Storage Foundation 7.4.2 Configuration and Upgrade Guide - AIX

VERITAS™

Last updated: 2020-07-30

Legal Notice

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

Technical Support

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

You can manage your Veritas account information at the following URL:
https://my.veritas.com

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

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

Documentation

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

Documentation feedback

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

You can also see documentation information or ask a question on the Veritas community site:
http://www.veritas.com/community/

Veritas Services and Operations Readiness Tools (SORT)

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

Contents

Section 1

Introduction and configuration of Storage Foundation

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.

About Veritas Services and Operations Readiness Tools (SORT)

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

SORT can help you do the following:

Prepare for your next installation or upgrade	■ List product installation and upgrade requirements, including operating system versions, memory, disk space, and architecture. ■ Analyze systems to determine if they are ready to install or upgrade Veritas products. ■ Download the latest patches, documentation, and high availability agents from a central repository. ■ Access up-to-date compatibility lists for hardware, software, databases, and operating systems.
Manage risks	■ Get automatic email notifications about changes to patches, array-specific modules (ASLs/APMs/DDIs/DDLs), and high availability agents from a central repository. ■ Identify and mitigate system and environmental risks. ■ Display descriptions and solutions for hundreds of Veritas error codes.
Improve efficiency	■ Find and download patches based on product version and platform. ■ List installed Veritas products and license keys. ■ Tune and optimize your environment.

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

To access SORT, go to:

https://sort.veritas.com

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 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/filesystems
```

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

See the *Storage Foundation Administrator's Guide*.

Configuring DMP support for booting over a SAN

For DMP to work with an LVM root disk over a SAN, configure the system to use the boot device over all possible paths.

To configure DMP support for booting over a SAN

1. Verify that each path to the root device has the same physical volume identifier (PVID) and the same volume group. Use the `lspv` command for the root volume group to verify that the PVID and volume group entries are set correctly. The PVID and volume group entries in the second and third columns of the output should be identical for all the paths.

   In this example, the LVM root disk is multi-pathed with four paths. The output from the `lspv` command for the root volume group (rootvg) is as follows:

   ```
   # lspv | grep rootvg
   hdisk374 00cbf5ce56def54d rootvg active
   hdisk375 00cbf5ce56def54d rootvg active
   hdisk376 00cbf5ce56def54d rootvg active
   hdisk377 00cbf5ce56def54d rootvg active
   ```

2. If the PVID and volume group entries are not set correctly on any of the paths, use the `chdev` command to set the correct value.

   For example, the following output shows that the `hdisk377` path is not set correctly:

   ```
   # lspv
   hdisk374 00cbf5ce56def54d rootvg active
   hdisk375 00cbf5ce56def54d rootvg active
   hdisk376 00cbf5ce56def54d rootvg active
   hdisk377 none             None
   ```

   To set the PVID for the path, use the following command:

   ```
   # chdev -l hdisk377 -a pv=yes
   hdisk377 changed
   ```

   The output of the `lspv` command now shows the correct values:

   ```
   # lspv | grep rootvg
   hdisk374 00cbf5ce56def54d rootvg active
   hdisk375 00cbf5ce56def54d rootvg active
   hdisk376 00cbf5ce56def54d rootvg active
   hdisk377 00cbf5ce56def54d rootvg active
   ```

3 If any path to the target disk has SCSI reserve ODM attribute set, then change the attributes to release the SCSI reservation from the paths, on a restart.

- If a path has the `reserve_policy` attribute set, change the`reserve_policy` attribute to `no_reserve` for all the paths.

```
# lsattr -E1 hdisk557 | grep res

reserve_policy single_path
Reserve Policy True

# chdev -l hdisk557 -a reserve_policy=no_reserve -P

hdisk557 changed
```

- If a path has the `reserve_lock` attribute set, change the `reserve_lock`attribute to `no`.

```
#  lsattr -E1 hdisk558 | grep reserve_lock

reserve_lock yes
Reserve Device on open True

# chdev -l hdisk558 -a reserve_lock=no -P

hdisk558 changed
```

4 Set the boot list to include all the paths of current boot disk.

```
# bootlist -m normal hdisk374 hdisk375 hdisk376 hdisk377 blv=hd5
```

Verify that the boot list includes all paths and that each path shows the default boot volume hd5:

```
# bootlist -m normal -o
hdisk374 blv=hd5
hdisk375 blv=hd5
hdisk376 blv=hd5
hdisk377 blv=hd5
```

5 If the `blv` option is not set for a path to the disk, use the `bootlist` command to set it. For example:

```
# bootlist -m normal hdisk374 hdisk375 hdisk376 hdisk377 blv=hd5
```

6 Run one of the following commands to configure DMP on the root disk:

- The recommended method is to turn on DMP support for LVM volumes, including the root volume.

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

- The following command enables DMP support for LVM volumes only for the root disk.

  ```
  # vxdmpadm native enable vgname=rootvg
  ```

7 Reboot the system. DMP takes control of the SAN boot device to perform load balancing and failover.

8 Verify whether DMP controls the root disk.

```
# vxdmpadm native list vgname=rootvg
PATH            DMPNODENAME
============================
hdisk374        ams_wms0_491
hdisk375        ams_wms0_491
hdisk376        ams_wms0_491
hdisk377        ams_wms0_491

# lspv | grep rootvg
hdisk374 00cbf5ce56def54d rootvg active
hdisk375 00cbf5ce56def54d rootvg active
hdisk376 00cbf5ce56def54d rootvg active
hdisk377 00cbf5ce56def54d rootvg active
```

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

The installer supports the following types of upgrade:

- Full upgrade
- Automated upgrade using response files

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

Table 3-1 Veritas InfoScale product mapping after upgrade

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

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

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

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

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

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

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

   If your current installation uses a permanent license key, you will be prompted to update the license to 7.4.2. Ensure that the license key file is downloaded on the local host, where you want to upgrade the product. The license key file must not be saved in the root directory (`/`) or the default license directory on the local host (`/etc/vx/licenses/lic`). You can save the license key file inside any other directory on the local host.

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

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

5. Starts the configured components.

Supported upgrade paths

You can upgrade to Veritas InfoScale 7.4.2 only if your currently installed product has one of the base versions: 6.2.1, 7.2, 7.3.1, 7.4.1. If your existing installation does not have one of these base versions, you must first upgrade your current installation to one of these versions. Then, follow the procedures mentioned in the Configuration and Upgrade Guide for the component configured with your InfoScale product.

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
6.2.1	AIX 7.1 TL2, TL3, TL4, TL5	AIX 7.1 TL4, TL5 AIX 7.2 TL3, TL4	Veritas InfoScale Storage 7.4.2	SF
7.2	AIX 7.1 TL3, TL4, TL5 AIX 7.2 TL0, TL1, TL2	AIX 7.1 TL4, TL5 AIX 7.2 TL3, TL4	Veritas InfoScale Storage 7.4.2	SF
7.3.1	AIX 7.1 TL4, TL5 AIX 7.2 TL0, TL1, TL2, TL3, TL4	AIX 7.1 TL4, TL5 AIX 7.2 TL3, TL4	Veritas InfoScale Storage 7.4.2	SF
7.4.1	AIX 7.1 TL4, TL5 AIX 7.2 TL0, TL1, TL2, TL3	AIX 7.1 TL4, TL5 AIX 7.2 TL3, TL4	Veritas InfoScale Storage 7.4.2	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 7.4.2 Release Notes* for any late-breaking information on upgrading your system.
- Review the Veritas Technical Support website for additional information: https://www.veritas.com/support/en_US.html
- You can configure the Veritas Telemetry Collector while upgrading, if you have do not already have it configured. For more information, refer to the *About telemetry data collection in InfoScale* section in the *Veritas Installation guide*.
- Make sure that the administrator who performs the upgrade has root access and a good knowledge of the operating system's administration.
- Make sure that all users are logged off and that all major user applications are properly shut down.
- Make sure that you have created a valid backup.
  See "Creating backups" on page 22.
- Ensure that you have enough file system space to upgrade. Identify where you want to copy the filesets, 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 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.
- For any startup scripts in `/etc/init.d/`, 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 7.4.2 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.
- Make sure that the file systems are clean before upgrading.

See "Verifying that the file systems are clean" on page 26.

- 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.
- 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.

Preparing for an upgrade of Storage Foundation

Before the upgrade of Storage Foundation to a new release, synchronize snapshots.

To prepare for an upgrade of Storage Foundation

1. Log in as `root`.
2. Stop activity to all file systems and raw volumes, for example by unmounting any file systems that have been created on volumes.

   ```
   # umount mnt_point
   ```

3. Stop all the volumes by entering the following command for each disk group:

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

4. Upgrade AIX on your system to the required levels if applicable.

Creating backups

Save relevant system information before the upgrade.

To create backups

1. Log in as superuser.
2. Make a record of the mount points for VxFS file systems and the VxVM volumes that are defined in the `/etc/filesystems` file. You need to recreate these entries in the `/etc/filesystems` file on the freshly upgraded system.
3. Before the upgrade, ensure that you have made backups of all data that you want to preserve.
4. Installer verifies that recent backups of configuration files in VxVM private region have been saved in `/etc/vx/cbr/bk`.

   If not, a warning message is displayed.

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

   ---

5. Copy the `filesystems` file to `filesystems.orig`:

   ```
   # cp /etc/filesystems /etc/filesystems.orig
   ```

6. Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
7. If you install Veritas InfoScale Enterprise 7.4.2 software, follow the guidelines that are given in the *Cluster Server Configuration and Upgrade Guide* for information on preserving your VCS configuration across the installation procedure.
8. Back up the external `quotas` and `quotas.grp` files.

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

Pre-upgrade planning when VVR is configured

Before installing or upgrading Volume Replicator (VVR):

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

```
# vxdg list diskgroup
```

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

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

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

- If VCS is used to manage VVR replication, follow the preparation steps to upgrade VVR and VCS agents.

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

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

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

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

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

Planning an upgrade from the previous VVR version

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

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

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 7.4.2, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 7.4.2, 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 7.4.2 (DG version <= 140)	**VVR 7.4.2 (DG version >= 290)**	**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

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 -V vxfs filesystem
# /opt/VRTS/bin/mount -V vxfs filesystem 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 fileset 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 fileset 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 7.4.2 release includes all array support in a single fileset, `VRTSaslapm`. The array support fileset includes the array support previously included in the `VRTSvxvm` fileset. The array support fileset also includes support previously packaged as external Array Support Libraries (ASLs) and array policy modules (APMs).

See the 7.4.2 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` fileset exits with an error if external ASLs or APMs are detected.

After you have installed Veritas InfoScale 7.4.2, Veritas provides support for new disk arrays through updates to the `VRTSaslapm` fileset.

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

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

Beginning with version 6.2.1, you can easily install or upgrade your systems directly to a base, maintenance, patch level or a combination of multiple patches and packages together in one step using Install Bundles. With Install Bundles, the installer has the ability to merge so that customers can install or upgrade directly to maintenance or patch levels in one execution. The various scripts, filesets, 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	filesets	All products	Major, minor, Service Pack (SP), Platform Release (PR)	FileConnect
Maintenance	Fixes, new features	filesets	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	filesets	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 6.2.1 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 filesets and patches from different media paths, and merge fileset 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 filesets 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:

- 7.4.2 is the base version
- 7.4.2.1 is the maintenance version
- 7.4.2.1.100 is the patch version for 7.4.2.1
- 7.4.2.0.100 is the patch version for 7.4.2

1. Base + maintenance:

   This integration method can be used when you install or upgrade from a lower version to 7.4.2.1.

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 7.4.2.0.100.

   Enter the following command:

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

3. Maintenance + patch:

   This integration method can be used when you upgrade from version 7.4.2 to 7.4.2.1.100.

   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 7.4.2.1.100.

   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 with the product installer

This section describes upgrading from Storage Foundation products to 7.4.2.

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.

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

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

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

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

5 From the disc (or if you downloaded the software) , run the `installer` command.

```
# ./installer
```

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

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

```
Enter the system names separated by spaces: [q,?] sys1 sys2
```

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

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

9 Stop the product's processes.

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

If you select `y`, the installer stops the product processes and makes some configuration updates before it upgrades.

10 The installer stops, uninstalls, reinstalls, and starts specified filesets.

11 The Storage Foundation software is verified and configured.

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

13 Restart the systems if the installer prompts restart to enable DMP native support.

Upgrade Storage Foundation and AIX on a DMP-enabled rootvg

The following upgrade paths are supported to upgrade SF and AIX on a DMP-enabled rootvg

Table 4-1 Upgrade paths for SF on a DMP-enabled rootvg

Upgrade path	Procedure
Previous version of SF on AIX 7.1	See "Upgrading from prior version of SF on AIX 7.1 to SF 7.4.2 on a DMP-enabled rootvg" on page 33.
Upgrade from AIX 7.1 to AIX 7.2 in Veritas InfoScale 7.4.2	See "Upgrading the operating system from AIX 7.1 to AIX 7.2 in Veritas InfoScale 7.4.2 " on page 34.

Upgrading from prior version of SF on AIX 7.1 to SF 7.4.2 on a DMP-enabled rootvg

When you upgrade from a previous version of SF on a DMP-enabled rootvg to Veritas InfoScale Storage 7.4.2, you must disable DMP root support before performing the upgrade. Enable the DMP root support after the upgrade. If the AIX version is not supported by Veritas InfoScale Storage 7.4.2, an operating system upgrade is required.

To upgrade from an earlier release of SF to SF 7.4.2 on a DMP-enabled rootvg

1. Disable DMP support for the rootvg:

   For release SF 6.2.1 or later:

   ```
   # vxdmpadm native disable vgname=rootvg
   Please reboot the system to disable DMP support for LVM
   bootability
   ```

2. Restart the system.
3. Upgrade SF to 7.4.2.

   Run the installer command on the disc, and enter G for the upgrade task.

   See "Upgrading Storage Foundation with the product installer" on page 31.
4. Restart the system.

5 Enable DMP for rootvg.

```
# vxdmpadm native enable vgname=rootvg
Please reboot the system to enable DMP support for LVM bootability
```

6 Restart the system. After the restart, the system has DMP root support enabled.

Upgrading the operating system from AIX 7.1 to AIX 7.2 in Veritas InfoScale 7.4.2

In Veritas InfoScale 7.4.2, when you upgrade the operating system from AIX 7.1 to AIX 7.2, DMP root support is not automatically enabled.

To upgrade AIX and enable DMP support for rootvg

1 Disable DMP support for rootvg.

2 Restart the system.

3 Upgrade the operating system from AIX 7.1 to AIX 7.2.

4 Enable DMP support for rootvg.

5 Restart the system. After the restart, the system has DMP root support enabled.

Upgrading SF on a Virtual I/O server (VIOS) from 6.2.1 or later to 7.4.2

This section provides the instructions to upgrade SF on a Virtual I/O server (VIOS) from 6.2.1 or later to 7.4.2.

To upgrade SF on VIOS

1 Shut down all Virtual I/O clients not having a failover capability, and only dependent on the Virtual I/O server being upgraded.

2 Disable DMP support for the `rootvg`:

```
# vxdmpadm native disable vgname=rootvg
Please reboot the system to disable DMP support for LVM
bootability
```

3 Restart the system.

4 Log in to the VIO Server partition.

Use the following command to access the non-restricted root shell.

```
$ oem_setup_env
```

Note: In this procedure, invoke all subsequent commands from the non-restricted shell.

Veritas recommends that you take a backup, in case you want to revert back to the earlier version.

---

5 Unconfigure all virtual devices from all virtual adapters.

```
# rmdev -p  vhost0
vtscsi0 Defined
..
```

6 Follow the procedure to upgrade SF on a Virtual I/O server.

See "Upgrading Storage Foundation with the product installer" on page 31.

7 If required, reconfigure all the virtual target devices from all the virtual adapters.

```
# cfgmgr -p vhost0
```

8 Enable DMP for `rootvg`.

```
# vxdmpadm native enable vgname=rootvg
Please reboot the system to enable DMP support for LVM bootability
```

9 Restart the system. After the restart, the system has DMP root support enabled.

10 For all the Virtual I/O servers, repeat step 1 through step 5.

11 Restart all the Virtual I/O clients you had shut down, and verify the configuration.

Upgrading the AIX operating system

Use this procedure to upgrade the AIX operating system if OS upgrade is needed. You must upgrade to a version that Veritas InfoScale Enterprise 7.4.2 supports.

To upgrade the AIX operating system

1. Create the `install-db` file.

   ```
   # touch /etc/vx/reconfig.d/state.d/install-db
   ```

   Note: The AIX OS upgrade may involve single or multiple reboots. It is necessary to create this file to prevent Veritas Volume Manager from starting VxVM daemons or processes.

2. Stop activity to all file systems and raw volumes, for example by unmounting any file systems that have been created on volumes.

   ```
   # umount mnt_point
   ```

3. Stop all the volumes by entering the following command for each disk group:

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

4. Upgrade the AIX operating system. See the operating system documentation for more information.
5. Apply the necessary APARs.

   For information about APARs required for Veritas InfoScale Storage 7.4.2, refer to the *Veritas InfoScale 7.4.2 Release Notes*.
6. Restart the system.

   ```
   # shutdown -Fr
   ```

7. Enable SF to start after you restart.

   ```
   # rm /etc/vx/reconfig.d/state.d/install-db
   ```

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 37.

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 7.4.2 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 7.4.2 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 SFDB

While upgrading to 7.4.2, 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 7.4.2, 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 filesets 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 filesets that are needed during the install. The default location is /opt/VRTStmp. 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{edgeserver_host}	Use this parameter to configure the edge server. Enter **telemetry.veritas.com** to use the Veritas Cloud Receiver, which is a preconfigured, cloud-based edge server deployed by Veritas. Optional or required: required **Note:** An edge server is used to collect licensing and platform related information from InfoScale products as part of the Veritas Product Improvement Program. The information collected helps identify how customers deploy and use the product, and enables Veritas to manage customer licenses more efficiently.

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

Variable	Description
$CFG{edgeserver_port}	Use this parameter to configure the port number of the edge server. Enter **443**, which is the port number used by the Veritas Cloud Receiver. Optional or required: required **Note:** An edge server is used to collect licensing and platform related information from InfoScale products as part of the Veritas Product Improvement Program. The information collected helps identify how customers deploy and use the product, and enables Veritas to manage customer licenses more efficiently.
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 fileset 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

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

Variable	Description
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
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}="STORAGE742";
$CFG{opt}{upgrade}=1;
$CFG{systems}=[ qw(sys1) ];
$CFG{edgeserver_host}="telemetry.veritas.com";
$CFG{edgeserver_port}=443;
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}="STORAGE742";
$CFG{systems}=[ qw(sys1) ];
$CFG{edgeserver_host}="telemetry.veritas.com";
$CFG{edgeserver_port}=443;
1;
```

Chapter 6

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 upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.

See "Upgrading VxVM disk group versions" on page 48.

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 only file systems with disk layout version 12, 13, 14, 15, and 16. You can local mount disk layout version 6, 7, 8, 9, 10, and 11 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, 9, 10, and 11 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 12, using the `vxupgrade` utility:

```
# vxupgrade -n 7 /mnt
# vxupgrade -n 8 /mnt
# vxupgrade -n 9 /mnt
# vxupgrade -n 10 /mnt
# vxupgrade -n 11 /mnt
# vxupgrade -n 12 /mnt
# vxupgrade -n 13 /mnt
# vxupgrade -n 14 /mnt
# vxupgrade -n 15 /mnt
# vxupgrade -n 16 /mnt
```

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 7.4.2, 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 7.4.2, 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*.

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 7

Performing configuration tasks

This chapter includes the following topics:

Switching on Quotas

This turns on the group and user quotas once all the nodes are upgraded to 7.4.2, if it was turned off earlier.

To turn on the group and user quotas

- Switch on quotas:

```
# vxquotaon -av
```

Enabling DMP support for native devices

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

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

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

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

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

To enable DMP support for native devices

- Turn on the tunable parameter to enable DMP support:

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

  The `dmp_native_support` parameter is persistent.

About configuring authentication for SFDB tools

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

Configure the `vxdbd` daemon to require authentication	See "Configuring vxdbd for SFDB tools authentication" on page 53.
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

Support for AIX Live Update

This appendix includes the following topics:

Support for AIX Live Update (Technology preview)

Veritas InfoScale supports the AIX Live Update feature. Starting with AIX Version 7.2, the AIX operating system provides the AIX Live Update feature that aims to eliminate the workload downtime that is associated with the AIX kernel update operation.

The AIX Live Update feature provides an efficient way to apply the AIX updates, ifixes, service packs, and technology levels without restarting the system. You can trigger the AIX 7.2 Live Kernel Update using the `geninstall -k` command that updates the OS automatically without any manual intervention or downtime. Though the I/O operations are paused for a few seconds, the critical enterprise workloads remain almost during the Live Update operation. The LKU framework recognizes if InfoScale is installed on the server and takes appropriate action while performing live updates.

Note: If Live update operation fails due to any AIX specific error, Veritas does not guarantee sanity of machine after LKU operation is completed.

Prerequisites to use the LKU feature with InfoScale

- The systems with InfoScale running on it must be LKU compatible
- InfoScale is running on a platform where IBM supports LKU with InfoScale

- The Technology Level to which you want to upgrade must be supported by InfoScale
- LKU should not be executed with the array having 2Mb gatekeeper disk

How does Live Update work?

- The Live kernel update operation gets initiated using the `geninstall -k` command from the original partition where the workload is currently running.
- The LKU framework provisions another LPAR on-the-fly with updated kernel extensions. This partition is referred to as a surrogate partition.
- The surrogate partition is patched with the updated kernel versions while the workload is still running on the original partition.
- Once the surrogate partition is up and running, the workload is moved from the original partition to the surrogate partition using the checkpoint and restart mechanism.
- The workload resumes on the surrogate partition in a “chrooted” environment.

When you perform an LKU operation, the `geninstall` command uses the `lvupdate.data` configuration file that is available in the `/var/adm/ras/liveupdate` directory. This configuration file contains the data that is required for the LKU operation. You can use the `lvupdate.template` file from the `/var/adm/ras/liveupdate` directory to create the `lvupdate.data` file. The template file contains the descriptions of all possible fields required for the LKU operation. The following example shows a sample `lvupdate.data` file:

```
general:
       kext_check = yes
       aix_mpio = no
 disks:
      nhdisk  = <hdisk1>
      mhdisk  = <hdisk2>
 hmc:
      lpar_id  = <lparid>
      management_console = <management console ip>
      user = <user>
```

When you create this configuration file, ensure that:

- You set the value of **aix_mpio** field to **no** to disable the native Multi-Path I/O (MPIO).
- Provide **hdisk#** as values for the **nhdisk** and **mhdisk** fields.

  - **nhdisk**: The names of disks to be used to make a copy of the original rootvg which will be used to boot the Surrogate.
  - **mhdisk**: The names of disks to be used to temporarily mirror rootvg on the Original LPAR.
- The size of the specified disks must match the total size of the original rootvg.
- These disks should be free. Application or Administrator should not use these disks for any other operation during the Live update operation.
- These disks should not be a part of any active or disabled Logical Volume Manager (LVM) volume groups.
- These disks should not be a part of any VxVM disk group and should not have any VxVM tag.

Limitations of LKU with InfoScale

Consider the following restrictions for the AIX Live Update operation with InfoScale:

- LKU supports only the storage components of InfoScale
- LKU is not supported in a CVM environment
- LKU is not supported for setups with combined configuration of DMP and third-party driver. For example, native MPIO.
- LKU does not support the following InfoScale features:
  - Clustering for HA or DR
  - Support for 3rd party multipathing solution
  - VVR and VFR Replication
  - Snapshot
  - FSS
  - SmartIO
  - Deduplication
  - Compression
  - In-memory statistics handling
  - Power VC
  - User initiated VxVM operations during LKU
  - Read-Write clones (checkpoints)
  - Cluster Filesystem

    - Partition Directories
- InfoScale product upgrades are not supported through the LKU operation
- LKU operation is not supported in high availability configurations for InfoScale
- LKU operation is not supported in presence of VxVM swap devices
- LKU operation is not supported if any of the administrative tasks like fsadm, fsck is running
- LKU operation fails if any changes like volume creation, deletion and so on are made to the VxVM configuration within the LKU start and MCR phase
- LKU operation is not supported in presence of vSCSI disk
- The integration of InfoScale products and LKU framework is supported only for the Local Mount filesystem

Known issues

LKU operation fails with the "kernel extensions are not known to be safe for Live Update: vxglm.ext(vxglm.ext64)" error.

A Live Update operation fails if a loaded kernel extension is not marked as safe in the safe list.

If the Group Lock Manager (GLM) is installed on a system, but the VRTSglm package is not marked with the SYS_LUSAFE flag, the LKU operation fails with the "kernel extensions are not known to be safe for Live Update: vxglm.ext(vxglm.ext64)" error.

Workaround:

Mark the VRTSglm package SYS_LUSAFE before initiating the LKU operation.

To add the VRTSglm package to the safe list for the Live Update operation, use the following command:

```
# lvupdateSafeKE -a /usr/lib/drivers/vxglm.ext\(vxglm.ext64\)
```

LKU operation fails if the ODM file system is mounted

In the technology preview mode, LKU operation is not supported with VRTSodm.

Workaround:

1. Unmount the ODM file system using the umount /dev/odm command.
2. Initiate the LKU operation using the geninstall -k command.

Appendix B

Installation scripts

This appendix includes the following topics:

Installation script options

Table B-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 B-1 Available command line options

Command Line Option	Function
-allpkgs	Displays all filesets required for the specified product. The filesets are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-comcleanup	The `-comcleanup` option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
-comsetup	The `-comsetup` option is used to set up the ssh or rsh communication between systems without requests for passwords or passphrases.

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

Command Line Option	Function
-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 fileset 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.
-patch_path	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-patch2_path	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.

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

Command Line Option	Function
-patch3_path	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch4_path	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch5_path	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
–keyfile *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.
-nolic	Allows installation of product filesets without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
-pkgtable	Displays product's filesets 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.

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

Command Line Option	Function
-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.
-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 68.
-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.

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

Command Line Option	Function
-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 `/opt/VRTStmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
-tunables	Lists all supported tunables and create a tunables file template.
-tunables_file *tunables_file*	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-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.

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

Command Line Option	Function
-version	Checks and reports the installed products and their versions. Identifies the installed and missing filesets and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing filesets 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 filesets are installed.
- The versions of the required filesets are correct.
- There are no verification issues for the required filesets.

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/filesystems` .

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/filesystems` file are mounted.
- Whether all VxFS file systems present in `/etc/filesystems` are in disk layout 12 or higher.
- Whether all mounted VxFS file systems are in disk layout 12 or higher.

Checks for Cluster File System:

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

Appendix C

Configuring the secure shell or the remote shell for communications

This appendix includes the following topics:

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

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 C-1 illustrates this procedure.

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

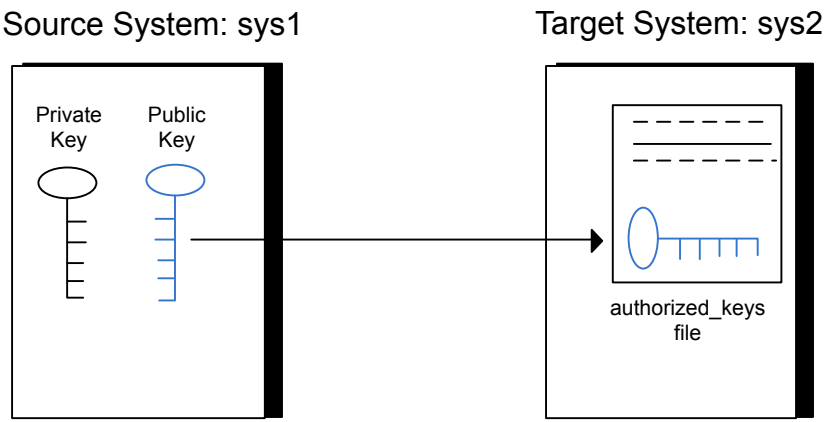

Read the ssh documentation and online manual pages before enabling ssh. Contact your operating system support provider for issues regarding ssh configuration.

Visit the Openssh website that is located at: http://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:

   Change the permissions of this directory, to secure it.

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):
   ```

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

   Use the secure file transfer program.

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

   Use the following command for secure file transfer:

   ```
   sys1 # sftp sys2
   ```

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

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

2. Enter `yes`.

   Output similar to the following is displayed:

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

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

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

   The following output is displayed:

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

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

   ```
   sftp> quit
   ```

6 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:
```

7 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
```

8 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
```

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

```
sys2 # exit
```

10 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 C-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 rsh for AIX

To enable `rsh`, create a `/.rhosts` file on each target system. Then add a line to the file specifying the full domain name of the source system. For example, add the line:

sysname.*domainname*.com root

Change permissions on the `/.rhosts` file to 600 by typing the following command:

```
# chmod 600 /.rhosts
```

After you complete an installation procedure, delete the `.rhosts` file from each target system to ensure security:

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

Appendix D

Changing NFS server major numbers for VxVM volumes

This appendix includes the following topics:

Changing NFS server major numbers for VxVM volumes

In a VCS cluster, block devices providing NFS service must have the same major and minor numbers on each cluster node. Major numbers identify required device drivers (such as AIX partition or VxVM volume). Minor numbers identify the specific devices themselves. NFS also uses major and minor numbers to identify the exported file system. Major and minor numbers must be verified to ensure that the NFS identity for the file system is the same when exported from each node.

Use the `haremajor` command to determine and reassign the major number that a system uses for shared VxVM volume block devices. For Veritas Volume Manager, the major number is set to the `vxio` driver number. To be highly available, each NFS server in a VCS cluster must have the same `vxio` driver number, or major number.

To list the major number currently in use on a system

- Use the command:

```
# haremajor -v
55
```

Run this command on each cluster node. If major numbers are not the same on each node, you must change them on the nodes so that they are identical.

To list the available major numbers for a system

- Use the command:

```
# haremajor -a
54,56..58,60,62..
```

The output shows the numbers that are not in use on the system where the command is issued.

To reset the major number on a system

- You can reset the major number to an available number on a system. For example, to set the major number to 75 type:

```
# haremajor -s 75
```