Storage Foundation 7.4 Configuration and Upgrade Guide - Solaris

VERITAS™

Last updated: 2019-04-17

Legal Notice

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

http://www.veritas.com

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

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

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

infoscaledocs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Section 1

Introduction and configuration of Storage Foundation

Chapter 1

Introducing Storage Foundation

This chapter includes the following topics:

About Storage Foundation

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

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

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

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

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

About Veritas Replicator Option

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

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

About Veritas InfoScale Operations Manager

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

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

You can download Veritas InfoScale Operations Manager from https://sort.veritas.com/.

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

The Veritas Enterprise Administrator (VEA) console is no longer packaged with Veritas InfoScale products. If you want to continue using VEA, a software version is available for download from https://www.veritas.com/product/storage-management/infoscale-operations-manager. Storage Foundation Management Server is deprecated.

Chapter 2

Configuring Storage Foundation

This chapter includes the following topics:

Configuring Storage Foundation using the installer

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

To configure Storage Foundation

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

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

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

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

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

Configuring SF manually

You can manually configure different products within SF.

Configuring Veritas Volume Manager

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

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

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

Starting and enabling the configuration daemon

The VxVM configuration daemon (`vxconfigd`) maintains VxVM disk and disk group configurations. The `vxconfigd` communicates configuration changes to the kernel and modifies configuration information stored on disk.

Startup scripts usually invoke `vxconfigd` at system boot time. The `vxconfigd` daemon must be running for VxVM to operate properly.

The following procedures describe how to check that `vxconfigd` is started, whether it is enabled or disabled, how to start it manually, or how to enable it as required.

To determine whether `vxconfigd` is enabled, use the following command:

```
# vxdctl mode
```

The following message indicates that the `vxconfigd` daemon is running and enabled:

```
mode: enabled
```

This message indicates that `vxconfigd` is not running:

```
mode: not-running
```

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

```
mode: disabled
```

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

```
# vxconfigd
```

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

```
# vxdctl enable
```

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

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

Starting the volume I/O daemon

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

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

```
# vxiod
```

The `vxiod` daemon is a kernel thread and is not visible using the `ps` command.

If, for example, 16 `vxiod` daemons are running, the following message displays:

```
16 volume I/O daemons running
```

where 16 is the number of `vxiod` daemons currently running. If no `vxiod` daemons are currently running, start some by entering this command:

```
# vxiod set no_of_daemons
```

where the number of daemons ranges from 1 to16. Veritas recommends that at least one `vxiod` daemon should be run for each CPU in the system.

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

Using vxinstall to configure Veritas Volume Manager

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

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

The utility provides the following functions:

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

To run the command, enter

```
# vxinstall
```

which will prompt you to enter a license key:

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

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

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

VxVM will continue with the question:

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

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

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

```
# vxdg defaultdg
```

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

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

Configuring Veritas File System

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

```
/etc/vfstab
```

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

See the *Storage Foundation Administrator's Guide*.

Loading and unloading the file system module

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

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

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

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

To determine if the modules successfully loaded, enter:

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

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

You can unload the module by entering:

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

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

```
# df -F vxfs
```

vxtunefs command permissions and Cached Quick I/O

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

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

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

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

Configuring SFDB

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

To enable SFDB tools

1. Log in as root.
2. Run the following command to configure and start the vxdbd daemon. After you perform this step, entries are made in the system startup so that the daemon starts on a system restart.

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

To disable SFDB tools

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

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

Section 2

Upgrade of Storage Foundation

Chapter 3

Planning to upgrade Storage Foundation

This chapter includes the following topics:

About the upgrade

This release supports upgrades from 6.1 and later versions. If your existing installation is from a pre-6.1 version, you must first upgrade to version 6.1, then follow the procedures mentioned in this document to upgrade the product.

The installer supports the following types of upgrade:

- Full upgrade
- Automated upgrade using response files

Table 3-1 describes the product mapping after an upgrade.

Table 3-1 Veritas InfoScale product mapping after upgrade

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

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

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

Note: From 7.0 onwards, the existing Veritas InfoScale product upgrades to the higher version of the same product. For example, Veritas InfoScale Enterprise 7.1 gets upgraded to Veritas InfoScale Enterprise 7.2.

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

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

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

   If your current installation uses a permanent license key, you will be prompted to update the license to 7.4. Ensure that the license key file is downloaded on the local host, where you want to upgrade the product. The license key file must not be saved in the root directory (`/`) or the default license directory on the local host (`/etc/vx/licesnes/lic`). You can save the license key file inside any other directory on the local host.

   If you choose not to update your license, you will be registered with a keyless license. Within 60 days of choosing this option, you must install a valid license key file corresponding to the entitled license level.
3. Restores the existing configuration.

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

Supported upgrade paths

Table 3-2 lists the supported upgrade paths.

Table 3-2 Supported upgrade paths

From product version	From OS version	To OS version	To product version	To component
6.1.1 6.2, 6.2.1	Solaris 11 Update 1, 2, 3	Solaris 11 Update 1, 2, 3	Veritas InfoScale Storage 7.4	SF
7.0 7.0.1 7.1 7.2 7.3 7.3.1	Solaris 11 Update 1, 2, 3	Solaris 11 Update 1, 2, 3	Veritas InfoScale Storage 7.4	SF

Preparing to upgrade SF

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

Getting ready for the upgrade

Complete the following tasks before you perform the upgrade:

- Review the *Veritas InfoScale 7.4 Release Notes* for any late-breaking information on upgrading your system.
- Review the Veritas Technical Support website for additional information: https://www.veritas.com/support/en_US.html
- Perform the following system-level settings:
  - Set `diag-level` to `min` to perform the minimum number of diagnostics when the system boots. Depending on the configuration of your systems you may want to turn it on after you perform the upgrade.

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

 diag-level=min
```

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

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

  auto-boot?=false
  ```

- Deactivate cron to make sure that extraneous jobs are not performed while you upgrade the systems.
  Solaris 11:

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

- If zones are present, make sure that all non-global zones are booted and are in the running state before you use the Veritas InfoScale product installer to upgrade the Storage Foundation products in the global zone so that any packages present inside non-global zones also gets updated automatically.
- Make sure that the administrator who performs the upgrade has root access and a good knowledge of the operating system's administration.
- Make sure that all users are logged off and that all major user applications are properly shut down.
- Make sure that you have created a valid backup.
  See "Creating backups" on page 23.
- Ensure that you have enough file system space to upgrade. Identify where you want to copy the packages, for example `/packages/Veritas` when the root file system has enough space or `/var/tmp/packages` if the `/var` file system has enough space.
  Do not put the files under `/tmp`, which is erased during a system restart.
  Do not put the files on a file system that is inaccessible before running the upgrade script.
  You can use a Veritas-supplied disc for the upgrade as long as modifications to the upgrade script are not required.
  If `/usr/local` was originally created as a slice, modifications are required.
- Unmount all the file systems not on the `root disk`. Comment out their entries in `/etc/vfstab`. Stop the associated volumes and deport the associated disk groups. Any file systems that the Solaris operating system or Storage Foundation assumes should be in `rootdg` but are not, must be unmounted, and the associated entry in `/etc/vfstab` commented out.

- For any startup scripts in `/usr/sbin/svcadm disable`, comment out any application commands or processes that are known to hang if their file systems are not present.
- Make sure that the current operating system supports version 7.4 of the product. If the operating system does not support it, plan for a staged upgrade.
- Schedule sufficient outage time and downtime for the upgrade and any applications that use the Veritas InfoScale products. Depending on the configuration, the outage can take several hours.
- Any swap partitions not in `rootdg` must be commented out of `/etc/vfstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/vfstab` and not mounted during the upgrade. The active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.
- Make sure that the file systems are clean before upgrading.
  See "Verifying that the file systems are clean" on page 26.
- Veritas recommends that you upgrade VxFS disk layouts to a supported version before installing VxFS 7.4. Unsupported disk layout versions 4, 5, and 6 can be mounted for the purpose of online upgrading in VxFS 7.4. You can upgrade unsupported layout versions online before installing VxFS 7.4.
- Upgrade arrays (if required).
  See "Upgrading the array support" on page 27.
- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. Removing the disk in this manner offers a failback point.
- Determine if the root disk is encapsulated.
- Make sure that DMP support for native stack is disabled (`dmp_native_support=off`). If DMP support for native stack is enabled (`dmp_native_support=on`), the installer may detect it and ask you to restart the system.
- If you want to upgrade the application clusters that use CP server based fencing to version 6.1 and later, make sure that you first upgrade VCS or SFHA on the CP server systems to version 6.1 and later. And then, from 7.0.1 onwards, CP server supports only HTTPS based communication with its clients and IPM-based communication is no longer supported. CP server needs to be reconfigured if you upgrade the CP server with IPM-based CP server configured.
  For instructions to upgrade VCS or SFHA on the CP server systems, refer to the relevant Configuration and Upgrade Guides.

Creating backups

Save relevant system information before the upgrade.

To create backups

1. Log in as superuser.
2. Before the upgrade, ensure that you have made backups of all data that you want to preserve.

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

   If not, a warning message is displayed.

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

   ```
   # cp /etc/vfstab /etc/vfstab.orig
   ```
5. Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
6. If you install Veritas InfoScale Enterprise 7.4 software, follow the guidelines that are given in the *Cluster Server Configuration and Upgrade Guide* for information on preserving your VCS configuration across the installation procedure.
7. Back up the external `quotas` and `quotas.grp` files.

   If you are upgrading from 6.0.3, you must also back up the `quotas.grp.64` and `quotas.64` files.
8. Verify that quotas are turned off on all the mounted file systems.

Pre-upgrade planning when VVR is configured

Before installing or upgrading Volume Replicator (VVR):

- Confirm that your system has enough free disk space to install VVR.
- Make sure you have root permissions. You must have root permissions to perform the install and upgrade procedures.
- If replication using VVR is configured, Veritas recommends that the disk group version is at least 110 prior to upgrading.

You can check the Disk Group version using the following command:

```
# vxdg list diskgroup
```

- If replication using VVR is configured, make sure the size of the SRL volume is greater than 110 MB.
  Refer to the *Veritas InfoScale™ Replication Administrator's Guide*.
- If replication using VVR is configured, verify that all the Primary RLINKs are up-to-date on all the hosts.

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

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

- If VCS is used to manage VVR replication, follow the preparation steps to upgrade VVR and VCS agents.
- 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 3-3, if either the Primary or Secondary are running a version of VVR prior to 7.4, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 7.4, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

Table 3-3 VVR versions and checksum calculations

VVR prior to 7.4 (DG version <= 140)	VVR 7.4 (DG version >= 150)	VVR calculates checksum TCP connections?
Primary	Secondary	Yes
Secondary	Primary	Yes
Primary and Secondary		Yes
	Primary and Secondary	No

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

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

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

Planning and upgrading VVR to use IPv6 as connection protocol

SF supports using IPv6 as the connection protocol.

This release supports the following configurations for VVR:

- VVR continues to support replication between IPv4-only nodes with IPv4 as the internet protocol
- VVR supports replication between IPv4-only nodes and IPv4/IPv6 dual-stack nodes with IPv4 as the internet protocol
- VVR supports replication between IPv6-only nodes and IPv4/IPv6 dual-stack nodes with IPv6 as the internet protocol

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

Additional settings for using VVR in a localized environment

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

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

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

Verifying that the file systems are clean

Verify that all file systems have been cleanly unmounted.

To make sure the file systems are clean

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

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

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

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

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

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

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

```
file system device busy
```

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

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

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

Upgrading the array support

The Veritas InfoScale 7.4 release includes all array support in a single package, `VRTSaslapm`. The array support package includes the array support previously included in the `VRTSvxvm` package. The array support package also includes support previously packaged as external Array Support Libraries (ASLs) and array policy modules (APMs).

See the 7.4 Hardware Compatibility List for information about supported arrays.

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

After you have installed Veritas InfoScale 7.4, Veritas provides support for new disk arrays through updates to the `VRTSaslapm` package.

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

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

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

Releases are divided into the following categories:

Table 3-4 Release Levels

Level	Content	Form factor	Applies to	Release types	Download location
Base	Features	packages	All products	Major, minor, Service Pack (SP), Platform Release (PR)	FileConnect
Maintenance	Fixes, new features	packages	All products	Maintenance Release (MR), Rolling Patch (RP)	Veritas Services and Operations Readiness Tools (SORT)

Table 3-4 Release Levels *(continued)*

Level	Content	Form factor	Applies to	Release types	Download location
Patch	Fixes	packages	Single product	P-Patch, Private Patch, Public patch	SORT, Support site

When you install or upgrade using Install Bundles:

- Veritas InfoScale products are discovered and assigned as a single version to the maintenance level. Each system can also have one or more patches applied.
- Base releases are accessible from FileConnect that requires customer serial numbers. Maintenance and patch releases can be automatically downloaded from SORT.
- Patches can be installed using automated installers from the 6.0.1 version or later.
- Patches can now be detected to prevent upgrade conflict. Patch releases are not offered as a combined release. They are only available from Veritas Technical Support on a need basis.

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

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

In the example below:

- 7.4 is the base version
- 7.4.1 is the maintenance version
- 7.4.1.100 is the patch version for 7.4.1
- 7.4.0.100 is the patch version for 7.4

1. Base + maintenance:

   This integration method can be used when you install or upgrade from a lower version to 7.4.1.

Enter the following command:

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

2. Base + patch:

   This integration method can be used when you install or upgrade from a lower version to 7.4.0.100.

   Enter the following command:

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

3. Maintenance + patch:

   This integration method can be used when you upgrade from version 7.4 to 7.4.1.100.

   Enter the following command:

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

4. Base + maintenance + patch:

   This integration method can be used when you install or upgrade from a lower version to 7.4.1.100.

   Enter the following command:

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

Note: From the 6.1 or later release, you can add a maximum of five patches using *-patch_path <path_to_patch> -patch2_path <path_to_patch> ... -patch5_path <path_to_patch>*

Chapter 4

Upgrading Storage Foundation

This chapter includes the following topics:

Upgrading Storage Foundation to 7.4 using the product installer

This section describes upgrading SF from a previous release to 7.4. Veritas recommends that you perform this upgrade from single-user mode.

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

If you plan to upgrade the operating system, or if the current InfoScale product is installed on an operating system which is no longer supported by 7.4, you must perform additional steps to upgrade.

Upgrading Storage Foundation with the product installer

This section describes upgrading to the current Storage Foundation, and you do not intend to upgrade your Solaris version. Only use this procedure if you are already running a version of Solaris that is supported with 7.4.

To upgrade Storage Foundation

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

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

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

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

   Note: Do not continue until the Primary RLINKs are up-to-date.
5. Load and mount the disc.
6. To invoke the common installer, run the `installer` command on the disc as shown in this example:

   ```
   # cd /cdrom/cdrom0
   # ./installer
   ```
7. Enter `G` to upgrade and press Enter.
8. You are prompted to enter the system names (in the following example, "host1"). Enter the system name or names and then press Return.

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

   Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.
9. Installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.

10 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.

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

12 Stop the product's processes.

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

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

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

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

Upgrading Volume Replicator

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

You have the option to upgrade without disrupting replication.

See "Upgrading VVR without disrupting replication" on page 33.

Upgrading VVR without disrupting replication

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

You may also need to set up replication between versions.

See "Planning an upgrade from the previous VVR version" on page 24.

When both the Primary and the Secondary have the previous version of VVR installed, the upgrade can be performed either on the Primary or on the Secondary. We recommend that the Secondary hosts be upgraded before the Primary host in the RDS. This section includes separate sets of steps, for the Primary upgrade and for the Secondary upgrade.

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

Upgrading VVR on the Secondary

Follow these instructions to upgrade the Secondary hosts.

To upgrade the Secondary

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

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

2. Upgrade from VVR 6.0 or later to VVR 7.4 on the Secondary.
3. Do one of the following:
   - Upgrade the disk group now. Enter the following:

     ```
     # vxdg upgrade dgname
     ```

   - Upgrade the disk group later.
     If you upgrade the disk group later, be sure to pause replication before you upgrade the disk group. Also, after pausing replication, upgrade the disk group on Primary as well as Secondary.
4. Resume the replication from the Primary using the following command:

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

Upgrading VVR on the Primary

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

To upgrade the Primary

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

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

2. Upgrade from VVR 6.0 or later to VVR 7.4 on the Secondary.
3. Do one of the following:
   - Upgrade the disk group now. Enter the following:

     ```
     # vxdg upgrade dgname
     ```

- Upgrade the disk group later.
  If you upgrade the disk group later, be sure to pause replication before you upgrade the disk group. Also, after pausing replication, upgrade the disk group on Primary as well as Secondary.

4 Resume the replication from the Primary using the following command:

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

See "Planning an upgrade from the previous VVR version" on page 24.

Upgrading language packages

If you want to upgrade Veritas InfoScale products in a language other than English, you must install the required language packages after installing the English packages. Verify that the English installation is correct before you proceed.

Install the language packages as for an initial installation.

Upgrading SFDB

While upgrading to 7.4, the SFDB tools are enabled by default, which implies that the vxdbd daemon is configured. You can enable the SFDB tools, if they are disabled.

To enable SFDB tools

1 Log in as root.

2 Run the following command to configure and start the vxdbd daemon.

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

Note: If any SFDB installation with authentication setup is upgraded to 7.4, the commands fail with an error. To resolve the issue, setup the SFDB authentication again. For more information, see the *Veritas InfoScale™ Storage and Availability Management for Oracle Databases* or *Veritas InfoScale™ Storage and Availability Management for DB2 Databases*.

Chapter 5

Performing an automated SF upgrade using response files

This chapter includes the following topics:

Upgrading SF using response files

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

To perform automated SF upgrade

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

5 Mount the product disc and navigate to the folder that contains the installation program.

6 Start the upgrade from the system to which you copied the response file. For example:

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

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

Response file variables to upgrade SF

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

Table 5-1 Response file variables for upgrading SF

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

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

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

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

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

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

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

Sample response file for SF upgrade

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

```
our %CFG;

our %CFG;
$CFG{accepteula}=1;
$CFG{keys}{keyless}=[ qw(STORAGE) ];
$CFG{prod}="STORAGE74";
$CFG{opt}{upgrade}=1;
$CFG{systems}=[ qw(sys1) ];
1;
```

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

```
our %CFG;

$CFG{accepteula}=1;
$CFG{keys}{licensefile}=["<path_to_license_key_file>"];
$CFG{opt}{noipc}=1;
$CFG{opt}{upgrade}=1;
$CFG{prod}="STORAGE74";
$CFG{systems}=[ qw(sys1) ];

1;
```

Chapter 6

Upgrading SF using Boot Environment upgrade

This chapter includes the following topics:

About ZFS Boot Environment (BE) upgrade

A Boot Environment (BE) is a bootable instance of the Oracle Solaris operating system image along with any other application software packages installed into that image. System administrators can maintain multiple BEs on their systems, and each BE can have different software versions installed. Upon the initial installation of the Oracle Solaris 11 release onto a system, a BE is created.

On Solaris 11, you can use the `beadm` utility to create and administer additional BEs on your system.

Figure 6-1 Boot Environment upgrade process

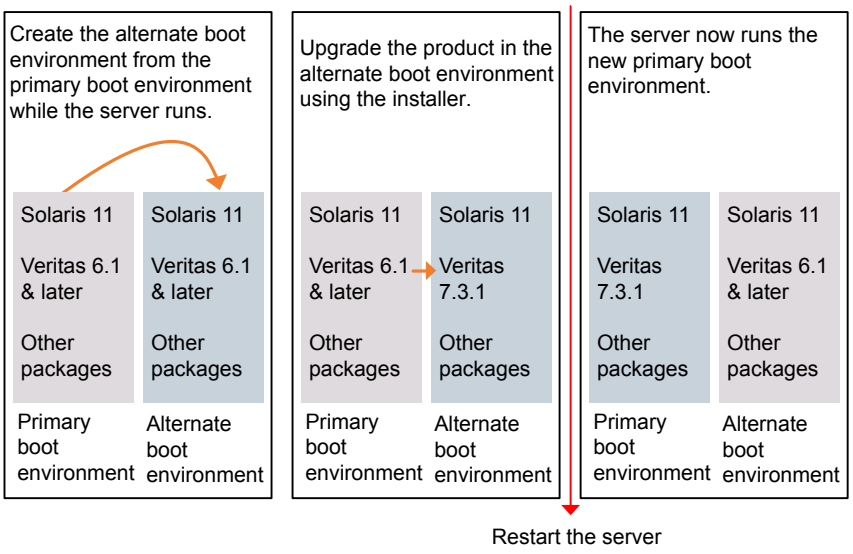

Supported upgrade paths for Boot Environment upgrade

Boot Environment upgrade can be used on Solaris 11 system only.

Veritas requires that both global and non-global zones run the same version of Veritas InfoScale products.

You can use Boot Environment upgrade in the following virtualized environments:

Table 6-1 Boot Environment upgrade support in virtualized environments

Environment	Procedure
Solaris native zones	Perform Boot Environment upgrade to upgrade both global and non-global zones. Use the standard procedure for the standby nodes. See "Performing Boot Environment upgrade on Solaris 11 systems" on page 43.

Table 6-1 Boot Environment upgrade support in virtualized environments *(continued)*

Environment	Procedure
Oracle VM Server for SPARC	Use Boot Environment upgrade procedure for Control domain as well as guest domains. See "Performing Boot Environment upgrade on Solaris 11 systems" on page 43.

Performing Boot Environment upgrade on Solaris 11 systems

Perform the Storage Foundation 7.4 Boot Environment (BE) upgrade using the installer.

Table 6-2 Upgrading SF using BE upgrade

Step	Description
Step 1	Create a new BE on the primary boot disk. See "Creating a new Solaris 11 BE on the primary boot disk" on page 43.
Step 2	Upgrade SF using the installer on the new BE. See "Upgrading SF using the installer for upgrading BE on Solaris 11" on page 44.
	To upgrade only Solaris See the Oracle documentation on Oracle Solaris 11 operating system.
Step 3	Switch the alternate BE to be the new primary. See "Completing the SF upgrade on BE on Solaris 11" on page 45.
Step 4	Verify Upgrade of SF. See "Verifying Solaris 11 BE upgrade " on page 46.

Creating a new Solaris 11 BE on the primary boot disk

Run the `beadm create` command to create a new BE on the primary boot disk.

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

To create a new BE on the primary boot disk

1. View the list of BE in the primary disk.

   ```
   # beadm list
   ```

2. Create a new BE in the primary boot disk.

   ```
   # beadm create beName
   ```

   ```
   # beadm mount beName mountpoint
   ```

If VVR is configured, it is recommended that <*beName*> should have the value *altroot.7.3* and <*mountpoint*> should have the value `/altroot.7.3`.

Upgrading SF using the installer for upgrading BE on Solaris 11

You can use the Veritas InfoScale product installer to upgrade SF on a BE.

At the end of the process, the Storage Foundation 7.4 is installed on the alternate BE.

To perform BE upgrade of SF using the installer

1. Insert the product disc with Storage Foundation 7.4 or access your copy of the software on the network.
2. Run the installer script specifying the root path as the alternate BE:

   ```
   # ./installer -upgrade -rootpath /altroot.7.3
   ```

3. Enter the names of the nodes that you want to upgrade to Storage Foundation 7.4.

   Note: Make sure that the installed version of VxFS uses the disk layout version 6 or later. If you are on a previous disk layout version, upgrade the version before you proceed with the SF installation.

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

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

5 Verify that the version of the Veritas packages on the alternate BE is 7.4.

```
# pkg -R /altroot.7.3 list VRTS\*
```

For example:

```
# pkg -R /altroot.7.3 list VRTSvxvm
```

Review the installation logs at `/altroot.7.3/opt/VRTS/install/logs`.

Completing the SF upgrade on BE on Solaris 11

At the end of the process:

- The alternate BE is activated.
- The system is booted from the alternate BE.

To complete the BE upgrade

1 Activate the alternate BE.

```
# beadm activate altroot.7.3
```

2 Restart the system. The BE on the alternate disk is activated when you restart it.

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

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

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

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

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

Verifying Solaris 11 BE upgrade

To ensure that BE upgrade has completed successfully, verify that the system have booted from the alternate BE.

To verify that BE upgrade is completed successfully

1. Verify that the alternate BE is active.

   ```
   # beadm list
   ```

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

   See "Reverting to the primary BE on a Solaris 11 system" on page 47.

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

3. In a zone environment, verify the zone configuration.

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

Administering BEs on Solaris 11 systems

Use the following procedures to perform relevant administrative tasks for BEs.

Switching the BE for Solaris SPARC

1. Display the status of Live Upgrade boot environments.

   ```
   # beadm list

   BE           Active Mountpoint Space  Policy Created
   --           ------ ---------- -----  ------ -------
   solaris      NR        /       13.08G static 2012-11-14 10:22
   altroot.7.3 -          -       3.68G static 2013-01-06 18:41
   ```

   In this example, the primary boot disk is currently *solaris*. You want to activate the alternate boot disk *altroot.7.3*.

2. Activate the Live Upgrade boot environment.

   ```
   # beadm activate altroot.7.3
   ```

3 Restart the system to complete the BE activation.

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

The system automatically selects the BE entry that was activated.

4 You can destroy an existing BE.

```
# beadm destroy altroot.7.3
```

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, enter the following:

```
# boot -Z <path to boot env>
```

For example:

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

To boot the selected entry, enter the following:

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

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

About Live Upgrade in a Volume Replicator (VVR) environment

This section provides an overview of the VVR upgrade process.

In an SF environment that uses Volume Replicator, the following scripts provide the means to upgrade the VVR configuration:

- vvr_upgrade_lu_start
- vvr_upgrade_lu_finish

The scripts are available in the scripts directory in the install media.

- Immediately before restarting the system to switch over to the alternate boot environment, run the `vvr_upgrade_lu_start` script.

  Note: Use the `vvr_upgrade_lu_start` script only when the applications are stopped and the next step is to switch over to the alternate boot environment.

- After the `vvr_upgrade_lu_start` script completes successfully, restart the system. This restart results in the system booting from the alternate boot environment.
- After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.

Chapter 7

Performing post-upgrade tasks

This chapter includes the following topics:

Optional configuration steps

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

You can perform the following optional configuration steps:

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

- To encapsulate and mirror the boot disk, follow the procedures in the "Administering Disks" chapter of the *Storage Foundation Administrator's Guide*.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
  See "Upgrading VxVM disk group versions" on page 51.

Recovering VVR if automatic upgrade fails

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

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

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

Resetting DAS disk names to include host name in FSS environments

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

To regenerate the disk names, run the following command:

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

Upgrading disk layout versions

In this release, you can create and mount file systems with disk layout version 10 and later. You can local mount disk layout version 6, 7, 8, and 9 to upgrade to a later disk layout version.

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

Disk layout version 6, 7, 8, and 9 are deprecated and you cannot cluster mount an existing file system that has any of these versions. To upgrade a cluster file system from any of these deprecated versions, you must local mount the file system and then upgrade it using the `vxupgrade` utility or the `vxfsconvert` utility.

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

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

For example, to upgrade from disk layout version 6 to a disk layout version 10, using the `vxupgrade` utility:

```
# vxupgrade -n 7 /mnt
# vxupgrade -n 8 /mnt
# vxupgrade -n 9 /mnt
# vxupgrade -n 10 /mnt
```

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

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

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

Use the following command to check your disk layout version:

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

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

Upgrading VxVM disk group versions

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

For 7.4, the Veritas Volume Manager disk group version is different than in previous VxVM releases. Veritas recommends that you upgrade the disk group version if you upgraded from a previous VxVM release.

After upgrading to SF 7.4, you must upgrade any existing disk groups that are organized by ISP. Without the version upgrade, configuration query operations continue to work fine. However, configuration change operations will not function correctly.

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

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

```
# vxdg list diskgroup
```

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

```
# vxdg upgrade diskgroup
```

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

Updating variables

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

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

Setting the default disk group

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

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

```
# vxdctl defaultdg diskgroup
```

See the *Storage Foundation Administrator's Guide*.

Upgrading the Array Support Library

VxVM provides support for new disk arrays in the form of Array Support Library (ASL) software package.

Adding JBOD support for storage arrays for which there is not an ASL available

If an array is of type A/A-A, A/P or ALUA and a suitable ASL is not available, the array must be claimed as a JBOD of type A/P. This is to prevent path delays and I/O failures arising. As JBODs are assumed to be type A/A by default, you must create appropriate JBOD entries for such arrays.

To configure an A/A-A, A/P or ALUA array as a JBOD

1. Stop all applications, such as databases, from accessing the VxVM volumes that are configured on the array, and unmount all VxFS file systems and Storage Checkpoints that are configured on the array.
2. Add the array as a JBOD of type A/P:

```
# vxddladm addjbod vid=SUN pid=T300 policy=ap
```

3. If you have not already done so, upgrade the Storage Foundation or VxVM software to 7.4. Device discovery is performed during the upgrade, and the array is claimed as a JBOD of appropriate type.

   If you have already upgraded your system to 7.4, run the following command to perform device discovery:

```
# vxdctl enable
```

4 Verify that the array has been added with the policy set to `APdisk`:

```
# vxddladm listjbod
VID    PID      Opcode Page Code Page Offset SNO length Policy
==========================================================
SUN    T300     18     -1        36          12         APdisk
```

5 Check that the correct devices are listed for the array:

```
# vxdisk list
DEVICE      TYPE           DISK     GROUP    STATUS
APdisk_0    auto:cdsdisk   -        -        online invalid
APdisk_1    auto:cdsdisk   -        -        online invalid
APdisk_2    auto:cdsdisk   -        -        online invalid
...
```

Unsuppressing DMP for EMC PowerPath disks

This section is only applicable if you want to upgrade a system that includes EMC PowerPath disks.

In releases of VxVM before 4.1, a combination of DMP subpaths and the controllers of DMP subpaths were usually suppressed to prevent interference between DMP and the EMC PowerPath multi-pathing driver. Suppression has the effect of hiding these subpaths and their controllers from DMP, and as a result VxVM cannot see the disks on these subpaths and controllers.

VxVM 4.1 and later releases have the ability to discover EMCpower disks, and configure them as autodiscovered disks that DMP recognizes are under the control of a separate multi-pathing driver. This has the benefit of allowing such disks to reconfigured in cluster-shareable disk groups. Before upgrading to VxVM 7.4, you must remove the suppression of the subpaths and controllers so that DMP can determine the association between EMCpower metadevices and `c#t#d#` disk devices.

In the following scenarios, you may need to unsuppress DMP subpaths and controllers:

- Converting a foreign disk
  See "Converting a foreign disk to auto:simple" on page 55.
- Converting a defined disk
  See "Converting a defined disk to auto:simple" on page 57.
- Converting a powervxvm disk
  See "Converting a powervxvm disk to auto:simple" on page 60.

Because EMCpower disks are auto-discovered, the `powervxvm` script should be disabled and removed from the startup script. To remove the `powervxvm` script, use the command:

```
# powervxvm remove
```

Converting a foreign disk to auto:simple

Release 4.0 of VxVM provides the `vxddladm addforeign` command to configure foreign disks with default disk offsets for the private regions and public regions, and to define them as simple disks. A foreign disk must be manually converted to `auto:simple` format before you upgrade to VxVM 7.4.

If the foreign disk is defined on a slice other than `s2`, you must copy the partition entry for that slice to that for `s0` and change the tag. If the tag of the original slice is changed, the status of the disk is seen as `online:aliased` after the upgrade.

The following example is used to illustrate the procedure. The `vxdisk list` command can be used to display the EMCpower disks that are known to VxVM:

```
# vxdisk list
DEVICE          TYPE           DISK    GROUP   STATUS
c6t0d12s2       auto:sliced    -       -       online
emcpower10c     simple         fdisk   fdg     online
...
```

The `vxprint` command is used to display information about the disk group, `fdg`:

```
# vxprint
Disk group: fdg
TY NAME    ASSOC        KSTATE  LENGTH    PLOFFS  STATE TUTIL0 PUTIL0
dg fdg     fdg          -       -         -       -     -      -
dm fdisk   emcpower10c  -       17673456  -       -     -      -
...
```

To convert a foreign disk to `auto:simple` format

1 Stop all the volumes in the disk group, and then deport it:

```
# vxvol -g fdg stopall
# vxdg deport fdg
```

2 Use the `vxddladm` command to remove definitions for the foreign devices:

```
# vxddladm rmforeign blockpath=/dev/dsk/emcpower10c \
  charpath=/dev/rdsk/emcpower10c
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
# vxdisk list
DEVICE       TYPE           DISK   GROUP   STATUS
c6t0d12s2    auto:sliced    -      -       online
...
```

3 Run the `vxprtvtoc` command to retrieve the partition table entry for the device:

```
# /etc/vx/bin/vxprtvtoc -f /tmp/vtoc /dev/rdsk/emcpower10c
```

4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
# /etc/vx/bin/vxedvtoc -f /tmp/vtoc /dev/rdsk/emcpower10c

# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
# SLICE     TAG  FLAGS   START  SIZE
  0         0x0  0x201   0      0
  1         0x0  0x200   0      0
  2         0x5  0x201   0      17675520

# THE NEW PARTITIONING WILL BE AS FOLLOWS:
# SLICE     TAG  FLAGS    START  SIZE
  0         0xf  0x201    0      17675520
  1         0x0  0x200    0      0
  2         0x5  0x201    0      17675520

DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

5 Upgrade to VxVM 7.4 using the appropriate upgrade procedure.

6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to `auto:simple` format:

```
# vxdisk list
DEVICE          TYPE            DISK    GROUP   STATUS
c6t0d12s2       auto:sliced     -       -       online
emcpower10s2    auto:simple     -       -       online
...
```

To display the physical device that is associated with the metadevice, `emcpower10s2`, enter the following command:

```
# vxdmpadm getsubpaths dmpnodename=emcpower10s2
```

7 Import the disk group and start the volumes:

```
# vxdg import fdg
# vxvol -g fdg startall
```

You can use the `vxdisk list` command to confirm that the disk status is displayed as `online:simple`:

```
# vxdisk list
DEVICE          TYPE            DISK    GROUP   STATUS
c6t0d12s2       auto:sliced     -       -       online
emcpower10s2    auto:simple     fdisk   fdg     online
```

Converting a defined disk to auto:simple

In VxVM 4.0, and particularly in prior releases, EMCpower disks could be defined by a persistent disk access record (`darec`), and identified as simple disks. If an EMCpower disk is defined with a persistent `darec`, it must be manually converted to `auto:simple` format before upgrading to VxVM 7.4.

If the defined disk is defined on a slice other than `s2`, you must copy the partition entry for that slice to that for `s0` and change the tag. If the tag of the original slice is changed, the status of the disk is seen as `online:aliased` after the upgrade.

The following example is used to illustrate the procedure. The `ls` command shows the mapping of the EMC disks to persistent disk access records:

```
# ls -l /dev/vx/dmp/emcdisk1
lrwxrwxrwx 1 root other 36 Sep 24 17:59 /dev/vx/dmp/emcdisk1->
/dev/dsk/c6t0d11s5
# ls -l /dev/vx/rdmp/emcdisk1
```

```
lrwxrwxrwx 1 root other 40Sep 24 17:59 /dev/vx/rdmp/emcdisk1->
/dev/dsk/c6t0d11s5
```

Here the fifth partition of `c6t0d11s5` is defined as the persistent disk access record `emcdisk1`.

The `vxdisk list` command can be used to display the EMCpower disks that are known to VxVM:

```
# vxdisk list
DEVICE        TYPE           DISK   GROUP   STATUS
c6t0d12s2     auto:sliced    -      -       online
emcdisk1      simple         fdisk  fdg     online
...
```

The `vxprint` command is used to display information about the disk group, `fdg`:

```
# vxprint
Disk group: fdg
TY NAME    ASSOC     KSTATE  LENGTH    PLOFFS  STATE TUTIL0 PUTIL0
dg fdg     fdg       -       -         -       -     -      -
dm fdisk   emcdisk1  -       17673456  -       -     -      -
...
```

To convert a disk with a persistent disk access record to auto:simple format

1 Stop all the volumes in the disk group, and then deport it:

```
# vxvol -g fdg stopall
# vxdg deport fdg
```

2 Use the `vxdisk rm` command to remove the persistent record definitions:

```
# vxdisk rm emcdisk1
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
# vxdisk list
DEVICE       TYPE           DISK  GROUP   STATUS
c6t0d12s2    auto:sliced    -     -       online
...
```

3 Use the `vxprtvtoc` command to retrieve the partition table entry for the device:

```
# /etc/vx/bin/vxprtvtoc -f /tmp/hdisk /dev/rdsk/c6t0d11s2
```

4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
# /etc/vx/bin/vxedvtoc -f /tmp/hdisk /dev/rdsk/c6t0d11s2

# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
# SLICE     TAG  FLAGS    START   SIZE
  4         0x0  0x200    0       0
  5         0x0  0x200    3591000 2100375
  6         0x0  0x200    0       0

# THE NEW PARTITIONING WILL BE AS FOLLOWS:
# SLICE     TAG  FLAGS    START   SIZE
  4         0x0  0x200    0       0
  5         0xf  0x200    3591000 2100375
  6         0x0  0x200    0       0

DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

5 Upgrade to VxVM 7.4 using the appropriate upgrade procedure.

6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to `auto:simple` format:

```
# vxdisk list
DEVICE           TYPE           DISK  GROUP  STATUS
c6t0d12s2        auto:sliced    -     -      online
emcpower10s2     auto:simple    -     -      online:aliased
...
```

To display the physical device that is associated with the metadevice, `emcpower10s2`, enter the following command:

```
# vxdmpadm getsubpaths dmpnodename=emcpower10s2
```

7 Import the disk group and start the volumes:

```
# vxdg import fdg
# vxvol -g fdg startall
```

You can use the `vxdisk list` command to confirm that the disk status is displayed as `online:simple`:

```
# vxdisk list
DEVICE            TYPE           DISK   GROUP  STATUS
c6t0d12s2         auto:sliced    -      -      online
emcpower10s2      auto:simple    fdisk  fdg    online:aliased
```

To allow DMP to receive correct enquiry data, the common Serial Number (C-bit) Symmetrix Director parameter must be set to enabled.

Converting a powervxvm disk to auto:simple

In VxVM 4.0, and particularly in previous releases, EMCpower disks can be defined by a persistent disk access record (darec) using `powervxvm` script, and identified as simple disks. If an EMCpower disk is used using `powervxvm`, it must be manually converted to auto:simple format before you upgrade to VxVM 7.4.

If there are any controllers or devices that are suppressed from VxVM as `powervxvm` requirement, then such controllers or disks must be unsuppressed. This is required for Veritas DMP to determine the association between PowerPath metanodes and their subpaths. After the conversion to auto:simple is complete, the `powervxvm` script is no longer useful, and should be disabled from startup script.

The following example is used to illustrate the procedure. The `ls` command shows the mapping of the EMC disks to persistent disk access records:

```
# ls -l /dev/vx/rdmp/
crw-------   1 root     root     260, 76 Feb  7 02:36 emcpower0c

# vxdisk list
DEVICE       TYPE            DISK         GROUP        STATUS
c6t0d12s2    auto:sliced       -            -          online
emcpower0c   simple          ppdsk01      ppdg         online

# vxprint
Disk group: fdg
TY NAME     ASSOC      KSTATE  LENGTH   PLOFFS STATE TUTIL0 PUTIL0
dg ppdg     ppdg       -       -        -      -     -      -
dm ppdsk01  emcpower0c -       2094960  -      -     -      -
```

To convert an EMCpower disk (defined using `powervxvm`) to auto:simple format

1 Stop all the volumes in the disk group, and then deport it:

```
# vxvol -g ppdg stopall
# vxdg deport ppdg
```

2 Use the `vxdisk rm` command to remove all emcpower disks from VxVM:

```
# vxdisk rm emcpower0c
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
# vxdisk list
DEVICE       TYPE            DISK         GROUP        STATUS
c6t0d12s2    auto:sliced       -            -          online
```

3 Use the `vxprtvtoc` command to retrieve the partition table entry for this device:

```
# /etc/vx/bin/vxprtvtoc -f /tmp/vtoc /dev/vx/rdmp/emcpower0c
```

4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
# /etc/vx/bin/vxedvtoc -f /tmp/vtoc /dev/vx/rdmp/emcpower0c
# THE ORIGINAL PARTITIONING IS AS FOLLOWS:
# SLICE     TAG  FLAGS   START  SIZE
  0         0x0  0x201   0      0
  1         0x0  0x200   0      0
  2         0x5  0x201   0      17675520

# THE NEW PARTITIONING WILL BE AS FOLLOWS:
# SLICE     TAG  FLAGS   START  SIZE
  0         0xf  0x201   0      17675520
  1         0x0  0x200   0      0
  2         0x5  0x201   0      17675520

DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

5 Upgrade to VxVM 7.4 using the appropriate upgrade procedure.

6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to auto:simple format:

```
# vxdisk list
DEVICE       TYPE            DISK         GROUP        STATUS
c6t0d12s2    auto:sliced       -            -          online
emcpower0s2  auto:simple       -            -          online
```

7 Import the disk group and start the volumes.

```
# vxdg import ppdg
# vxvol -g ppdg startall
# vxdisk list

DEVICE       TYPE            DISK         GROUP        STATUS
c6t0d12s2    auto:sliced       -            -          online
emcpower0s2  auto:simple     ppdsk01       ppdg         online
```

Converting from QuickLog to Multi-Volume support

The Version 6 and later disk layouts do not support QuickLog. The functionality provided by the Veritas Multi-Volume Support (MVS) feature replaces most of the functionality provided by QuickLog.

The following procedure describes how to convert from QuickLog to MVS. Unlike QuickLog, which allowed logging of up to 31 VxFS file systems to one device, MVS allows intent logging of only one file system per device. Therefore, the following procedure must be performed for each file system that is logged to a QuickLog device if the Version 6 or later disk layout is used.

The QuickLog device did not need to be related to the file system. For MVS, the log volume and the file system volume must be in the same disk group.

To convert Quicklog to MVS

1. Select a QuickLog-enabled file system to convert to MVS and unmount it.

   ```
   # umount myfs
   ```

2. Detach one of the QuickLog volumes from the QuickLog device that the file system had been using. This volume will be used as the new intent log volume for the file system.

   ```
   # qlogdetach -g diskgroup log_vol
   ```

3. Create the volume set.

   ```
   # vxvset make myvset myfs_volume
   ```

4. Mount the volume set.

   ```
   # mount -F vxfs /dev/vx/dsk/rootdg/myvset /mnt1
   ```

5. Upgrade the volume set's file system to the Version 7 or later disk layout.

   For example:

   ```
   # vxupgrade -n 9 /mnt1
   ```

6. Add the log volume from step 2 to the volume set.

   ```
   # vxvset addvol myvset log_vol
   ```

7 Add the log volume to the file system. The size of the volume must be specified.

```
# fsvoladm add /mnt1 log_vol 50m
```

8 Move the log to the new volume.

```
# fsadm -o logdev=log_vol,logsize=16m /mnt1
```

Verifying the Storage Foundation upgrade

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

Section 3

Post configuration tasks

Chapter 8

Performing configuration tasks

This chapter includes the following topics:

Changing root user into root role

On Oracle Solaris 11, you need to create root user to perform installation. This means that a local user cannot assume the root role. After installation, you may want to turn root user into root role for a local user, who can log in as root.

1. Log in as root user.
2. Change the root account into role.

```
# rolemod -K type=role root

# getent user_attr root

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

3. Assign the root role to a local user who was unassigned the role.

```
# usermod -R root admin
```

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

Switching on Quotas

This turns on the group and user quotas once all the nodes are upgraded to 7.4, if it was turned off earlier.

To turn on the group and user quotas

- Switch on quotas:

```
# vxquotaon -av
```

Enabling DMP support for native devices

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

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

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

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

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

Starting with Solaris 11.1, enabling DMP support for native devices also enables support for ZFS root on DMP devices. If DMP native support is enabled with an earlier Solaris version, ZFS root devices are not supported on DMP. Upgrading the operating system to version 11.1 or later does not enable support for ZFS root devices by default. To enable DMP support for the ZFS root devices, use the following procedure to enable DMP support for native devices again.

To enable DMP support for native devices

1 Turn on the tunable parameter to enable DMP support:

```
# vxdmpadm settune dmp_native_support=on
```

The `dmp_native_support` parameter is persistent.

2 If the system has Solaris version 11.1 or later installed, turning on DMP support also enables support for the ZFS root device. Reboot the system for the changes to take effect.

About configuring authentication for SFDB tools

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

Configure the `vxdbd` daemon to require authentication	See "Configuring vxdbd for SFDB tools authentication" on page 68.
Add a node to a cluster that is using authentication for SFDB tools	

Configuring vxdbd for SFDB tools authentication

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

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

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

2 Stop the `vxdbd` daemon.

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

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

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

4 Start the `vxdbd` daemon.

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

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

Section 4

Configuration and Upgrade reference

Appendix A

Installation scripts

This appendix includes the following topics:

Installation script options

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

Table A-1 Available command line options

Command Line Option	Function
-ai	The `-ai` option is supported on Solaris 11 only, and is used to generate Automated Installation manifest. This can be used by Solaris Automated Installation Server to install the Veritas InfoScale product, along with the Solaris 11 operation system.
-allpkgs	Displays all packages required for the specified product. The packages are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.

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

Command Line Option	Function
-comcleanup	The `-comcleanup` option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
-comsetup	The `-comsetup` option is used to set up the ssh or rsh communication between systems without requests for passwords or passphrases.
-configcps	The `-configcps` option is used to configure CP server on a running system or cluster.
-configure	Configures the product after installation.
-disable_dmp_native_support	Disables Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases package upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.
-fencing	Configures I/O fencing in a running cluster.
-fips	The `-fips` option is used to enable or disable security with fips mode on a running VCS cluster. It could only be used together with `-security` or `-securityonenode` option.
–hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.
-install	Used to install products on system
-online_upgrade	Used to perform online upgrade. Using this option, the installer upgrades the whole cluster and also supports customer's application zero down time during the upgrade procedure. Now this option only supports VCS and ApplicationHA.

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

Command Line Option	Function
-patch_path	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-patch2_path	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch3_path	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch4_path	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch5_path	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
–keyfile *ssh_key_file*	Specifies a key file for secure shell (SSH) installs. This option passes `-I ssh_key_file` to every SSH invocation.
-license	Registers or updates product licenses on the specified systems.
–logpath *log_path*	Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved.
-noipc	Disables the installer from making outbound networking calls to Veritas Services and Operations Readiness Tool (SORT) in order to automatically obtain patch and release information updates.

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

Command Line Option	Function
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
-pkgtable	Displays product's packages in correct installation order by group.
–postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Veritas recommends doing a precheck before installing a product.
-prod	Specifies the product for operations.
-component	Specifies the component for operations.
-redirect	Displays progress details without showing the progress bar.
-require	Specifies an installer patch file.
–responsefile *response_file*	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
–rootpath *root_path*	Specifies an alternative root directory on which to install packages. On Solaris operating systems, `-rootpath` passes `-R path` to `pkgadd` command.
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See "About configuring secure shell or remote shell communication modes before installing products" on page 79.

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

Command Line Option	Function
-security	The -security option is used to convert a running VCS cluster between secure and non-secure modes of operation.
-securityonenode	The `-securityonenode` option is used to configure a secure cluster node by node.
-securitytrust	The `-securitytrust` option is used to setup trust with another broker.
–serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-settunables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the `-tunablesfile` option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-timeout	The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the `-timeout` option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 prevents the script from timing out. The `-timeout` option does not work with the `-serial option`
–tmppath *tmp_path*	Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.

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

Command Line Option	Function
-tunables	Lists all supported tunables and create a tunables file template.
-tunables_file *tunables_file*	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-uninstall	This option is used to uninstall the products from systems
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.
-version	Checks and reports the installed products and their versions. Identifies the installed and missing packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable. Lists the installed patches, patches, and available updates for the installed product if an Internet connection is available.

About using the postcheck option

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

Note: This command option requires downtime for the node.

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

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

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

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

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

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

General checks for all products:

- All the required packages are installed.

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

Checks for Volume Manager (VM):

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

Checks for File System (FS):

- Lists the VxFS kernel modules which are not loaded (`vxfs/fdd/vxportal`.).
- Whether all VxFS file systems present in `/etc/vfstab` file are mounted.
- Whether all VxFS file systems present in `/etc/vfstab` are in disk layout 9 or higher.
- Whether all mounted VxFS file systems are in disk layout 9 or higher.

Checks for Cluster File System:

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

Appendix B

Configuring the secure shell or the remote shell for communications

This appendix includes the following topics:

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

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

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

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

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

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

Note: The product installer supports establishing passwordless communication.

Manually configuring passwordless ssh

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

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

Figure B-1 illustrates this procedure.

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

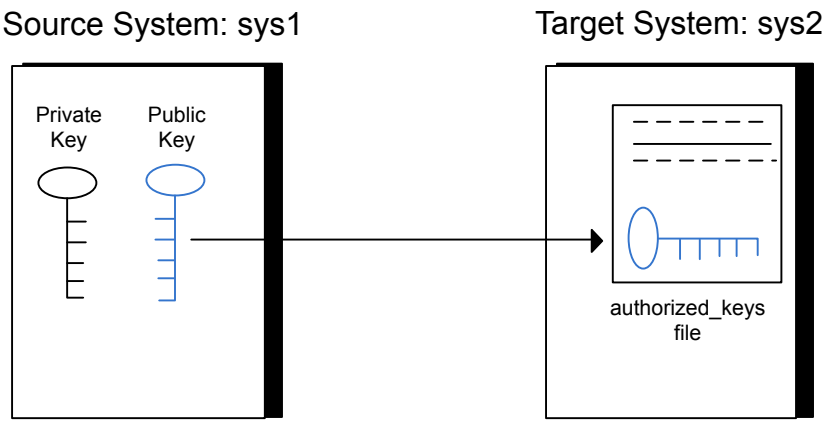

Read the ssh documentation and online manual pages before enabling ssh. Contact your operating system support provider for issues regarding ssh configuration.

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

To create the DSA key pair

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

```
sys1 # cd /
```

2 Make sure the `/.ssh` directory is on all the target installation systems (sys2 in this example). If that directory is not present, create it on all the target systems and set the write permission to root only:

Solaris 11:

```
sys2 # mkdir /root/.ssh
```

Change the permissions of this directory, to secure it.

Solaris 11:

```
sys2 # chmod go-w /root/.ssh
```

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

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

System output similar to the following is displayed:

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

For Solaris 11:

```
Your identification has been saved in /root/.ssh/id_dsa.
Your public key has been saved in /root/.ssh/id_dsa.pub.
```

4 Press Enter to accept the default location of `/.ssh/id_dsa`.

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

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

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

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

1 Make sure the secure file transfer program (SFTP) is enabled on all the target installation systems (sys2 in this example).

To enable SFTP, the `/etc/ssh/sshd_config` file must contain the following two lines:

```
PermitRootLogin           yes
  Subsystem          sftp      /usr/lib/ssh/sftp-server
```

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

To restart ssh on Solaris 11, type the following command:

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

Use the secure file transfer program.

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

Use the following command for secure file transfer:

```
sys1 # sftp sys2
```

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

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

4 Enter `yes`.

Output similar to the following is displayed:

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

5 Enter the root password of sys2.

6 At the `sftp` prompt, type the following command:

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

The following output is displayed:

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

7 To quit the SFTP session, type the following command:

```
sftp> quit
```

8 To begin the `ssh` session on the target system (sys2 in this example), type the following command on sys1:

```
sys1 # ssh sys2
```

Enter the root password of sys2 at the prompt:

```
password:
```

9 After you log in to sys2, enter the following command to append the `id_dsa.pub` file to the `authorized_keys` file:

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

10 After the `id_dsa.pub` public key file is copied to the target system (sys2), and added to the authorized keys file, delete it. To delete the `id_dsa.pub` public key file, enter the following command on sys2:

```
sys2 # rm /id_dsa.pub
```

11 To log out of the `ssh` session, enter the following command:

```
sys2 # exit
```

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

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

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

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

To verify that you can connect to a target system

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

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

where sys2 is the name of the target system.

2 The command should execute from the source system (sys1) to the target system (sys2) without the system requesting a passphrase or password.

3 Repeat this procedure for each target system.

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

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

Enter the following:

```
# ./installer -comsetup

Input the name of the systems to set up communication:
Enter the <platform> system names separated by spaces:
[q,?] sys2
Set up communication for the system sys2:

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

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

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

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

Enter the superuser password for system sys2:

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

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

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

 Checking communication on sys2 ..................... Done

Successfully set up communication for the system sys2
```

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

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

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

Command syntax with simple format:

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

```
Command syntax with advanced format:

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

    pwdutil.pl -h | -?
```

Table B-1 Options with pwdutil.pl utility

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

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

- To check ssh connection for only one host:

  pwdutil.pl check ssh hostname

- To configure ssh for only one host:

  pwdutil.pl configure ssh hostname user password

- To unconfigure rsh for only one host:

  pwdutil.pl unconfigure rsh hostname

- To configure ssh for multiple hosts with same user ID and password:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Restarting the ssh session

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

- After a terminal session is closed
- After a new terminal session is opened
- After a system is restarted

- After too much time has elapsed, to refresh ssh

To restart ssh

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

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

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

```
sys1 # ssh-add
```

Enabling and disabling rsh for Solaris

The following section describes how to enable remote shell on Solaris system.

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

See "Manually configuring passwordless ssh" on page 80.

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

To enable rsh

1 To determine the current status of `rsh` and `rlogin`, type the following command:

```
# inetadm | grep -i login
```

If the service is enabled, the following line is displayed:

```
enabled online svc:/network/login:rlogin
```

If the service is not enabled, the following line is displayed:

```
disabled disabled svc:/network/login:rlogin
```

2 To enable a disabled `rsh`/`rlogin` service, type the following command:

```
# inetadm -e rlogin
```

3 To disable an enabled `rsh`/`rlogin` service, type the following command:

```
# inetadm -d rlogin
```

4 Modify the `.rhosts` file. A separate `.rhosts` file is in the `$HOME` directory of each user. This file must be modified for each user who remotely accesses the system using rsh. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system having access to the local system. For example, if the root user must remotely access `sys1` from `sys2`, you must add an entry for `sys2.`*`companyname`*`.com` in the `.rhosts` file on `sys1`.

```
# echo "sys2.companyname.com" >> $HOME/.rhosts
```

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

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