboot, halt, or uadmin commands to reboot the system. Use either the init or the shutdown commands to enable the system to boot using the alternate boot environment.

---

shutdown -g0 -y -i6

6   Relink the SF Oracle RAC libraries with the Oracle RAC libraries:

See *Veritas™ Storage Foundation for Oracle RAC 7.0 Installation and Configuration Guide*.

7   Start the database group on all nodes:

```
# hagrp -online oradb_grpname -any
```

8   If the application database is managed by VCS, modify the VCS configuration file (`/etc/VRTSvcs/conf/config/main.cf`) to set the AutoStart value to 1.

```
group oradb_grp (
        SystemList = { sys1 = 0, sys2 = 1 }
        AutoStart = 1
        AutoFailOver = 0
        Parallel = 1
        AutoStartList = { sys1, sys2 }
        )
.
.
```

If the database is not managed by VCS, change the management policy for the database to automatic on the primary boot disk:

```
$ srvctl modify database -d db-name -y AUTOMATIC
```

9   Complete the post-upgrade tasks.

See *Veritas™ Storage Foundation for Oracle RAC 7.0 Installation and Configuration Guide*.

10   If you are on an unsupported version of Oracle RAC, upgrade Oracle RAC.

For instructions, see the chapter *Upgrading Oracle RAC* in this document.

Note: 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.

Verifying Live Upgrade of SF Oracle RAC 7.0.1

To ensure that Live Upgrade has completed successfully, verify that all the nodes have booted from the alternate boot environment and joined the cluster.

To verify that Live Upgrade completed successfully

1 Verify that the alternate boot environment is active.

    ```
    # lustatus
    ```

    If the alternate boot environment is not active, you can revert to the primary boot environment.

    See *Veritas™ Storage Foundation for Oracle RAC 7.0 Installation and Configuration Guide*.

2 In a cluster environment, make sure that all the GAB ports are up. Note different ports appear for different products.

    ```
    # gabconfig -a
    Port a gen   d77c08 membership 0123
    Port b gen   d77c0a membership 0123
    Port d gen   d77c0c membership 0123
    Port f gen   d77c2d membership 0123
    Port h gen   d77c3d membership 0123
    Port o gen   d77c0b membership 0123
    Port u gen   d77c2f membership 0123
    Port v gen   d77c28 membership 0123
    Port w gen   d77c2a membership 0123
    Port y gen   d77c26 membership 0123
    ```

3 Perform other verification as required to ensure that the new boot environment is configured correctly. The non-global zones must be brought to configured state and then attached with -U option so that packages are upgraded inside the non-global zone also.

    For example, verify the version in the /etc/release file and verify the VRTSdbac version.

Upgrading Solaris only using Live Upgrade

Perform the following steps to upgrade only Solaris using Live Upgrade.

To upgrade only Solaris using Live Upgrade

1 Prepare to upgrade using Solaris Live Upgrade.

   See "Before you upgrade SF Oracle RAC using Solaris Live Upgrade"
   on page 83.

2 Create a new boot environment on the alternate boot disk.

   See "Creating a new boot environment on the alternate boot disk" on page 87.

3 Complete the Live Upgrade.

   See "Completing the Live Upgrade" on page 92.

4 Verify Live Upgrade of SF Oracle RAC.

   See "Verifying Live Upgrade of SF Oracle RAC 7.0.1" on page 96.

Manually installing packages on Solaris brand non-global zones

With Oracle Solaris 11, you must manually install Veritas product 7.0.1 packages inside non-global zones. The native non-global zones are called Solaris brand zones.

To install packages manually on Solaris brand non-global zones:

1 On the global zone, ensure that the SMF service
   `svc:/application/pkg/system-repository:default` is online:

   ```
   global>#  svcs svc:/application/pkg/system-repository
   ```

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

3 Copy the VRTSpkgs.p5p package from the `pkgs` directory from the 7.0
   installation media to the non-global zone (for example at `/tmp/install`
   directory).

4 Copy the VRTSpatches.p5p package from the `patches` directory from the
   Veritas product 7.0.1 installation media to the non-global zone (for example at
   `/tmp/install` directory).

5 Disable the publishers that are not reachable, as the package install may fail if any of the recently added repositories are unreachable. For system publishers added from the global zone that are not reachable inside the non-global zone, disable them from global zone and reboot the zone.

```
local># pkg set-publisher --disable <publisher_name>
```

where *publisher_name* is the name of the publisher.

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

```
local># pkg set-publisher -g /tmp/install/VRTSpkgs.p5p Veritas
local># pkg set-publisher -g /tmp/install/VRTSpatches.p5p Veritas
```

7 Install the required packages as per the product you have installed on the global zone.

```
local># pkg install --accept VRTSperl VRTSvlic VRTSvxfs VRTSvcs \
VRTSvcsag VRTSvcsea VRTSodm
```

8 Verify that required packages are installed.

```
local># pkg list | grep VRTS
```

9 Remove the publisher on the non-global zone.

```
local># pkg unset-publisher Veritas
```

10 Enable the publishers that were disabled earlier.

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

11 Repeat steps 2 through 9 on each non-global zone.

Verifying software versions

To verify the version of the software, enter the following command:

```
# installer -version
```

Rolling back Veritas InfoScale

This chapter includes the following topics:

- About rolling back Veritas InfoScale 7.0.1

- Rolling back using the uninstallmr script on Solaris 10

- Rolling back to previous boot environment on Solaris 11

- Rolling back manually

About rolling back Veritas InfoScale 7.0.1

This section describes how to roll back either by using the `uninstallmr` script or manually.

The `uninstallmr` script uninstalls all the patches associated with packages installed, and starts the processes.

The uninstallmr script uninstalls all the 7.0.1 patches. A scenario wherein the product is upgraded from 7.0 to 7.0.1, after you run the uninstallmr script, all the 7.0.1 patches are uninstalled while the 7.0 packages are retained.

Note: The uninstallmr script is available only on Solaris 10 and not on Solaris 11.

Note: On Solaris 10, the products can only roll back to 7.0.

Rolling back using the uninstallmr script on Solaris 10

Use the following procedure to roll back from any Veritas product to the previous version using the `uninstallmr` script.

Note: If any of the systems that you plan to roll back have encapsulated boot disks, you must reboot them after rollback.

To roll back

- For SFRAC:

1 On each node, take the Oracle resources in the VCS configuration file (main.cf) offline.

   ```
   # hagrp -offline oracle_group -sys node_name
   ```

   If the database is not managed by VCS, stop the Oracle database as follows:

   For Oracle RAC 12c:

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

   For Oracle RAC 10g and Oracle RAC 11g:

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

2 If CRS is not under VCS Control, then enter the following command on each node of the cluster to stop CRS.

   - For 11gR2 and later verions:

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

   - For 10gR2:

     ```
     # CRS_HOME/bin/crsctl stop crs
     ```

3 Stop the applications that use CVM or CFS that are not under VCS control.

   - Using native application commands, stop the applications that use CVM or CFS on all nodes.

   - Verify that no processes use the CFS mount point:

     ```
     # fuser -c /mount_point
     ```

4 Unmount CFS file systems that are not under VCS control.

- Determine the file systems that need to be unmounted by checking the output of mount command.

  ```
  # mount -v | grep vxfs | grep cluster
  ```

- Unmount each file system that is not controlled by VCS on each node:

  ```
  # umount /mount_point
  ```

5 Stop VCS to take the service groups on all nodes offline

On any node execute following command to stop VCS:

```
# hastop -all
```

6 Stopping the applications that use VxVM or VxFS that are not under VCS control

- Using native application commands, stop the applications that use VxVM or VxFS.

- Verify that no processes use the VxFS mount point:

  ```
  # fuser -c /mount_point
  ```

7 Unmounting VxFS file systems that are not under VCS control.

- Determine the file systems that need to be unmounted by checking the output of mount command.

  ```
  # mount -v | grep vxfs
  ```

- Unmount each file system that is not controlled by VCS on each node:

  ```
  # umount /mount_point
  ```

8 For Solaris 10, on nodes that run non-global zones managed by VCS, make sure the non-global zones are in the running state. If there are non-global zones managed by VCS, but not in the running state, boot those non-global zones.

9 Run the `uninstallmr` command, type:

```
# ./uninstallmr node A
        node B
        node C...
```

10 If you performed a roll back on a system that has an encapsualted boot disk, you must reboot the system. After reboot, you may need to run `hagrp -list Frozen=1` to get the frozen SG list . Then run `hagrp -unfreeze <group> -persistent` to unfreeze all the frozen SGs manually.

Rolling back to previous boot environment on Solaris 11

On Solaris 11, Veritas InfoScale 7.0.1 contains only packages, so rolling back to versions early than 7.0.1 is not supported. You can only uninstall the entire SFHA stack with the `uninstallmr` script. Instead, if you have already created a boot environment for installing 7.0.1, you can roll back to the previous boot environment available before installing 7.0.1.

To roll back to previous boot environment

1 Activate the boot environment available before installing 7.0.1:

```
# beadm activate bename
```

2 Reboot the node so that new be is active now:

```
# reboot
```

3 You may optionally destroy the boot environment on which 7.0.1 is installed:

```
# beadm destroy bename
```

For example,

```
# beadm destroy pre_sfha_7.0.1
```

Rolling back manually

Use one of the following procedures to roll back to 7.0 manually.

■ Rolling back Storage Foundation or Storage Foundation and High Availability manually

- Rolling back Storage Foundation Cluster File System High Availability manually

- Rolling back SF for Oracle RAC manually

- Rolling back Veritas Cluster Server manually

Note: You must reboot systems when you roll back manually at the end of the roll back procedure.

Rolling back Storage Foundation or Storage Foundation and High Availability manually

Use the following procedure to roll back to 7.0 manually.

To roll back SF or SFHA

1 Log in as superuser.

2 Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.

3 Unmount all Storage Checkpoints and file systems:

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

4 Check if the root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

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

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

Note: Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the rootdg disk group in addition to the root disk for vxunroot to succeed.

  # `/etc/vx/bin/vxunroot`

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

5 Enter the following command to check if any VxFS file systems are mounted:

  # `df -F vxfs`

  If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

  # `umount /filesystem`

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

  - Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

  - Use the vxrvg stop command to stop each RVG individually:

    # `vxrvg -g diskgroup stop rvg_name`

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

    # `vxrlink -g diskgroup status rlink_name`

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

7 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.

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

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

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

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

9   Stop VCS along with all its resources. Then, stop the remaining resources manually:

```
# svcadm disable -t vcs
```

10   If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

```
# rm /etc/vxfenmode
```

11   Unmount /dev/odm:

```
# umount /dev/odm
```

12   Unload the ODM module:

```
# svcadm disable -t odm
# modinfo | grep odm
# modunload -i odm_mod_id
```

13   Unload the cluster fencing (vxfen) module:

```
# svcadm disable -t vxfen
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

14   Stop GAB and LLT in the following order:

```
# svcadm disable -t gab
# svcadm disable -t llt
```

15   Remove the SF 7.0.1 patches.

- Get the list of 7.0.1 patches, type:

  ```
  # ./installmr -listpatches
  ```

- Remove each patch from the patch list. For example, on Solaris 10:

```
# patchrm 143287-07
```

Rolling back Storage Foundation Cluster File System High Availability manually

Use the following procedure to roll back to 7.0 manually.

To roll back SFCFSHA manually

1   Log in as superuser.

2   Verify that `/opt/VRTS/bin` is in your `PATH` so you can execute all product commands.

3   Unmount all Storage Checkpoints and file systems:

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

4   Check if the root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

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

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

  Note: Do not remove the plexes on the root disk that correspond to the original disk partitions.

- Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices. There must be at least one other disk in the rootdg disk group in addition to the root disk for vxunroot to succeed.

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

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

5 Enter the following command to check if any VxFS file systems are mounted:

```
# df -F vxfs
```

If any VxFS file systems are present, unmount all of the VxFS file systems that are not under VCS control::

```
# umount /filesystem
```

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

- Stop all applications that are involved in replication. For example, if a data volume contains a file system, unmount it.

- Use the vxrvg stop command to stop each RVG individually:

  ```
  # vxrvg -g diskgroup stop rvg_name
  ```

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

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

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

7 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.

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

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

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

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

9 Stop VCS along with all its resources. Then, stop the remaining resources manually:

```
# svcadm disable -t vcs
```

10 If cluster fencing was originally configured in enabled mode, type the following on all the nodes:

```
# rm /etc/vxfenmode
```

11 Unmount /dev/odm:

```
# umount /dev/odm
```

12 Unload the ODM module:

```
# svcadm disable -t odm
# modinfo | grep odm
# modunload -i odm_mod_id
```

13 Unload the cluster fencing (vxfen) module:

```
# svcadm disable -t vxfen
# modinfo | grep vxfen
# modunload -i vxfen_mod_id
```

14 Stop GAB and LLT in the following order:

```
# svcadm disable -t gab
# svcadm disable -t llt
```

15 Remove the SFCFSHA 7.0.1 patches.

- Get the list of 7.0.1 patches, type:

  ```
  # ./installmr -listpatches
  ```

- Remove each patch from the patch list. For example, on Solaris 10:

  ```
  # patchrm 143287-07
  ```

Rolling back SF for Oracle RAC manually

Use the following procedure to roll back to 7.0 manually.

To roll back SF for Oracle RAC manually

1 On each node, take the Oracle resources in the VCS configuration file (main.cf) offline.

hagrp -offline oracle_group -sys *node_name*

If the database is not managed by VCS, stop the Oracle database as follows:

For Oracle RAC 12c:

$ **srvctl stop database -db *db_name***

For Oracle RAC 11g and Oracle RAC 10g:

$ **srvctl stop database -d *db_name***

2 If CRS is not under VCS Control, then enter the following command on each node of the cluster to stop CRS.

- For 10gR2 or 11gR1:

  # ***CRS_HOME*/bin/crsctl stop crs**

- For 11gR2 and later versions:

  # ***GRID_HOME*/bin/crsctl stop crs**

3 Stop the applications that use CVM or CFS that are not under VCS control.

- Using native application commands, stop the applications that use CVM or CFS on all nodes.

- Verify that no processes use the CFS mount point:

fuser -c */mount_point*

4 Unmount CFS file systems that are not under VCS control.

- Determine the file systems that need to be unmounted by checking the output of mount command.

  # **mount -v | grep vxfs | grep cluster**

- Unmount each file system that is not controlled by VCS on each node:

  # **umount */mount_point***

5 Stop VCS to take the service groups on all nodes offline

On any node execute following command to stop VCS:

```
# hastop -all
```

6 Stopping the applications that use VxVM or VxFS that are not under VCS control

- Using native application commands, stop the applications that use VxVM or VxFS.

- Verify that no processes use the VxFS mount point:

```
# fuser -c /mount_point
```

7 Unmounting VxFS file systems that are not under VCS control.

- Determine the file systems that need to be unmounted by checking the output of mount command.

```
# mount -v | grep vxfs
```

- Unmount each file system that is not controlled by VCS on each node:

```
# umount /mount_point
```

8 To stop the process, type:

```
# ./installsfrac61 -stop <sys1> <sys2> ... <nodeN>
```

9 Remove the SF for Oracle RAC 7.0.1 patches.

- Get the list of 7.0.1 patches, type:

```
# ./installmr -listpatches
```

- Remove each patch from the patch list. For example, on Solaris 10:

```
# patchrm 143287-07
```

10 Verify that the patches have been remove on all the nodes.

11 Reboot the nodes.

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

Rolling back Veritas Cluster Server manually

Use the following procedure to roll back VCS 7.0.1 to VCS 7.0 on your cluster manually. To uninstall VCS, see the *Veritas Cluster Server 7.0 Installation Guide*.

Note: Use this procedure only when rolling back VCS. Do not roll back VCS when it is part of other products that rely on VCS, for example Storage Foundation Clustered File System High Availability or Storage Foundation for Oracle RAC.

To roll back VCS manually

1   List the service groups in your cluster and their status. On any node, type:

```
# hagrp -state
```

2   Take the ClusterService service group offline if it is running. On any node, type:

```
# hagrp -offline -force ClusterService -sys system
```

3   Make the VCS configuration writable. On any node, type:

```
# haconf -makerw
```

4   Freeze all service groups. On any node, type:

```
# hagrp -freeze service_group -persistent
```

where *service_group* is the name of the service group. Note that the ClusterService group cannot be frozen.

5   Save the configuration (main.cf) file with the groups frozen. On any node, type:

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

6   Make a backup copy of the current main.cf and all types.cf configuration files. For example, on one node in the cluster, type:

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

7   Shut down VCS. On any node, type:

```
# /opt/VRTSvcs/bin/hastop -all -force
```

8 Shut down CmdServer. On each node, type:

```
# /opt/VRTSvcs/bin/CmdServer -stop
```

9 Verify that VCS has shut down. On any node, type:

```
# /sbin/gabconfig -a
```

The output resembles: GAB Port Memberships Port a gen 23dc0001 membership 01 The output shows no membership for port h.

10 For Solaris 10, on nodes that run non-global zones, check if the non-global zones are in the running state. Boot the non-global zones that are not in the running state.

- Check the zone's state. On each node, type:

```
# zoneadm list -icv
```

- Boot the zone if it is not in the running state. On each node, type:

```
# zoneadm -z zone boot
```

where *zone* is the name of the non-global zone.

Note: Do not configure one or more Solaris zones to boot from the shared storage.

---

11 Unconfigure vxfen if the VCS cluster uses the fencing option. On each node, type:

```
# /sbin/vxfenconfig -U
```

12 Unload vxfen. On each node, perform the following steps:

- Identify the vxfen kernel module, for example:

```
# modinfo|grep vxfen
210 7ba44000 39488 258 1 vxfen (VRTS Fence 7.0.1)
```

- Unload vxfen using the module number.

```
# modunload -i module ID
```

13 Unconfigure GAB. On each node, type:

```
# /sbin/gabconfig -U
```

14 Unload GAB. On each node, perform the following steps:

- Identify the GAB kernel module. For example:

  ```
  # modinfo | grep gab
  149 50cc6000 2b451 112 1 gab (GAB device 7.0.1)
  ```

- Unload GAB using the module number:

  ```
  # modunload -i module ID
  ```

15 Unconfigure LLT. On each node, perform the following steps:

- Type:

  ```
  # /sbin/lltconfig -U
  ```

- Type **y** on each node in response to the message.

16 Unload LLT. On each node, perform the following steps:

- Identify the LLT kernel module. For example:

  ```
  # modinfo | grep llt
  218 7b386000  42d38 313   1  llt (LLT 7.0.1)
  ```

- Unload LLT using the module number:

  ```
  # modunload -i module ID
  ```

17 Remove the VCS 7.0.1 patches. On each node, perform the following steps:

- Get the list of 7.0.1 patches, type:

  ```
  # ./installmr -listpatches
  ```

- Remove each patch from the patch list. For example:

  ```
  # patchrm 148492-02
  ```

18 Verify that the patches have been removed. On each node, type:

```
# showrev -p | grep VRTS
```

19 If the LLT, GAB, or VXFEN modules cannot be stopped or unloaded following the patch removal, reboot all nodes in the cluster.

20 If you do not perform step 19, start the VCS components manually. On each node, type:

```
# /sbin/lltconfig -c
# /sbin/gabconfig -cx
# /sbin/vxfenconfig -c
# /opt/VRTSvcs/bin/hastart
```

You do not have to start vxfen unless you use the fencing option.

21 After VCS has started, perform the following steps:

- Verify all resources have been probed. On any node, type:

  ```
  # hastatus -summary
  ```

- Unfreeze all service groups. On any node, type:

  ```
  # haconf -makerw
  # hagrp -unfreeze service_group -persistent
  # haconf -dump -makero
  ```

  where *service_group* is the name of the service group.

22 Bring online the ClusterService service group, if necessary. On any node type:

```
# hagrp -online ClusterService -sys system
```

where *system* is the node name.

About the installation and the uninstallation scripts

This appendix includes the following topics:

- About the installation and the uninstallation scripts

About the installation and the uninstallation scripts

Veritas™ Veritas InfoScale 7.0.1 provides an installation and upgrade script. To install or upgrade the patches that are included in this release, you can use the `installmr` script. The `installmr` script lets you install or upgrade all the patches that are associated with the packages installed.

For more information regarding installation,

Veritas has introduced a new Install Bundles feature to help you install or upgrade directly to maintenance level with one execution. You can use the `-base_path` option to install or upgrade base and maintenance bundles. There are a few prerequisites for using Install Bundles feature for installation and upgrade of 7.0.1 mentioned below:

The installmr script options

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

Table A-1   The command line options for the product installmr script

Command Line Option	Function
system1 system2...	Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.
-base_path	The *-base_path* option is used to define the path of a base level release to be integrated with a maintenance level release in order for the two releases to be simultaneously installed.
-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 patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch3_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.
-patch4_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.
-patch5_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.
–precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Veritas recommends doing a precheck before installing a product.
–postcheck	Checks any issues after installation or upgrading on the system.

Table A-1   The command line options for the product installmr script
(continued)

Command Line Option	Function
–responsefile *response_file*	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
–logpath *log_path*	Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved.
–tmppath *tmp_path*	Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
–timeout *timeout_value*	The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the -timeout option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 will prevent the script from timing out. The -timeout option does not work with the -serial option
–keyfile *ssh_key_file*	Specifies a key file for secure shell (SSH) installs. This option passes `-i ssh_key_file` to every SSH invocation.
–hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.
–patchpath *patch_path*	Designates the path of a directory that contains all patches to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.

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

Command Line Option	Function
–jumpstart *dir_path*	Produces a sample finish file for Solaris JumpStart installation. The *dir_path* indicates the path to the directory in which to create the finish file. **Note:** This option is supported only on Solaris 10. **Note:** The –jumpstart option is not supported with `-base_path` option.
–rootpath *root_path*	Specifies an alternative root directory on which to install packages. **Note:** This option is supported only on Solaris 10.
-flash_archive*<flash_archive_path>*	The `-flash_archive` option is used to generate Flash archive scripts which can be used by Solaris Jumpstart Server for automated Flash archive installation of all packages and patches for every product, an available location to store the post deployment scripts should be specified as a complete path. The -flash_archive option is supported on Solaris only. **Note:** This option is supported only on Solaris 10. **Note:** The `-flash_archive` option is not supported with `-base_path` option.
-require	The `-require` option is used to specify a installer patch file.
–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.

Table A-1 The command line options for the product installmr script
(continued)

Command Line Option	Function
–rsh	Specifies this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.
–redirect	Displays progress details without showing the progress bar.
–pkgset	Discovers and displays the package group (minimum, recommended, all) and packages that are installed on the specified systems.
–pkgtable	Displays product's packages in correct installation order by group.
–listpatches	The `-listpatches` option displays product patches in correct installation order.
–comcleanup	The `-comcleanup` option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
–version	Checks and reports the installed products and their versions. Identifies the installed and missed packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missed packages and patches where applicable. Lists the installed patches and available updates for the installed product if an Internet connection is available.
–nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.

Table A-1   The command line options for the product installmr script
(continued)

Command Line Option	Function
-disable_dmp_native_support	Disables Dynamic multi-pathing support for native the LVM volume groups/ZFS pools during an upgrade. Retaining Dynamic multi-pathing support for the native LVM volume groups/ZFS pools during an upgrade increases package upgrade time depending on the number of LUNs and native LVM volume groups/ZFS pools configured on the system. The `-disable_dmp_native_support` option is supported in upgrade scenario only.
-noipc	Disables the installer from making outbound networking calls to Veritas Operations Readiness Tool (SORT) in order to automatically obtain patches and release information updates.

The uninstallmr script options

The following table lists the command line options for uninstallmr script:

Table A-2   The command line options for the product uninstallmr script

Command Line Option	Function
system1 system2...	Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.
–responsefile *response_file*	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.

Table A-2   The command line options for the product uninstallmr script
(continued)

Command Line Option	Function
–logpath *log_path*	Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved.
–tmppath *tmp_path*	Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
–timeout *timeout_value*	The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the -timeout option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 will prevent the script from timing out. The -timeout option does not work with the -serial option
–keyfile *ssh_key_file*	Specifies a key file for secure shell (SSH) installs. This option passes `-i ssh_key_file` to every SSH invocation.
–hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.
–rootpath *root_path*	Specifies an alternative root directory on which to install packages.
–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.
–rsh	Specifies this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.
–redirect	Displays progress details without showing the progress bar.

Table A-2 The command line options for the product uninstallmr script
(continued)

Command Line Option	Function
–listpatches	The `-listpatches` option displays product patches in correct installation order.
–comcleanup	The `-comcleanup` option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
–version	Checks and reports the installed products and their versions. Identifies the installed and missed packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missed packages and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available.
-require	The `-require` option is used to specify a installer patch file.
-noipc	Disables the installer from making outbound networking calls to Veritas Operations Readiness Tool (SORT) in order to automatically obtain patches and release information updates.
-comsetup	Sets up the ssh or rsh communication between systems without requests for passwords or passphrases.

Index