Veritas InfoScale™ Availability 7.0.1 Installation Guide - Solaris 10 x64

Platform Release

January 2016

VERITAS™

Veritas InfoScale Availability Installation Guide

Product version: 7.0.1

Document version: 7.0.1 Rev 1

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

Our support offerings include the following:

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

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

www.veritas.com/support

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

Contacting Technical Support

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

www.veritas.com/support

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information

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

Licensing and registration

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

www.veritas.com/support

Customer service

Customer service information is available at the following URL:

www.veritas.com/support

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

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

Support agreement resources

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

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

Contents

Technical Support 4

Section 1 Introduction to Veritas InfoScale Availability 14

Chapter 1 About Veritas InfoScale Availability 15

About Veritas InfoScale Availability 15

Chapter 2 Licensing Veritas InfoScale Availability 16

About Veritas InfoScale product licensing 16
Registering Veritas InfoScale using product license keys 17
Registering Veritas InfoScale product using keyless licensing 18
Updating your product licenses 20
Using the `vxlicinstupgrade` utility 20
About the `VRTSvlic` package 21

Section 2 Planning and preparation 23

Chapter 3 System requirements 24

Important release information 24
Disk space requirements 25
Hardware requirements 25
VCS hardware requirements 25
Supported operating systems and database versions 26
Number of nodes supported 26

Chapter 4 Preparing to install 27

Mounting the ISO image 27
Setting up ssh or rsh for inter-system communications 28
Obtaining installer patches 28
Disabling external network connection attempts 29
Verifying the systems before installation 30

Setting up the private network 30
Optimizing LLT media speed settings on private NICs 33
Guidelines for setting the media speed for LLT interconnects 34
Synchronizing time settings on cluster nodes 34
Creating a root user 34
Preparing zone environments 35

Section 3 Installation of Veritas InfoScale Availability 37

Chapter 5 Installing Veritas InfoScale using the installer 38

Installing Veritas InfoScale using the installer 38
Installing language packages 39

Chapter 6 Installing Veritas InfoScale Availability using response files 40

About response files 40
Syntax in the response file 41
Installing InfoScale Availability using response files 41
Response file variables to install Veritas InfoScale Availability 42
Sample response file for Veritas InfoScale Availability installation 44

Chapter 7 Installing Veritas InfoScale using operating system-specific methods 45

About installing InfoScale Availability using operating system-specific methods 45
Installing InfoScale Availability on Solaris 10 using JumpStart 46
Overview of JumpStart installation tasks 46
Generating the finish scripts 46
Preparing installation resources 47
Adding language pack information to the finish file 48
Creating the Veritas InfoScale post-deployment scripts 49
Using a Flash archive to install InfoScale Availability and the operating system 50
Manually installing InfoScale Availability using the system command 51
Installing InfoScale Availability on Solaris 10 using the pkgadd command 51
Manually installing packages on solaris10 brand zones 53

Chapter 8 Verifying the Veritas InfoScale installation 54

Verifying product installation 54
Installation log files 54
Using the installation log file 55
Using the summary file 55
Setting environment variables 55
Checking installed product versions and downloading maintenance releases and patches 56

Section 4 Upgrading Veritas InfoScale Availability 57

Chapter 9 Preparing to upgrade to Veritas InfoScale Availability 7.0.1 58

Supported upgrade types for Veritas InfoScale Availability 7.0.1 58
Supported upgrade paths for Veritas InfoScale Availability 7.0.1 59

Chapter 10 Performing full upgrade of VCS with 2048 bit key and SHA256 signature certificates 60

Stronger security with 2048 bit key and SHA256 signature certificates 60
Upgrading VCS using the product installer 61
Tasks to perform after upgrading to 2048 bit key and SHA256 signature certificates 63
Deleting certificates of non-root users after upgrading to 2048 bit key and SHA256 signature certificates 63
Re-establishing WAC communication in global clusters after upgrading to 2048 bit key and SHA256 signature certificates 64
Re-establishing trust with Steward after upgrading to 2048 bit key and SHA256 signature certificates 65
Re-establish trust between VOM and SFMH/VBS after upgrading to 2048 bit key and SHA256 signature certificates 66
Upgrading Steward to 2048 bit key and SHA256 signature certificates 67

Chapter 11 Performing Online Upgrade 69

Limitations of online upgrade 69
Upgrading VCS online using the installer 69

Tasks to perform after upgrading to 2048 bit key and SHA256 signature certificates 71

Chapter 12 Upgrading InfoScale Availability using Live Upgrade 72

Supported upgrade paths for Live Upgrade and Boot Environment upgrade 72
Performing Live Upgrade in a Solaris zone environment on Solaris 10 72
Performing Live Upgrade on Solaris 10 systems 73
Before you upgrade InfoScale Availability using Solaris Live Upgrade 74
Creating a new Solaris 10 boot environment on the alternate boot disk 75
Upgrading InfoScale Availability using the installer for Solaris 10 Live Upgrade 75
Completing the Solaris 10 Live Upgrade 76
Tasks to perform after upgrading to 2048 bit key and SHA256 signature certificates 77
Verifying the Solaris 10 Live Upgrade of InfoScale Availability 78
Administering boot environments in Solaris 10 Live Upgrade 78

Chapter 13 Performing InfoScale Availability upgrade using response files 79

Upgrading InfoScale Availability using response files 79
Response file variables to upgrade InfoScale Availability 80
Sample response file for upgrading InfoScale Availability 82

Section 5 Configuration of Veritas InfoScale Availability 83

Chapter 14 Configuring InfoScale Availability 84

Overview of tasks to configure InfoScale Availability using the product installer 85
Starting the software configuration 85
Specifying systems for configuration 86
Configuring the cluster name 87
Configuring private heartbeat links 87
Configuring the virtual IP of the cluster 90

Configuring InfoScale Availability in secure mode 92
Setting up trust relationships for your InfoScale Availability cluster 92
Configuring a secure cluster node by node 94
Configuring the first node 94
Configuring the remaining nodes 95
Completing the secure cluster configuration 96
Adding InfoScale Availability users 99
Configuring SMTP email notification 100
Configuring SNMP trap notification 101
Configuring global clusters 103
Completing the InfoScale Availability configuration 103

Chapter 15 Configuring InfoScale Availability clusters for data integrity 105

Setting up disk-based I/O fencing using installer 105
Configuring disk-based I/O fencing using installer 105
Refreshing keys or registrations on the existing coordination points for disk-based fencing using the installer 108
Initializing disks as VxVM disks 109
Checking shared disks for I/O fencing 110
Setting up non-SCSI-3 I/O fencing in virtual environments using installer 113
Setting up majority-based I/O fencing using installer 115
Enabling or disabling the preferred fencing policy 117

Chapter 16 Manually configuring InfoScale Availability 120

About configuring InfoScale Availability manually 120
Configuring LLT manually 121
Setting up /etc/llthosts for a manual installation 121
Setting up /etc/lltab for a manual installation 122
LLT directives for a manual installation 122
About LLT directives in /etc/lltab file 123
Additional considerations for LLT for a manual installation 125
Configuring GAB manually 125
Configuring InfoScale Availability manually 126
Configuring the cluster UUID when creating a cluster manually 127
Configuring InfoScale Availability in single node mode 127
Starting LLT, GAB, and InfoScale Availability after manual configuration 128

About configuring cluster using InfoScale Availability Cluster Configuration wizard 130
Before configuring a InfoScale Availability cluster using the InfoScale Availability Cluster Configuration wizard 130
Launching the InfoScale Availability Cluster Configuration wizard 131
Configuring a cluster by using the InfoScale Availability cluster configuration wizard 133
Adding a system to a InfoScale Availability cluster 136
Modifying the InfoScale Availability configuration 139
Configuring the ClusterService group 139

Chapter 17 Manually configuring the clusters for data integrity 140

Setting up disk-based I/O fencing manually 140
Setting up coordinator disk groups 141
Creating I/O fencing configuration files 142
Modifying InfoScale Availability configuration to use I/O fencing 143
Verifying I/O fencing configuration 144
Identifying disks to use as coordinator disks 145
Setting up non-SCSI-3 fencing in virtual environments manually 145
Sample /etc/vxfenmode file for non-SCSI-3 fencing 147
Setting up majority-based I/O fencing manually 151
Creating I/O fencing configuration files 151
Modifying InfoScale Availability configuration to use I/O fencing 151
Verifying I/O fencing configuration 153
Sample /etc/vxfenmode file for majority-based fencing 154

Section 6 Uninstallation of Veritas InfoScale Availability 155

Chapter 18 Uninstalling Veritas InfoScale Availability using the installer 156

About removing Veritas InfoScale Availability 156
Preparing to uninstall 156
Uninstalling InfoScale Availability packages using the product installer 157
Uninstalling Veritas InfoScale Availability using the `pkgrm` command 159

Uninstalling the language packages using the `pkgrm` command 159

Chapter 19 Uninstalling Veritas InfoScale Availability using response files 161

Uninstalling InfoScale Availability using response files 161
Response file variables to uninstall Veritas InfoScale Availability 162
Sample response file for Veritas InfoScale Availability uninstallation 163

Section 7 Installation reference 164

Appendix A Installation scripts 165

Installation script options 165

Appendix B Troubleshooting installation issues 171

Restarting the installer after a failed connection 171
About the VRTSspt package troubleshooting tools 171
Incorrect permissions for root on remote system 172
Inaccessible system 173

Index 174

Section 1

Introduction to Veritas InfoScale Availability

- Chapter 1. About Veritas InfoScale Availability
- Chapter 2. Licensing Veritas InfoScale Availability

Chapter 1

About Veritas InfoScale Availability

This chapter includes the following topics:

- About Veritas InfoScale Availability

About Veritas InfoScale Availability

Veritas InfoScale™ Availability helps keep organizations' information available and critical business services up and running with a robust software-defined approach. Organizations can innovate and gain cost benefits of physical and virtual across commodity server deployments. Maximum IT service continuity is ensured at all times, moving resiliency from the infrastructure layer to the application layer.

Chapter 2

Licensing Veritas InfoScale Availability

This chapter includes the following topics:

- About Veritas InfoScale product licensing
- Registering Veritas InfoScale using product license keys
- Registering Veritas InfoScale product using keyless licensing
- Updating your product licenses
- Using the vxlicinstupgrade utility
- About the VRTSvlic package

About Veritas InfoScale product licensing

You must obtain a license to install and use Veritas InfoScale products.

You can choose one of the following licensing methods when you install a product:

- Install with a license key for the product
  When you purchase a Veritas InfoScale product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased.
  See "Registering Veritas InfoScale using product license keys" on page 17.
- Install without a license key (keyless licensing)
  Installation without a license does not eliminate the need to obtain a license. The administrator and company representatives must ensure that a server or cluster is entitled to the license level for the products installed. Veritas reserves the right to ensure entitlement and compliance through auditing.

See "Registering Veritas InfoScale product using keyless licensing" on page 18.

If you encounter problems while licensing this product, visit the Veritas licensing support website.

www.veritas.com/licensing/process

Registering Veritas InfoScale using product license keys

You can register your product license key in the following ways:

Using the `installer` — The installer automatically registers the license at the time of installation or upgrade.

- You can register your license keys during the installation process.
  During the installation, you will get the following prompt:

  ```
  1) Enter a valid license key
  2) Enable keyless licensing and complete system licensing later

  How would you like to license the systems? [1-2,q] (2)
  ```

  Enter **1** to register your license key.
  See "Installing Veritas InfoScale using the installer" on page 38.
- You can also register your license keys using the installer menu.
  Run the following command:

  ```
  ./installer
  ```

  Select the **L) License a Product** option in the installer menu. The following menu is displayed:

  ```
  1)  Enter a valid license key
  2)  Enable keyless licensing and complete system licensing later
  ```

  Select **1)** to register license key.

Manual

If you are performing a fresh installation, run the following commands on each node:

```
# cd /opt/VRTS/bin

# ./vxlicinst -k license key

# vxdctl license init
```

If you are performing an upgrade, run the following commands on each node:

```
# cd /opt/VRTS/bin

# ./vxlicinstupgrade -k license key
```

For more information:

See "Using the `vxlicinstupgrade` utility" on page 20.

Even though other products are included on the enclosed software discs, you can only use the Veritas InfoScale software products for which you have purchased a license.

Registering Veritas InfoScale product using keyless licensing

The keyless licensing method uses product levels to determine the Veritas InfoScale products and functionality that are licensed.

You can register a Veritas InfoScale product in the following ways:

Using the `installer`

- Run the following command:

  ```
  ./installer
  ```

  The installer automatically registers the license at the time of installation or upgrade.
- You can also register your license keys using the installer menu.
  Run the following command:

  ```
  ./installer
  ```

  Select the **L) License a Product** option in the installer menu. The following menu is displayed:

  ```
  1)  Enter a valid license key
  2)  Enable keyless licensing and complete system licensing later
  ```

  Select **1)** to register license key.

Manual

Perform the following steps after installation or upgrade:

1. Change your current working directory:

   ```
   # export PATH=$PATH:/opt/VRTSvlic/bin
   ```

2. View the possible settings for the product level:

   ```
   # vxkeyless displayall
   ```

3. Register the desired product:

   ```
   # vxkeyless set prod_levels
   ```

   where *prod_levels* is a comma-separated list of keywords. The keywords are the product levels as shown by the output of step 2.

Warning: Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled, or continue with keyless licensing by managing the systems with Veritas InfoScale Operation Manager. If you fail to comply with the above terms, continuing to use the Veritas InfoScale product is a violation of your End User License Agreement, and results in warning messages.

For more information about keyless licensing, see the following URL:

http://www.veritas.com/community/blogs/introducing-keyless-feature-enablement-storage-foundation-ha-51

For more information to use keyless licensing and to download the Veritas InfoScale Operation Manager, see the following URL:

https://www.veritas.com/product/storage-management/infoscale-operations-manager

Updating your product licenses

At any time, you can update your product licenses in any of the following ways:

Move from one product to another	Perform the following steps:
	`# export PATH=$PATH:/opt/VRTSvlic/bin`
	`# vxkeyless set prod_levels`

Move from keyless licensing to key-based licensing	You will need to remove the keyless licenses by using the NONE keyword. **Note:** Clearing the keys disables the Veritas InfoScale products until you install a new key or set a new product level. `# vxkeyless [-q] set NONE` Register a Veritas InfoScale product using a license key:

Using the `vxlicinstupgrade` utility

The `vxlicinstupgrade` utility enables you to perform the following tasks:

- Upgrade to another Veritas InfoScale product
- Update a temporary license to a permanent license
- Manage co-existence of multiple licenses

On executing the `vxlicinstupgrade` utility, the following checks are done:

- If the current license key is keyless or user-defined and if the user is trying to install the keyless or user defined key of the same product.
  Example: If the 7.0.1 Foundation Keyless license key is already installed on a system and the user tries to install another 7.0.1 Foundation Keyless license key, then `vxlicinstupgrade` utility shows an error message:

  ```
  vxlicinstupgrade WARNING: The input License key and Installed key are same.
  ```

- If the current key is keyless and the newly entered license key is user-defined of the same product
  Example: If the 7.0.1 Foundation Keyless license key is already installed on a system and the user tries to install 7.0.1 Foundation user-defined license, then the `vxlicinstupgrade` utility installs the new licenses at `/etc/vx/licenses/lic` and all the 7.0.1Foundation Keyless keys are deleted and backed up at `/var/vx/licenses/lic<date-timestamp>`.
- If the current key is of higher version and the user tries to install a lower version license key.
  Example: If the 7.0.1 Enterprise license key is already installed on a system and the user tries to install the 6.0 SFSTD license key, then the `vxlicinstupgrade` utility shows an error message:

  ```
  vxlicinstupgrade WARNING: The input License key is lower than the
  Installed key.
  ```

- If the current key is of a lower version and the user tries to install a higher version license key.
  Example: If 6.0 SFSTD license key is already installed on a system and the user tries to install 7.0.1 Storage license key, then the `vxlicinstupgrade` utility installs the new licenses at `/etc/vx/licenses/lic` and all the 6.0 SFSTD keys are deleted and backed up at `/var/vx/licenses/lic<date-timestamp>`.
- Supported Co-existence scenarios:
- InfoScale Foundation and InfoScale Availability
- InfoScale Storage and InfoScale Availability

Example: If the 7.0.1 Foundation or 7.0.1 Storage license key is already installed and the user tries to install 7.0.1 Availability license key or vice -versa, then the `vxlicinstupgrade` utility installs the new licenses and both the keys are preserved at `/etc/vx/licenses/lic`.

Note: When registering license keys manually during upgrade, you have to use the `vxlicinstupgrade` command. When registering keys using the installer script, the same procedures are performed automatically.

About the `VRTSvlic` package

The `VRTSvlic` package enables product licensing. After the `VRTSvlic` is installed, the following commands and their manual pages are available on the system:

`vxlicinst`	Installs a license key for a Veritas InfoScale product See the `vxlicinst(1m)` manual page
`vxlicinstupgrade`	Upgrades your license key when you have a product or older license already present on the system. See the `vxlicinstupgrade(1m)` manual page
`vxlicrep`	Displays the currently installed licenses
`vxlictest`	Retrieves the features and their descriptions that are encoded in a license key

Section 2

Planning and preparation

- Chapter 3. System requirements
- Chapter 4. Preparing to install

Chapter 3

System requirements

This chapter includes the following topics:

- Important release information
- Disk space requirements
- Hardware requirements
- Supported operating systems and database versions
- Number of nodes supported

Important release information

Review the Release notes for the latest information before you install the product.

Review the current compatibility lists to confirm the compatibility of your hardware and software:

- For important updates regarding this release, review the Late-Breaking News TechNote on the Veritas Technical Support website:
  Placeholder
- For the latest patches available for this release, go to:
  https://sort.veritas.com
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware, visit the following URL:
  https://www.veritas.com/support/en_US/article.TECH230646
- The software compatibility list summarizes each Veritas InfoScale product stack and the product features, operating system versions, and third-party products it supports. For the latest information on supported software, visit the following URL:

https://www.veritas.com/support/en_US/article.TECH230619

Disk space requirements

Table 3-1 lists the minimum disk space requirements for Veritas InfoScale Availability.

Table 3-1 Disk space requirements

Product name	Solaris 10 (MB)
Veritas InfoScale Availability	701

Hardware requirements

VCS hardware requirements

Table 3-2 lists the hardware requirements for a VCS cluster.

Table 3-2 Hardware requirements for a VCS cluster

Item	Description
DVD drive	One drive in a system that can communicate to all the nodes in the cluster.
Disks	Typical configurations require that the applications are configured to use shared disks/storage to enable migration of applications between systems in the cluster. The SFHA I/O fencing feature requires that all data and coordinator disks support SCSI-3 Persistent Reservations (PR).
Ethernet controllers	In addition to the built-in public Ethernet controller, VCS requires at least one more Ethernet interface per system. Veritas recommends two additional network interfaces for private interconnects. You can also configure aggregated interfaces. Veritas recommends that you turn off the spanning tree algorithm on the switches used to connect private network interfaces..
Fibre Channel or SCSI host bus adapters	Typical VCS configuration requires at least one SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.
RAM	Each VCS node requires at least 1024 megabytes.

Supported operating systems and database versions

For information on supported operating systems and database versions for InfoScale Availability, see the *Veritas InfoScale Availability Release Notes*.

Number of nodes supported

Veritas InfoScale Availability supports cluster configurations up to 64 nodes. At the time of product release, cluster configurations have been qualified and tested with up to 32 nodes.

Chapter 4

Preparing to install

This chapter includes the following topics:

- Mounting the ISO image
- Setting up ssh or rsh for inter-system communications
- Obtaining installer patches
- Disabling external network connection attempts
- Verifying the systems before installation
- Setting up the private network
- Synchronizing time settings on cluster nodes
- Creating a root user
- Preparing zone environments

Mounting the ISO image

An ISO file is a disc image that must be mounted to a virtual drive for use. You must have superuser (root) privileges to mount the Veritas InfoScale ISO image.

To mount the ISO image

1. Log in as superuser on a system where you want to install Veritas InfoScale.
2. Associate the ISO image to a block device:

   ```
   # lofiadm -a <ISO_image_path> <block_device>
   ```

   Where:

   <ISO_image_path> is the complete path to the ISO image

   <block_device> is the complete path to the block device
3. Mount the image:

   ```
   # mount -F hsfs -o ro <block_device> /mnt
   ```

Setting up ssh or rsh for inter-system communications

The installer uses passwordless Secure Shell (ssh) or Remote Shell (rsh) communications among systems. During an installation, you choose the communication method that you want to use. Or, you can run the `installer -comsetup` command to set up ssh or rsh explicitly. When the installation process completes, the installer asks you if you want to remove the password-less connection. If installation terminated abruptly, use the installation script's `-comcleanup` option to remove the ssh or rsh configuration from the systems.

In most installation, configuration, upgrade (where necessary), and uninstallation scenarios, the installer configures ssh or rsh on the target systems. When you perform installation using a response file, you need to set up ssh or rsh manually, or use the`installer -comsetup` option to set up an ssh or rsh configuration from the systems.

Obtaining installer patches

You can access public installer patches automatically or manually on the Veritas Services and Operations Readiness Tools (SORT) website's Patch Finder page at:

https://sort.veritas.com/patch/finder

To download installer patches automatically

- If you are running InfoScale Availability version 7.0.1 or later, and your system has Internet access, the installer automatically imports any needed installer patch, and begins using it.

Automatically downloading installer patches requires the installer to make outbound networking calls. You can also disable external network connection attempts.

If your system does not have Internet access, you can download installer patches manually.

To download installer patches manually

1. Go to the Veritas Services and Operations Readiness Tools (SORT) website's Patch Finder page, and save the most current patch on your local system.
2. Navigate to the directory where you want to unzip the file you downloaded in step 3.
3. Unzip the patch tar file. For example, run the following command:

```
# gunzip cpi-701P2-patches.tar.gz
```

4. Untar the file. For example, enter the following:

```
# tar -xvf cpi-701P2-patches.tar
patches/
patches/CPI701P2.pl
README
```

5. Navigate to the installation media or to the installation directory.
6. To start using the patch, run the `installer` command with the `-require` option. For example, enter the following:

```
# ./installer -require /target_directory/patches/CPI701P2.pl
```

Disabling external network connection attempts

When you execute the `installer` command, the installer attempts to make an outbound networking call to get information about release updates and installer patches. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network connection attempts by the installer.

To disable external network connection attempts

- Disable inter-process communication (IPC).

  To disable IPC, run the installer with the `-noipc` option.

  For example, to disable IPC for system1 (sys1) and system2 (sys2) enter the following:

  ```
  # ./installer -noipc sys1 sys2
  ```

Verifying the systems before installation

Use any of the following options to verify your systems before installation:

- Option 1: Run Veritas Services and Operations Readiness Tools (SORT).
  For information on downloading and running SORT:
  https://sort.veritas.com

  Note: You can generate a pre-installation checklist to determine the pre-installation requirements: Go to the SORT installation checklist tool. From the drop-down lists, select the information for the Veritas InfoScale product you want to install, and click Generate Checklist.

- Option 2: Run the installer with the "-precheck" option as follows:
  Navigate to the directory that contains the installation program.
  Start the preinstallation check:

  ```
  # ./installer -precheck sys1 sys2
  ```

  where *sys1, sys2* are the names of the nodes in the cluster.
  The program proceeds in a non-interactive mode, examining the systems for licenses, packages, disk space, and system-to-system communications. The program displays the results of the check and saves them in a log file. The location of the log file is displayed at the end of the precheck process.

Setting up the private network

VCS requires you to set up a private network between the systems that form a cluster. You can use either NICs or aggregated interfaces to set up private network.

You can use network switches instead of hubs. However, Oracle Solaris systems assign the same MAC address to all interfaces by default. Thus, connecting two or more interfaces to a network switch can cause problems.

For example, consider the following case where:

- The IP address is configured on one interface and LLT on another
- Both interfaces are connected to a switch (assume separate VLANs)

The duplicate MAC address on the two switch ports can cause the switch to incorrectly redirect IP traffic to the LLT interface and vice versa. To avoid this issue, configure the system to assign unique MAC addresses by setting the `eeprom(1M)` parameter local-mac-address to true.

Refer to the *InfoScale Availability Administrator's Guide* to review VCS performance considerations.

Figure 4-1 shows two private networks for use with VCS.

Figure 4-1 Private network setups: two-node and four-node clusters

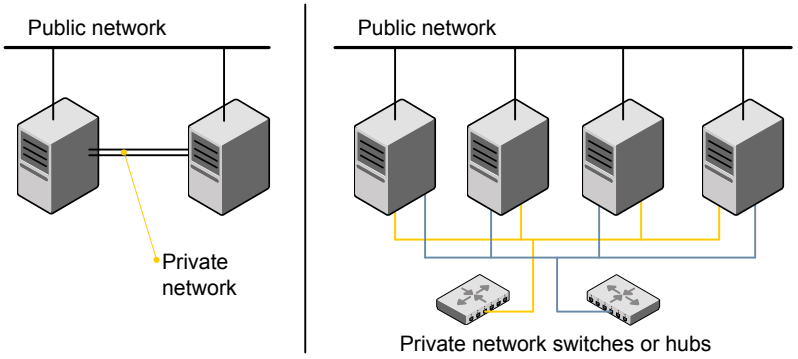

You need to configure at least two independent networks between the cluster nodes with a network switch for each network. You can also interconnect multiple layer 2 switches for advanced failure protection. Such connections for LLT are called cross-links.

Figure 4-2 shows a private network configuration with crossed links between the network switches.

Figure 4-2 Private network setup with crossed links

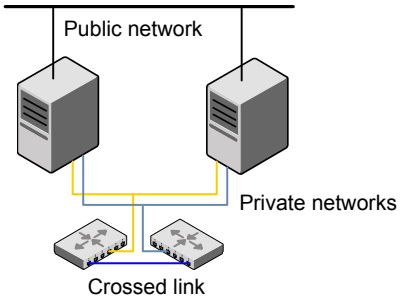

Veritas recommends one of the following two configurations:

- Use at least two private interconnect links and one public link. The public link can be a low priority link for LLT. The private interconnect link is used to share cluster status across all the systems, which is important for membership arbitration and high availability. The public low priority link is used only for heartbeat communication between the systems.
- If your hardware environment allows use of only two links, use one private interconnect link and one public low priority link. If you decide to set up only two links (one private and one low priority link), then the cluster must be configured to use I/O fencing, either disk-based or server-based fencing configuration. With only two links, if one system goes down, I/O fencing ensures that other system can take over the service groups and shared file systems from the failed node.

To set up the private network

1. Install the required network interface cards (NICs).

   Create aggregated interfaces if you want to use these to set up private network.

2. Connect the InfoScale Availability private Ethernet controllers on each system.
3. Use crossover Ethernet cables, switches, or independent hubs for each InfoScale Availability communication network. Note that the crossover Ethernet cables are supported only on two systems.

   Ensure that you meet the following requirements:

   - The power to the switches or hubs must come from separate sources.
   - On each system, you must use two independent network cards to provide redundancy.
   - If a network interface is part of an aggregated interface, you must not configure the network interface under LLT. However, you can configure the aggregated interface under LLT.

- When you configure Ethernet switches for LLT private interconnect, disable the spanning tree algorithm on the ports used for the interconnect.

During the process of setting up heartbeat connections, consider a case where a failure removes all communications between the systems.

Note that a chance for data corruption exists under the following conditions:

- The systems still run, and
- The systems can access the shared storage.

4 Configure the Ethernet devices that are used for the private network such that the autonegotiation protocol is not used. You can achieve a more stable configuration with crossover cables if the autonegotiation protocol is not used.

To achieve this stable configuration, do one of the following:

- Edit the /etc/system file to disable autonegotiation on all Ethernet devices system-wide.
- Create a qfe.conf or bge.conf file in the /kernel/drv directory to disable autonegotiation for the individual devices that are used for private network.

Refer to the Oracle Ethernet driver product documentation for information on these methods.

5 Test the network connections. Temporarily assign network addresses and use `telnet` or `ping` to verify communications.

LLT uses its own protocol, and does not use TCP/IP. So, you must ensure that the private network connections are used only for LLT communication and not for TCP/IP traffic. To verify this requirement, unplumb and unconfigure any temporary IP addresses that are configured on the network interfaces.

The installer configures the private network in the cluster during configuration.

You can also manually configure LLT.

Optimizing LLT media speed settings on private NICs

For optimal LLT communication among the cluster nodes, the interface cards on each node must use the same media speed settings. Also, the settings for the switches or the hubs that are used for the LLT interconnections must match that of the interface cards. Incorrect settings can cause poor network performance or even network failure.

If you use different media speed for the private NICs, Veritas recommends that you configure the NICs with lesser speed as low-priority links to enhance LLT performance.

Guidelines for setting the media speed for LLT interconnects

Review the following guidelines for setting the media speed for LLT interconnects:

- Veritas recommends that you manually set the same media speed setting on each Ethernet card on each node.
  If you use different media speed for the private NICs, Veritas recommends that you configure the NICs with lesser speed as low-priority links to enhance LLT performance.
- If you have hubs or switches for LLT interconnects, then set the hub or switch port to the same setting as used on the cards on each node.

Details for setting the media speeds for specific devices are outside of the scope of this manual. Consult the device's documentation or the operating system manual for more information.

Synchronizing time settings on cluster nodes

Make sure that the time settings on all cluster nodes are synchronized. If the nodes are not in sync, timestamps for change (`ctime`) and modification (`mtime`) may not be consistent with the sequence in which operations actually happened.

For instructions, see the operating system documentation.

Creating a root user

On Oracle Solaris 11, you need to change the root role into a user as you cannot directly log in as root user.

To change root role into a user

1. Log in as local user and assume the root role.

```
% su  - root
```

2. Remove the root role from local users who have been assigned the role.

```
# roles admin

root

# usermod -R " " admin
```

3 Change the root role into a user.

```
# rolemod -K type=normal root
```

4 Verify the change.

- `# getent user_attr root`

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

  If the `type` keyword is not present in the output or is equal to normal, the account is not a role.

- `# userattr type root`

  If the output is empty or lists normal, the account is not a role.

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

Note: After installation, you may want to change root user into root role to allow local users to assume the root role.

Preparing zone environments

You need to keep the following items in mind when you install or upgrade VCS in a zone environment on an Oracle Solaris 10 operating system.

- When you install or upgrade InfoScale Availability using the `installer` program, all zones are upgraded (both global and non-global) unless they are detached and unmounted.
- Make sure that all non-global zones are booted and in the running state before you install or upgrade the InfoScale Availability packages in the global zone. If the non-global zones are not mounted and running at the time of upgrade, you must attach the zone with **-U** option to install or upgrade the InfoScale Availability packages inside the non-globle zone.
- If you install InfoScale Availability on Solaris 10 systems that run non-global zones, you need to make sure that non-global zones do not inherit the `/opt` directory. Run the following command to make sure that the `/opt` directory is not in the inherit-pkg-dir clause:

```
# zonecfg -z zone_name info
zonepath: /export/home/zone1
autoboot: false
pool: yourpool
inherit-pkg-dir:
dir: /lib
inherit-pkg-dir:
dir: /platform
inherit-pkg-dir:
dir: /sbin
inherit-pkg-dir:
dir: /usr
```

If the /opt directory appears in the output, remove the /opt directory from the zone's configuration and reinstall the zone.

Section 3

Installation of Veritas InfoScale Availability

- Chapter 5. Installing Veritas InfoScale using the installer
- Chapter 6. Installing Veritas InfoScale Availability using response files
- Chapter 7. Installing Veritas InfoScale using operating system-specific methods
- Chapter 8. Verifying the Veritas InfoScale installation

Chapter 5

Installing Veritas InfoScale using the installer

This chapter includes the following topics:

- Installing Veritas InfoScale using the installer
- Installing language packages

Installing Veritas InfoScale using the installer

The product installer is the recommended method to license and install Veritas InfoScale.

To install Veritas Infoscale

1. Load and mount the software disc. If you downloaded the software, navigate to the top level of the download directory and skip the next step.

2. Move to the top-level directory on the disc.

   ```
   # cd /mnt/cdrom
   ```

3. From this directory, type the following command to start the installation on the local system.

   ```
   # ./installer
   ```

4. Press **I** to install and press **Enter**.

5 The installer asks whether you wants to configure the product.

```
Would you like to configure InfoScale Availability after installation?
[y,n,q]
```

If you enter **y**, the installer configures the product after installation. If you enter **n**, the installer quits after the installation is complete.

6 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA).

```
Do you agree with the terms of the End User License Agreement as
specified in the EULA/en/EULA_InfoScale_Ux_7.0.pdf file
present on media? [y,n,q,?] y
```

7 Choose the licensing method. Answer the licensing questions and follow the prompts.

```
1)  Enter a valid license key
2)  Enable keyless licensing and complete system licensing later
How would you like to license the systems? [1-2,q] (2)
```

Note: You can also register your license using the installer menu by selecting the **L) License a Product** option.

---

8 Check the log file to confirm the installation. The log files, summary file, and response file are saved at: `/opt/VRTS/install/logs` directory.

Installing language packages

To install InfoScale Availability in a language other than English, install the required language packages after installing the English packages.

To install the language packages on the server

1 Insert the "Language" disc into the DVD-ROM or CD-ROM drive. With Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.

2 Install the language packages using the `install_lp` command.

```
# cd /cdrom/cdrom0
# ./install_lp
```

Chapter 6

Installing Veritas InfoScale Availability using response files

This chapter includes the following topics:

- About response files
- Installing InfoScale Availability using response files
- Response file variables to install Veritas InfoScale Availability
- Sample response file for Veritas InfoScale Availability installation

About response files

The installer script or product installation script generates a response file during any installation, configuration, upgrade, or uninstall procedure. The response file contains the configuration information that you entered during the procedure. When the procedure completes, the installation script displays the location of the response files.

You can use the response file for future installation procedures by invoking an installation script with the `-responsefile` option. The response file passes arguments to the script to automate the installation of that product. You can edit the file to automate installation and configuration of additional systems.

Note: Veritas recommends that you use the response file created by the installer and then edit it as per your requirement.

Syntax in the response file

The syntax of the Perl statements that is included in the response file variables varies. It can depend on whether the variables require scalar or list values.

For example, in the case of a string value:

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=["value 1 ", "value 2 ", "value 3 "];
```

Installing InfoScale Availability using response files

Typically, you can use the response file that the installer generates after you perform InfoScale Availability installation on a system to install InfoScale Availability on other systems..

To install InfoScale Availability using response files

1. Make sure the systems where you want to install InfoScale Availability meet the installation requirements.
2. Make sure that the preinstallation tasks are completed.
3. Copy the response file to the system where you want to install InfoScale Availability.
4. Edit the values of the response file variables as necessary.
5. Mount the product disc and navigate to the directory that contains the installation program.

6 Start the installation 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.

7 Complete the InfoScale Availability post-installation tasks.

For instructions, see the chapter *Performing post-installation tasks* in this document.

Response file variables to install Veritas InfoScale Availability

Table 6-1 lists the response file variables that you can define to install InfoScale Availability.

Table 6-1 Response file variables for installing InfoScale Availability

Variable	Description
CFG{opt}{install}	Installs InfoScale Availability packages. Configuration can be performed at a later time using the `-configure` option. List or scalar: scalar Optional or required: optional
CFG{activecomponent}	Specifies the component for operations like precheck, configure, addnode, install and configure(together). List or scalar: list Optional or required: required
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required

Table 6-1 Response file variables for installing InfoScale Availability *(continued)*

Variable	Description
CFG{keys}{vxkeyless} CFG{keys}{license}	`CFG{keys}{vxkeyless}` gives the list of keyless keys to be registered on the system. `CFG{keys}{license}` gives the list of user defined keys to be registered on the system List of Scalar: List 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{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required
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}{rsh}	Defines that *rsh* must be used instead of ssh as the communication method between systems. 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

Sample response file for Veritas InfoScale Availability installation

The following example shows a response file for installing Veritas InfoScale Availability.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{keys}{keyless}=[ qw(AVAILABILITY) ];
$CFG{opt}{gco}=1;
$CFG{opt}{install}=1;
$CFG{prod}="AVAILABILITY70";
$CFG{systems}=[ qw(system1 system2) ];

1;
```

Chapter 7

Installing Veritas InfoScale using operating system-specific methods

This chapter includes the following topics:

- About installing InfoScale Availability using operating system-specific methods
- Installing InfoScale Availability on Solaris 10 using JumpStart
- Using a Flash archive to install InfoScale Availability and the operating system
- Manually installing InfoScale Availability using the system command
- Manually installing packages on solaris10 brand zones

About installing InfoScale Availability using operating system-specific methods

On Solaris, you can install InfoScale Availability using the following methods:

- The procedure to manually install InfoScale Availability differs depending on the Solaris version.
- You can install InfoScale Availability on Solaris 10 systems using Solaris JumpStart.
- You can install InfoScale Availability using Flash archive on the Solaris 10 operating system.

Installing InfoScale Availability on Solaris 10 using JumpStart

This installation method applies only to Solaris 10. These JumpStart instructions assume a working knowledge of JumpStart. See the JumpStart documentation that came with your operating system for details on using JumpStart.

Upgrading is not supported. The following procedure assumes a standalone configuration.

For the language pack, you can use JumpStart to install packages. You add the language packages in the script, and put those files in the JumpStart server directory.

You can use a Flash archive to install InfoScale Availability and the operating system with JumpStart.

Overview of JumpStart installation tasks

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

Summary of tasks

1. Add a client (register to the JumpStart server). See the JumpStart documentation that came with your operating system for details.
2. Read the JumpStart installation instructions.
3. Generate the finish scripts.

   See “Generating the finish scripts” on page 46.
4. Prepare shared storage installation resources.

   See “Preparing installation resources” on page 47.
5. Modify the rules file for JumpStart.

   See the JumpStart documentation that came with your operating system for details.
6. Install the operating system using the JumpStart server.
7. When the system is up and running, run the installer command from the installation media to configure the Veritas InfoScale software.

   ```
   # /opt/VRTS/install/installer -configure
   ```

Generating the finish scripts

Perform these steps to generate the finish scripts to install InfoScale Availability.

To generate the script

1. Run the product installer program to generate the scripts for all products.

   ```
   ./installer -jumpstart directory_to_generate_scripts
   ```

2. JumpStart finish scripts are generated in the directory you specified in step 1.

   List the js_scripts directory.

   ```
   # ls /js_scripts
   ```

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

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

Preparing installation resources

Prepare resources for the JumpStart installation.

To prepare the resources

1. Copy the `pkgs` directory of the installation media to the shared storage.

   ```
   # cd /path_to_installation_media
   # cp -r pkgs BUILDSRC
   ```

2. Generate the response file with the list of packages.

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

3 Create the `adminfile` file under *`BUILDSRC`*`/pkgs/` directory.

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

4 If you want to encapsulate the root disk automatically when you perform the JumpStart installation, copy the scripts `encap_bootdisk_vm.fin` generated previously to *ENCAPSRC*.

Adding language pack information to the finish file

To add the language pack information to the finish file, perform the following procedure.

To add the language pack information to the finish file

1. For the language pack, copy the language packages from the language pack installation disc to the shared storage.

   ```
   # cd /cdrom/cdrom0/pkgs
   # cp -r * BUILDSRC/pkgs
   ```

   If you downloaded the language pack:

   ```
   # cd /path_to_language_pack_installation_media/pkgs
   # cp -r * BUILDSRC/pkgs
   ```

2. In the finish script, copy the product package information and replace the product packages with language packages.
3. The finish script resembles:

   ```
   . . .
   for PKG in product_packages
   do
   ...
   done. . .
   for PKG in language_packages
   do
   ...
   done. . .
   ```

Creating the Veritas InfoScale post-deployment scripts

The generated files vrts_deployment.sh and vrts_post-deployment.cf are customized Flash archive post-deployment scripts. These files clean up Veritas InfoScale product settings on a cloned system before you reboot it for the first time. Include these files in your Flash archives.

To create the post-deployment scripts

1. Mount the product disc.
2. From the prompt, run the `-flash_archive` option for the installer. Specify a directory where you want to create the files.

   ```
   # ./installer -flash_archive /tmp
   ```

3. Copy the vrts_postedeployment.sh file and the vrts_postedeployment.cf file to the golden system.

4 On the golden system perform the following:

- Put the vrts_postdeployment.sh file in the /etc/flash/postdeployment directory.
- Put the vrts_postdeployment.cf file in the /etc/vx directory.

5 Make sure that the two files have the following ownership and permissions:

```
# chown root:root /etc/flash/postdeployment/vrts_postdeployment.sh
# chmod 755 /etc/flash/postdeployment/vrts_postdeployment.sh
# chown root:root /etc/vx/vrts_postdeployment.cf
# chmod 644 /etc/vx/vrts_postdeployment.cf
```

Note that you only need these files in a Flash archive where you have installed Veritas InfoScale products.

Using a Flash archive to install InfoScale Availability and the operating system

You can only use Flash archive on the Solaris 10 operating system. In the following outline, refer to Solaris documentation for Solaris-specific tasks.

Note: Veritas does not support Flash Archive installation if the root disk of the master system is encapsulated.

The following is an overview of the creation and installation of a Flash archive with Veritas InfoScale software.

- If you plan to start flar (flash archive) creation from bare metal, perform step 1 through step 10.
- If you plan to start flar creation from a system where you have installed, but not configured the product, perform step 1 through step 4. Skip step 5 and finish step 6 through step 10.
- If you plan to start flar creation from a system where you have installed and configured the product, perform step 5 through step 10.

Flash archive creation overview

1 Ensure that you have installed Solaris 10 on the master system.

2 Use JumpStart to create a clone of a system.

3 Restart the cloned system.

4 Install the Veritas InfoScale products on the master system.

Perform one of the installation procedures from this guide.

5 If you have configured the product on the master system, create the `vrts_deployment.sh` file and the `vrts_deployment.cf` file and copy them to the master system.

6 Use the `flarcreate` command to create the Flash archive on the master system.

7 Copy the archive back to the JumpStart server.

8 Use JumpStart to install the Flash archive to the selected systems.

9 Configure the Veritas InfoScale product on all nodes in the cluster.

```
/opt/VRTS/install/installer -configure
```

10 Perform post-installation and configuration tasks.

See the *Configuration and Upgrade guides* for performing configuration tasks.

Manually installing InfoScale Availability using the system command

The procedure to manually install InfoScale Availability differs depending on the Solaris version.

Installing InfoScale Availability on Solaris 10 using the pkgadd command

On Solaris 10, the packages must be installed while in the global zone.

To install InfoScale Availability on Solaris 10 using the pkgadd command

1 Mount the software disc.

2 Copy the supplied VRTS* files from the installation media to a temporary location. Modify them if needed.

```
# cp /cdrom/cdrom0/pkgs/VRTS* \
    /tmp/pkgs
```

3 Create the admin file in the current directory. Specify the `-a adminfile` option when you use the `pkgadd` command:

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

4 Use the product-specific install command with one of the following options to get a list of packages in the order to be installed:

- minpkgs
- recpkgs
- allpkgs

5 Install the packages listed in step 4.

```
# pkgadd -a adminfile -d /tmp/pkgs pkgname.pkg
```

On Solaris 10, these packages must be installed while in the global zone. If a package's `pkginfo` file contains the variable SUNW_PKG_ALLZONES set not equal to true, the `-G` option should additionally be specified to the `pkgadd` command.

6 Verify that the packages are installed:

```
# pkginfo -l packagename
```

7 Start the processes.

Manually installing packages on solaris10 brand zones

You need to manually install InfoScale Availability 7.0.1 packages inside the solaris10 brand zones.

1. Boot the zone.
2. Logon to the solaris10 brand zone as a super user.
3. Copy the Solaris 10 packages from the pkgs directory from the installation media to the non-global zone (such as `/tmp/install` directory).
4. Install the following InfoScale Availability packages on the brand zone.

Note: Perform all the above steps on each Solaris 10 brand zone.

For more information on the support for Branded Zones, refer the *Veritas InfoScale™ 7.0 Virtualization Guide*.

Chapter 8

Verifying the Veritas InfoScale installation

This chapter includes the following topics:

- Verifying product installation
- Installation log files
- Setting environment variables
- Checking installed product versions and downloading maintenance releases and patches

Verifying product installation

Verify that the InfoScale Availability is installed.

You can verify the version of the installed product. Use the following command:

```
# /opt/VRTS/install/installer -version
```

Use the following sections to further verify the product installation.

Installation log files

After every product installation, the installer creates three text files:

- Installation log file
- Response file
- Summary file

The name and location of each file is displayed at the end of a product installation, and are always located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

Using the installation log file

The installation log file contains all commands executed during the procedure, their output, and errors generated by the commands. This file is for debugging installation problems and can be used for analysis by Veritas Support.

Using the summary file

The summary file contains the results of the installation by the installer or product installation scripts. The summary includes the list of the packages, and the status (success or failure) of each package. The summary also indicates which processes were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

Setting environment variables

Most of the commands which are used in the installation are present in the `/sbin` or `/usr/sbin` directory. Add these directories to your `PATH` environment variable as necessary.

After installation, InfoScale Availability commands are in `/opt/VRTS/bin`. InfoScale Availability manual pages are stored in `/opt/VRTS/man`.

Some VCS custom scripts reside in `/opt/VRTSvcs/bin`. If you want to install a high availability product, add `/opt/VRTSvcs/bin` to the `PATH` also.

Add the following directories to your `PATH` and `MANPATH` environment variable:

- If you want to use `Bourne` or `Korn` shell (`sh` or `ksh`), enter the following:

```
$ PATH=$PATH:/usr/sbin:/sbin:/usr/bin:/opt/VRTS/bin
$ MANPATH=/usr/share/man:/opt/VRTS/man:$MANPATH
$ export PATH MANPATH
```

- If you want to use a C shell (`csh` or `tcsh`), enter the following:

```
% set path = ( $path /usr/sbin /sbin/ /usr/bin/ /opt/VRTS/bin )
% setenv MANPATH /usr/share/man:/opt/VRTS/man:$MANPATH
```

Checking installed product versions and downloading maintenance releases and patches

Use the `installer` command with the `-version` option to:

- Determine the product packages that are installed on your system.
- Download required maintenance releases or patches .

The `version` option or the `showversion` script in the `/opt/VRTS/install` directory checks the specified systems and discovers the following:

- InfoScale Availability product version that is installed on the system
- All the required packages and the optional packages installed on the system
- Any required or optional packages (if applicable) that are not present
- Installed patches
- Available base releases (major or minor)
- Available maintenance releases
- Available patch releases

To check your systems and download maintenance releases and patches

1. Mount the media, or navigate to the installation directory.
2. Start the installer with the `-version` option.

   ```
   # ./installer -version sys1 sys2
   ```

   For each system, the installer lists all of the installed base releases, maintenance releases, and patches, followed by the lists of available downloads.
3. If you have Internet access, follow the prompts to download the available maintenance releases and patches to the local system.
4. If you do not have Internet access, you can download any needed maintenance releases and patches from the Services Operations Readiness Tools (SORT) Patch Finder page at:

   https://sort.veritas.com/patch/finder

   You can obtain installer patches automatically or manually.

Downloading maintenance releases and patches requires the installer to make outbound networking calls. You can also disable external network connection attempts.

Section 4

Upgrading Veritas InfoScale Availability

- Chapter 9. Preparing to upgrade to Veritas InfoScale Availability 7.0.1
- Chapter 10. Performing full upgrade of VCS with 2048 bit key and SHA256 signature certificates
- Chapter 11. Performing Online Upgrade
- Chapter 12. Upgrading InfoScale Availability using Live Upgrade
- Chapter 13. Performing InfoScale Availability upgrade using response files

Chapter 9

Preparing to upgrade to Veritas InfoScale Availability 7.0.1

This chapter includes the following topics:

- Supported upgrade types for Veritas InfoScale Availability 7.0.1
- Supported upgrade paths for Veritas InfoScale Availability 7.0.1

Supported upgrade types for Veritas InfoScale Availability 7.0.1

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

Table 9-1 lists the supported types of upgrade.

Table 9-1 Supported types of upgrade

Type of upgrade	Abstract
Full upgrade	A full upgrade involves upgrading all the nodes in the cluster at the same time. All components are upgraded during the process. The cluster remains unavailable for the duration of the upgrade.

Table 9-1 Supported types of upgrade *(continued)*

Type of upgrade	Abstract
Online upgrade	The online upgrade involves upgrading the whole cluster and supporting customer's application zero down time during the upgrade procedure.
Solaris Live Upgrade or Boot Environment upgrade	Solaris Live Upgrade or Boot Environment upgrade provides a method of upgrading a system while the system continues to operate.

Supported upgrade paths for Veritas InfoScale Availability 7.0.1

Table 9-2 lists the supported upgrade paths for Solaris 10 x64.

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

Table 9-2 Supported upgrade paths for Solaris 10 x64

Current version	Solaris 10
6.0.5	Upgrade OS to Sol 10 U9 or later. Upgrade to 7.0.1 using the `installer` script with product installer.

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

Chapter 10

Performing full upgrade of VCS with 2048 bit key and SHA256 signature certificates

This chapter includes the following topics:

- Stronger security with 2048 bit key and SHA256 signature certificates
- Upgrading VCS using the product installer
- Tasks to perform after upgrading to 2048 bit key and SHA256 signature certificates
- Upgrading Steward to 2048 bit key and SHA256 signature certificates

Stronger security with 2048 bit key and SHA256 signature certificates

Cluster Server (VCS) InfoScale Availability 7.0.1 uses 2048 bit key and SHA256 signature certificates. The vcsauthserver generates certificates with 2048 bit key and SHA256 signature. The enhancement provides stronger security to VCS users. The 2048 bit and SHA256 certificates are used by default during a fresh VCS InfoScale Availability 7.0.1 install. While upgrading VCS, you can either choose to upgrade the certificates during the VCS upgrade itself or you can upgrade the certificates later.

Upgrading VCS using the product installer

You can use the product installer to upgrade VCS.

To upgrade VCS using the product installer

1. Log in as superuser and mount the product disc.
2. On Solaris 10, if zones are configured, you need to set AutoStart attribute to 0 for zone groups, then make zone group offline.

   Set AutoStart attribute to 0 for zone groups:

   ```
   # haconf -makerw

   # hagrp -modify <zonegroup> AutoStart 0

   # haconf -dump -makero
   ```

3. Start the installer.

   ```
   # ./installer -sys -A sys -B
   ```

   The installer starts the product installation program with a copyright message. It then specifies where it creates the logs. Note the log's directory and name.
4. Select the product you want to upgrade.
5. Enter the names of the nodes that you want to upgrade. Use spaces to separate node names. Press the Enter key to proceed.
6. When the verification checks are complete, the installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.
7. The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

   If you are upgrading from 6.1 or earlier releases to 7.0.1, the following question is displayed:

   - Do you want to grant read access to everyone? [y,n,q,?]
     - To grant read access to all authenticated users, type **y**.
     - To grant usergroup specific permissions, type **n**.
   - Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
     - To specify usergroups and grant them read access, type **y**

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

- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

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

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

The installer lists the nodes that Veritas recommends you restart.

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

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

10 If you want to upgrade CP server systems that use VCS to InfoScale Availability 7.0.1, make sure that you first upgrade all application clusters to version InfoScale Availability 7.0.1. Then, upgrade VCS on the CP server systems.

For instructions to upgrade VCS, see the *Cluster Server Configuration and Upgrade Guide* or the *Storage Foundation and High Availability Configuration and Upgrade Guide*.

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

Sync non-global zone with the global:

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

Reset the AutoStart attribute to 1:

```
# haconf -makerw

# hagrp -modify <zonegroup> AutoStart 1

# haconf -dump -makero
```

Tasks to perform after upgrading to 2048 bit key and SHA256 signature certificates

Note that you must perform the following tasks after upgrading to 2048 bit key and SHA256 signature certificates:

- Delete certificates of non-root users after upgrading to 2048 bit key and SHA256 signature certificates.
- Re-establish WAC communication in global clusters after upgrading to 2048 bit key and SHA256 signature certificates.
  See "Re-establishing WAC communication in global clusters after upgrading to 2048 bit key and SHA256 signature certificates" on page 64.
- Re-establish communication between Steward and upgraded clusters.
  See "Re-establishing trust with Steward after upgrading to 2048 bit key and SHA256 signature certificates" on page 65.
- Re-establish trust between VOM and SFMH/VBS.
  See "Re-establish trust between VOM and SFMH/VBS after upgrading to 2048 bit key and SHA256 signature certificates" on page 66.

Note: The cluster communication between global clusters breaks if the cluster on one site is running VCS InfoScale Availability 7.0.1 and the other is on VCS version lower than 6.0.5. If you upgrade to 2048 bit key and SHA256 signature certificates in such a configuration, the communication will not be restored even after performing the essential tasks after the upgrade. The only workaround for this is to upgrade VCS to version 6.0.5 or above on clusters which are running on VCS versions lower than 6.0.5.

Deleting certificates of non-root users after upgrading to 2048 bit key and SHA256 signature certificates

After upgrading to 2048 bit key and SHA256 signature certificates, the certificates of non-root users become obsolete and must be deleted in order to allow non-root users to log in.

Perform the following steps to delete the non-root certificates for each non-root user:

1. Delete .VRTSat from the home directory of the users.

   ```
   # rm -rf /user_home_directory/.VRTSat
   ```

2. If the non-root user was not a cluster user prior to upgrade, add user to Cluster with appropriate Privilege.

   ```
   # hauser -add <user name> -priv <Privilege>
   ```

3. Log in with non-root user and create new certificates.

   ```
   # /opt/VRTS/bin/halogin <non_root_user> <password>
   ```

4. You must delete certificates related to Zone since they also become obsolete after the upgrade. You must delete them using the following command on each Zone.

   ```
   # rm -rf /var/VRTSvcs/vcsauth/data/CLIENT
   ```

5. Recreate the certificates using the following command.

   ```
   # /opt/VRTSvcs/bin/hawparsetup.pl <SG> <res> <WPAR> <passwd> <Systems>
   ```

   Where:

   - SG: Name of the service group you want to configure
   - res: Name of the WPAR resource in VCS configuration
   - WPAR: Name of the WPAR configured on the system
   - passwd: Password for the VCS user you want to create for WPAR
   - Systems: Names of systems on which service group can run

Re-establishing WAC communication in global clusters after upgrading to 2048 bit key and SHA256 signature certificates

During the upgrade, the vcsauthserver gets 2048 bit SHA256 certificates and the trust information gets deleted, which causes the WAC communication to break. To establish the communication again, you must set up trust for WAC on each node of every cluster. The remote site has to set up trust with the local site as a new broker certificate is created on the local site. The local site also has to set up trust with the remote site as the trust certificate gets deleted during the upgrade.

Perform the following steps to establish trust between the clusters:

1 On each node of the first cluster, run the following command:

```
# export EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/WAC;
/opt/VRTSvcs/bin/vcsat setuptrust -b
IP_address_of_any_node_from_the_second_cluster:14149 -s high
```

The command obtains and displays the security certificate and other details of the root broker of the second cluster. If the details are correct, enter **y** at the command prompt to establish trust.

For example: The hash of above credential is b36a2607bf48296063068e3fc49188596aa079bb

```
Do you want to trust the above?(y/n) y
```

2 On each node of the second cluster, run the following command:

```
# export EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/WAC;
/opt/VRTSvcs/bin/vcsat setuptrust -b
IP_address_of_any_node_from_the_first_cluster:14149 -s high
```

The command obtains and displays the security certificate and other details of the root broker of the first cluster. If the details are correct, enter **y** at the command prompt to establish trust.

Re-establishing trust with Steward after upgrading to 2048 bit key and SHA256 signature certificates

In case of Steward, when you upgrade the cluster to use the enhanced security, the vcsauthserver gets 2048 bit SHA256 certificates and the trust information gets deleted. This breaks the communication between cluster and Steward.

To reinstate the communication, you must setup trust between each node of the upgraded cluster and Steward.

1 Set up trust on all nodes of the GCO clusters:

```
# export EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/WAC
# vcsat setuptrust -b IP_of_Steward:14149 -s high
```

2 Set up trust on the Steward:

```
# export EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/STEWARD
# vcsat setuptrust -b VIP_of_upgraded_cluster:14149 -s high
```

Re-establish trust between VOM and SFMH/VBS after upgrading to 2048 bit key and SHA256 signature certificates

When clusters upgrade to OpenSSL 1.0.1 in InfoScale Availability 7.0.1, the trust setup between Storage Foundation Managed Hosts (SFMH) and VOM and the trust setup between Veritas Business Services (VBS) and VOM are broken. To setup the communication again, you must re-establish the trust for SFMH/VBS.

Perform the following steps on each node of the upgraded cluster to re-establish the trust between SFMH and VOM:

1 Create a temporary directory:

```
# mkdir -p /var/tmp
```

2 Setup trust with VOM and get the trusted certificate:

```
# EAT_DATA_DIR='/var/tmp' EAT_HOME_DIR='/opt/VRTSsfmh' \
# /opt/VRTSsfmh/bin/vssat setuptrust -b \
# <VOM_IP OR VOMHOST_NAME>:14145 -s low
```

3 Copy the trusted certificate to a desired location:

```
# cp /var/tmp/root/.VRTSat/profile/certstore/trusted/*.0 \
# /var/opt/VRTSsfmh/trusted/
```

4 Delete the temporary directory:

```
# rm -rf /var/tmp/root/.VRTSat
```

If the cluster is part of VBS, perform the following steps on each node of the upgraded cluster to re-establish the trust between VBS and VOM:

1. Get the VOM hostname:

```
# /opt/VRTSvbs/bin/vbsdeploy.pl -verify
```

2. Setup trust with VOM and get the trusted certificate:

```
# EAT_DATA_DIR='/var/VRTSvbs/vbsauth/data' EAT_HOME_DIR='/opt/VRTSvbs' \
# /opt/VRTSvbs/bin/vssatbin setuptrust -b \
# <VOM_IP OR VOMHOST_NAME>:14145 -s low
```

3. Restart VBS to load the new trusted certificate:
   - Offline the vbsapp resource on each upgraded cluster:

```
# hares -offline vbsapp -sys <system_name>
```

   - Online the vbsapp resource on each upgraded cluster:

```
# hares -online vbsapp -sys <system_name>
```

Upgrading Steward to 2048 bit key and SHA256 signature certificates

To upgrade Steward configured on Solaris 10 systems in secure mode:

1. Log on to the Steward system as a root user.
2. Stop the Steward process.

```
# steward -stop -secure
```

3. Uninstall the existing VRTS binaries.

4 Install the VRTSperl, VRTSvlic VRTSveki, VRTSvcs, VRTSllt, VRTSsfcpi, and VRTSvcsea binaries for InfoScale Availability InfoScale Availability 7.0.1.

Note: If you are upgrading from release prior to 7.0, you cannot directly upgrade to 7.0.1. You need to install 7.0 binaries and then 7.0.1 binaries.

---

5 Remove the old certificates using the following command:

```
# rm -rf /var/VRTSvcs/vcsauth/data/STEWARD
```

6 Run the install scripts:

```
# ./opt/VRTS/install/bin/UXRT70/add_install_scripts
```

7 Run the following command:

```
# ./opt/VRTS/install/installer -securityonenode
```

The installer prompts for a confirmation if InfoScale Availability is not configured or if InfoScale Availability is not running on all nodes of the cluster.

8 Enter **y** when the installer prompts whether you want to continue configuring security.

9 Run the following command:

```
# /opt/VRTSvcs/bin/steward_secure.pl
```

10 Set up trust on all nodes of the GCO clusters.

```
# export EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/WAC
# vcsat setuptrust -b IP_of_Steward:14149 -s high
```

11 Set up trust on the Steward for every GCO cluster.

```
# export EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/STEWARD
# vcsat setuptrust -b VIP_of_remote_cluster1:14149 -s high
# vcsat setuptrust -b VIP_of_remote_cluster2:14149 -s high
```

12 Start the Steward process.

```
# /opt/VRTSvcs/bin/steward -start -secure
```

Chapter 11

Performing Online Upgrade

This chapter includes the following topics:

- Limitations of online upgrade
- Upgrading VCS online using the installer
- Tasks to perform after upgrading to 2048 bit key and SHA256 signature certificates

Limitations of online upgrade

- Online upgrade is available only for VCS.
- The non-Veritas applications running on the node have zero down time during the online upgrade.
- VCS does not monitor the applications when online upgrade is in progress.
- For upgrades from VCS versions lower than 6.1 and CPS-based fencing is configured, upgrade the CP server before performing the online upgrade.

See “Upgrading VCS online using the installer” on page 69.

Upgrading VCS online using the installer

You can use the product installer to upgrade VCS online. The supported upgrade paths are same as those for the installer. The online upgrade uses 2048 bit key and SHA256 signature certificates. For more information,See “Stronger security with 2048 bit key and SHA256 signature certificates” on page 60.

To upgrade VCS online using the product installer

1. Log in as superuser and mount the product disc.
2. Start the installer.

   ```
   # ./installer
   ```

   The installer starts the product installation program with a copyright message. It then specifies where it creates the logs.

   Note the directory name and path where the logs get stored.
3. Choose **2** for **Online Upgrade** from the upgrade options.
4. After selecting the online upgrade method, enter any one system name from the cluster on which you want to perform the online upgrade.

   Even if you specify a single node from the cluster, the installer asks whether you want to perform online upgrade of VCS on the entire cluster, keeping your applications online. After you enter the system name, the installer performs some verification checks and asks the following question:

   ```
   Online upgrade supports application zero downtime.
   Would you like to perform online upgrade on the
   whole cluster? [y,n,q](y)
   ```
5. Enter **y** to initiate the online upgrade.

   Note: You can either exit the installer with the option **q** or cancel the upgrade using **n** and select any other cluster to upgrade in this step.

   The installer runs some verification checks on the nodes and subsequently asks if you agree with the terms of the End User License Agreement.
6. Enter **y** to agree and continue.
7. The installer displays the following question before the installer stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.
   - Do you want to grant read access to everyone? [y,n,q,?]
     - To grant read access to all authenticated users, type **y**.
     - To grant usergroup specific permissions, type **n**.
   - Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
     - To specify usergroups and grant them read access, type **y**

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

8 The installer asks if you want to stop VCS processes. Enter **y** to stop the VCS process.

It stops the VCS processes, uninstalls packages, reinstalls or upgrades packages, again configures VCS, and starts the processes.

9 The installer asks if you want to stop VCS processes. Enter **y** to stop the VCS process.

Tasks to perform after upgrading to 2048 bit key and SHA256 signature certificates

Note that you must perform the following tasks after upgrading to 2048 bit key and SHA256 signature certificates:

- Delete certificates of non-root users after upgrading to 2048 bit key and SHA256 signature certificates.
- Re-establish WAC communication in global clusters after upgrading to 2048 bit key and SHA256 signature certificates.
  See "Re-establishing WAC communication in global clusters after upgrading to 2048 bit key and SHA256 signature certificates" on page 64.
- Re-establish communication between Steward and upgraded clusters.
  See "Re-establishing trust with Steward after upgrading to 2048 bit key and SHA256 signature certificates" on page 65.
- Re-establish trust between VOM and SFMH/VBS.
  See "Re-establish trust between VOM and SFMH/VBS after upgrading to 2048 bit key and SHA256 signature certificates" on page 66.

Note: The cluster communication between global clusters breaks if the cluster on one site is running VCS InfoScale Availability 7.0.1 and the other is on VCS version lower than 6.0.5. If you upgrade to 2048 bit key and SHA256 signature certificates in such a configuration, the communication will not be restored even after performing the essential tasks after the upgrade. The only workaround for this is to upgrade VCS to version 6.0.5 or above on clusters which are running on VCS versions lower than 6.0.5.

Chapter 12

Upgrading InfoScale Availability using Live Upgrade

This chapter includes the following topics:

- Supported upgrade paths for Live Upgrade and Boot Environment upgrade
- Performing Live Upgrade on Solaris 10 systems

Supported upgrade paths for Live Upgrade and Boot Environment upgrade

The systems where you plan to use Live Upgrade must run Solaris 10.

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

Performing Live Upgrade in a Solaris zone environment on Solaris 10

If you have a zone root that reside on a VxVM volume, for the purpose of Live Upgrade, create another VxVM volume of same or bigger size than that of the existing zone root for copying the file system contents to alternate boot environment. Use VxVM commands for creating the volume.

Use the standard procedure for the other standby nodes.

See “Performing Live Upgrade on Solaris 10 systems” on page 73.

By default, Zone agent `BootState` is set to "multi-user." After you complete the upgrade, you may need to adjust this attribute to the appropriate value before you start your zone through VCS.

Note: Veritas recommends that you set `BootState` to "multi-user-server" to run applications inside non-global zones.

For Solaris 10, make sure that all non-global zones are either in the running or configured state before you use the Veritas product installer to upgrade the Storage Foundation products in the global zone. If the non-global zones are not mounted and running at the time of the upgrade, you must attach each non-global zone with update option manually after upgrade.

For Live Upgrade, if the alternative root environment also has a zone, you cannot install `VRTSodm`. You must remove the `VRTSodm` package first then install the Storage Foundation product. After you restart the alternative root, you can install `VRTSodm`.

Performing Live Upgrade on Solaris 10 systems

Perform the Live Upgrade using the installer. Live Upgrade uses 2048 bit key and SHA256 signature certificates. For more information,See “Stronger security with 2048 bit key and SHA256 signature certificates” on page 60.

Table 12-1 Upgrading InfoScale Availability using Solaris 10 Live Upgrade

Step	Description
Step 1	Prepare to upgrade using Solaris Live Upgrade. See “Before you upgrade InfoScale Availability using Solaris Live Upgrade” on page 74.
Step 2	Create a new boot environment on the alternate boot disk.

Table 12-1 Upgrading InfoScale Availability using Solaris 10 Live Upgrade *(continued)*

Step	Description
Step 3	Upgrade InfoScale Availability using the installer. See "Upgrading InfoScale Availability using the installer for Solaris 10 Live Upgrade" on page 75.
	To upgrade only Solaris See the Oracle documentation on Solaris 10 operating system **Note:** A new boot environment is created on the alternate boot disk by cloning the primary boot environment. If you choose to upgrade the operating system, the Solaris operating system on the alternate boot environment is upgraded.
Step 4	Switch the alternate boot environment to be the new primary. See "Completing the Solaris 10 Live Upgrade" on page 76.
Step 5	Perform the tasks after upgrading to 2048 bit key and SHA256 signature certificates. See "Tasks to perform after upgrading to 2048 bit key and SHA256 signature certificates" on page 77.
Step 6	Verify Live Upgrade of InfoScale Availability. See "Verifying the Solaris 10 Live Upgrade of InfoScale Availability" on page 78.

Before you upgrade InfoScale Availability using Solaris Live Upgrade

Before you upgrade, perform the following procedure.

To prepare for the Live Upgrade

1. Make sure that the InfoScale Availability installation media and the operating system installation images are available and on hand.
2. On the nodes to be upgraded, select an alternate boot disk that is at least the same size as the root partition of the primary boot disk
3. On the primary boot disk, patch the operating system for Live Upgrade.

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

- Remove the installed Live Upgrade packages for the current operating system version:
  All Solaris versions: `SUNWluu`, `SUNWlur` packages.
  Solaris 10 update 7 or later also requires: `SUNWlucfg` package.
  Solaris 10 zones or Branded zones also requires: `SUNWluzone` package.
- From the new Solaris installation image, install the new versions of the following Live Upgrade packages:
  All Solaris versions: `SUNWluu`, `SUNWlur`, and `SUNWlucfg` packages.
  Solaris 10 zones or Branded zones also requires: `SUNWluzone` package.

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

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

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

Creating a new Solaris 10 boot environment on the alternate boot disk

For more information about using live upgrade to create a boot environment, see http://docs.oracle.com/cd/E23823_01/html/E23801/lucreate-1.html#scrolltoc.

Upgrading InfoScale Availability using the installer for Solaris 10 Live Upgrade

You can use the Veritas product installer to upgrade InfoScale Availability as part of the Live Upgrade.

At the end of the process, InfoScale Availability 7.0.1 is installed on the alternate boot disk.

To perform Live Upgrade of InfoScale Availability using the installer

1. Insert the product disc with InfoScale Availability 7.0.1 or access your copy of the software on the network.
2. Run the installer script specifying the root path as the alternate boot disk:

   ```
   # ./installer -upgrade -rootpath /altroot.5.10
   ```

3 Enter the names of the nodes that you want to upgrade to InfoScale Availability 7.0.1.

Note: Make sure that the installed version of VxFS uses the disk layout version 6 or later. If you are on a previous disk layout version, upgrade the version before you proceed with the InfoScale Availability 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 boot disk is 7.0.1.

```
# pkginfo -R /altroot.5.10 -l VRTSpkgname
```

For example:

```
# pkginfo -R /altroot.5.10 -l VRTSvxvm
```

Review the installation logs at `/altroot.5.10/opt/VRTS/install/logs`.

Completing the Solaris 10 Live Upgrade

At the end of the process:

- If the original primary boot disk was encapsulated, the alternate boot disk is encapsulated.
- The alternate boot environment is activated.
- The system is booted from the alternate boot disk.

To complete the Live Upgrade

1 Complete the Live upgrade process. Enter the following command on all nodes in the cluster.

2 For Solaris x64 systems, restore the original `PrivNIC.cf` file:

```
# cp /mnt/etc/VRTSvcs/conf/config/PrivNIC.cf
/tmp/PrivNIC.cf.save
```

3 If you want to upgrade VVR, run the `vvr_upgrade_lu_start` command.

Note: Only run the `vvr_upgrade_lu_start` command when you are ready to restart the nodes and switch over to the alternate boot environment.

---

4 Restart the system. The boot environment 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 boot environment.

You can ignore the following error if it appears: Error: boot environment *<dest.13445>* already mounted on *</altroot.5.10>*.

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

5 After the alternate boot environment is activated, you can switch boot environments. If the root disk is encapsulated, refer to the procedure to switch the boot environments manually.

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

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

Tasks to perform after upgrading to 2048 bit key and SHA256 signature certificates

Note that you must perform the following tasks after upgrading to 2048 bit key and SHA256 signature certificates:

- Delete certificates of non-root users after upgrading to 2048 bit key and SHA256 signature certificates.
- Re-establish WAC communication in global clusters after upgrading to 2048 bit key and SHA256 signature certificates.
  See "Re-establishing WAC communication in global clusters after upgrading to 2048 bit key and SHA256 signature certificates" on page 64.
- Re-establish communication between Steward and upgraded clusters.
  See "Re-establishing trust with Steward after upgrading to 2048 bit key and SHA256 signature certificates" on page 65.
- Re-establish trust between VOM and SFMH/VBS.
  See "Re-establish trust between VOM and SFMH/VBS after upgrading to 2048 bit key and SHA256 signature certificates" on page 66.

Note: The cluster communication between global clusters breaks if the cluster on one site is running VCS InfoScale Availability 7.0.1 and the other is on VCS version lower than 6.0.5. If you upgrade to 2048 bit key and SHA256 signature certificates in such a configuration, the communication will not be restored even after performing the essential tasks after the upgrade. The only workaround for this is to upgrade VCS to version 6.0.5 or above on clusters which are running on VCS versions lower than 6.0.5.

Verifying the Solaris 10 Live Upgrade of InfoScale Availability

To ensure that Live Upgrade has completed successfully, verify that the system have booted from the alternate boot environment.

To verify that Live Upgrade completed successfully

1. Verify that the alternate boot environment is active.

   ```
   # lustatus
   ```

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

2. Perform other verification as required to ensure that the new boot environment is configured correctly.
3. In a zone environment, verify the zone configuration.

Administering boot environments in Solaris 10 Live Upgrade

Use the following procedures to perform relevant administrative tasks for boot environments.

Reverting to the primary boot environment on a Solaris 10 system

If the alternate boot environment fails to start, you can revert to the primary boot environment.

Start the system from the primary boot environment in the PROM monitor mode.

```
ok> boot disk0
```

where *disk0* is the primary boot disk.

Chapter 13

Performing InfoScale Availability upgrade using response files

This chapter includes the following topics:

- Upgrading InfoScale Availability using response files
- Response file variables to upgrade InfoScale Availability
- Sample response file for upgrading InfoScale Availability

Upgrading InfoScale Availability using response files

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

To perform automated InfoScale Availability upgrade

1. Make sure the systems where you want to upgrade InfoScale Availability 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 InfoScale Availability.

   See "Sample response file for upgrading InfoScale Availability" on page 82.

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

See “Response file variables to upgrade InfoScale Availability” on page 80.

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 InfoScale Availability

Table 13-1 lists the response file variables that you can define to upgrade InfoScale Availability.

Table 13-1 Response file variables specific to upgrading InfoScale Availability

Variable	List or Scalar	Description
CFG{opt}{upgrade}	Scalar	Upgrades InfoScale Availability packages. (Required)
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{systems}	List	List of systems on which the product is to be upgraded. (Required)
CFG{defaultaccess}	Scalar (optional)	Defines if the user chooses to grant read access for VCS cluster information to everyone.
CFG{key}	Scalar (optional)	Stores the keyless key you want to register.

Table 13-1 Response file variables specific to upgrading InfoScale Availability *(continued)*

Variable	List or Scalar	Description
CFG{vcs_allowcomms}	Scalar	Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{pkgpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product packages. The location must be accessible from all target systems. (Optional)
CFG{opt}{tmppath}	Scalar	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. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. **Note:** The installer copies the response files and summary files also to the specified *logpath* location. (Optional)
CFG{opt}{rsh}	Scalar	Defines that *rsh* must be used instead of ssh as the communication method between systems. (Optional)

Table 13-1 Response file variables specific to upgrading InfoScale Availability *(continued)*

Variable	List or Scalar	Description
CFG{opt}{online_upgrade}	Scalar	Set the value to 1 for online upgrades.

Sample response file for upgrading InfoScale Availability

Review the response file variables and their definitions.

See "Response file variables to upgrade InfoScale Availability" on page 80.

```
#
# Configuration Values:
#
our %CFG;

$CFG{accepteula}=1;
$CFG{secusrgrps}=qw{staff usergroup@hostname.cdc.veritas.com}
$CFG{vcs_allowcomms}=1;
$CFG{opt}{upgrade}=1;
$CFG{prod}="Availability701";
$CFG{opt}{online_upgrade}=1;
$CFG{systems}=[ qw( sys1 sys2) ];
1;
```

Section 5

Configuration of Veritas InfoScale Availability

- Chapter 14. Configuring InfoScale Availability
- Chapter 15. Configuring InfoScale Availability clusters for data integrity
- Chapter 16. Manually configuring InfoScale Availability
- Chapter 17. Manually configuring the clusters for data integrity

Chapter 14

Configuring InfoScale Availability

This chapter includes the following topics:

- Overview of tasks to configure InfoScale Availability using the product installer
- Starting the software configuration
- Specifying systems for configuration
- Configuring the cluster name
- Configuring private heartbeat links
- Configuring the virtual IP of the cluster
- Configuring InfoScale Availability in secure mode
- Setting up trust relationships for your InfoScale Availability cluster
- Configuring a secure cluster node by node
- Adding InfoScale Availability users
- Configuring SMTP email notification
- Configuring SNMP trap notification
- Configuring global clusters
- Completing the InfoScale Availability configuration

Overview of tasks to configure InfoScale Availability using the product installer

Table 14-1 lists the tasks that are involved in configuring SFHA using the script-based installer.

Table 14-1 Tasks to configure InfoScale Availability using the script-based installer

Task	Reference
Start the software configuration	See "Starting the software configuration" on page 85.
Specify the systems where you want to configure SFHA	See "Specifying systems for configuration" on page 86.
Configure the basic cluster	See "Configuring the cluster name" on page 87. See "Configuring private heartbeat links" on page 87.
Configure virtual IP address of the cluster (optional)	See "Configuring the virtual IP of the cluster" on page 90.
Configure the cluster in secure mode (optional)	See "Configuring InfoScale Availability in secure mode" on page 92.
Add InfoScale Availability users (required if you did not configure the cluster in secure mode)	See "Adding InfoScale Availability users" on page 99.
Configure SMTP email notification (optional)	See "Configuring SMTP email notification" on page 100.
Configure SNMP email notification (optional)	See "Configuring SNMP trap notification" on page 101.
Complete the software configuration	See "Completing the InfoScale Availability configuration" on page 103.

Starting the software configuration

You can configure SFHA using the product installer.

Note: If you want to reconfigure SFHA, before you start the installer you must stop all the resources that are under InfoScale Availability control using the `hastop` command or the `hagrp -offline` command.

To configure SFHA using the product installer

1. Confirm that you are logged in as a superuser.
2. Start the configuration using the installer.

   ```
   # /opt/VRTS/install/installer -configure
   ```

   The installer starts the product installation program with a copyright message and specifies the directory where the logs are created.
3. Select the component to configure.
4. Continue with the configuration procedure by responding to the installer questions.

Specifying systems for configuration

The installer prompts for the system names on which you want to configure SFHA. The installer performs an initial check on the systems that you specify.

To specify system names for configuration

1. Enter the names of the systems where you want to configure SFHA.

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

2. Review the output as the installer verifies the systems you specify.

   The installer does the following tasks:

   - Checks that the local node running the installer can communicate with remote nodes

     If the installer finds ssh binaries, it confirms that ssh can operate without requests for passwords or passphrases. If ssh binaries cannot communicate with remote nodes, the installer tries rsh binaries. And if both ssh and rsh binaries fail, the installer prompts to help the user to setup ssh or rsh binaries.
   - Makes sure that the systems are running with the supported operating system
   - Makes sure the installer started from the global zone

- Checks whether Veritas InfoScale Enterprise is installed
- Exits if Veritas InfoScale Enterprise7.0.1 is not installed

3 Review the installer output about the I/O fencing configuration and confirm whether you want to configure fencing in enabled mode.

```
Do you want to configure I/O Fencing in enabled mode? [y,n,q,?] (y)
```

Configuring the cluster name

Enter the cluster information when the installer prompts you.

To configure the cluster

1 Review the configuration instructions that the installer presents.

2 Enter a unique cluster name.

```
Enter the unique cluster name: [q,?] clus1
```

Configuring private heartbeat links

You now configure the private heartbeat links that LLT uses.

InfoScale Availability provides the option to use LLT over Ethernet or LLT over UDP (User Datagram Protocol). Veritas recommends that you configure heartbeat links that use LLT over Ethernet for high performance, unless hardware requirements force you to use LLT over UDP. If you want to configure LLT over UDP, make sure you meet the prerequisites.

You must not configure LLT heartbeat using the links that are part of aggregated links. For example, link1, link2 can be aggregated to create an aggregated link, aggr1. You can use aggr1 as a heartbeat link, but you must not use either link1 or link2 as heartbeat links.

The following procedure helps you configure LLT heartbeat links.

To configure private heartbeat links

1 Choose one of the following options at the installer prompt based on whether you want to configure LLT over Ethernet or LLT over UDP.

- Option 1: Configure the heartbeat links using LLT over Ethernet (answer installer questions)
  Enter the heartbeat link details at the installer prompt to configure LLT over Ethernet.

Skip to step 2.

- Option 2: Configure the heartbeat links using LLT over UDP (answer installer questions)
  Make sure that each NIC you want to use as heartbeat link has an IP address configured. Enter the heartbeat link details at the installer prompt to configure LLT over UDP. If you had not already configured IP addresses to the NICs, the installer provides you an option to detect the IP address for a given NIC.
  Skip to step 3.
- Option 3: Automatically detect configuration for LLT over Ethernet
  Allow the installer to automatically detect the heartbeat link details to configure LLT over Ethernet. The installer tries to detect all connected links between all systems.
  Skip to step 5.

Note: Option 3 is not available when the configuration is a single node configuration.

---

2 If you chose option 1, enter the network interface card details for the private heartbeat links.

The installer discovers and lists the network interface cards.

Answer the installer prompts. The following example shows different NICs based on architecture:

- For Solaris x64:
  You must not enter the network interface card that is used for the public network (typically e1000g0.)

```
Enter the NIC for the first private heartbeat link on sys1:
[b,q,?] e1000g1
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat link on sys1:
[b,q,?] e1000g2
Would you like to configure a third private heartbeat link?
[y,n,q,b,?](n)
```

3 If you chose option 2, enter the NIC details for the private heartbeat links. This step uses examples such as *private_NIC1* or *private_NIC2* to refer to the available names of the NICs.

```
Enter the NIC for the first private heartbeat
link on sys1: [b,q,?] private_NIC1
Do you want to use address 192.168.0.1 for the
first private heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the first private heartbeat
link on sys1: [b,q,?] (50000)
Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat
link on sys1: [b,q,?] private_NIC2
Do you want to use address 192.168.1.1 for the
second private heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the second private heartbeat
link on sys1: [b,q,?] (50001)
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) y
Enter the NIC for the low priority heartbeat
link on sys1: [b,q,?] (private_NIC0)
Do you want to use address 192.168.3.1 for
the low priority heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the low priority heartbeat
link on sys1: [b,q,?] (50004)
```

4 Choose whether to use the same NIC details to configure private heartbeat links on other systems.

```
Are you using the same NICs for private heartbeat links on all
systems? [y,n,q,b,?] (y)
```

If you want to use the NIC details that you entered for sys1, make sure the same NICs are available on each system. Then, enter **y** at the prompt.

For LLT over UDP, if you want to use the same NICs on other systems, you still must enter unique IP addresses on each NIC for other systems.

If the NIC device names are different on some of the systems, enter **n**. Provide the NIC details for each system as the program prompts.

5 If you chose option 3 , the installer detects NICs on each system and network links, and sets link priority.

If the installer fails to detect heartbeat links or fails to find any high-priority links, then choose option 1 or option 2 to manually configure the heartbeat links.

See step 2 for option 1, or step 3 for option 2 or step 5 for option 3.

6 Verify and confirm the information that the installer summarizes.

Configuring the virtual IP of the cluster

You can configure the virtual IP of the cluster to use to connect from the Cluster Manager (Java Console), Veritas InfoScale Operations Manager, or to specify in the RemoteGroup resource.

See the *Cluster Server Administrator's Guide* for information on the Cluster Manager.

See the *Cluster Server Bundled Agents Reference Guide* for information on the RemoteGroup agent.

To configure the virtual IP of the cluster

1 Review the required information to configure the virtual IP of the cluster.

2 When the system prompts whether you want to configure the virtual IP, enter `y`.

3 Confirm whether you want to use the discovered public NIC on the first system.

Do one of the following:

- If the discovered NIC is the one to use, press `Enter`.
- If you want to use a different NIC, type the name of a NIC to use and press `Enter`.

```
Active NIC devices discovered on sys1: bge0
Enter the NIC for Virtual IP of the Cluster to use on sys1:
[b,q,?](bge0)
```

4 Confirm whether you want to use the same public NIC on all nodes.

Do one of the following:

- If all nodes use the same public NIC, enter `y`.
- If unique NICs are used, enter `n` and enter a NIC for each node.

```
Is bge0 to be the public NIC used by all systems
[y,n,q,b,?] (y)
```

5 Enter the virtual IP address for the cluster.

You can enter either an IPv4 address or an IPv6 address.

For IPv4:

- Enter the virtual IP address.

```
Enter the Virtual IP address for the Cluster:
[b,q,?] 192.168.1.16
```

- Confirm the default netmask or enter another one:

```
Enter the netmask for IP 192.168.1.16: [b,q,?]
(255.255.240.0)
```

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:

      NIC: bge0
      IP: 192.168.1.16
      Netmask: 255.255.240.0

Is this information correct? [y,n,q] (y)
```

For IPv6

- Enter the virtual IP address.

```
Enter the Virtual IP address for the Cluster:
[b,q,?] 2001:454e:205a:110:203:baff:feee:10
```

- Enter the prefix for the virtual IPv6 address you provided. For example:

```
Enter the Prefix for IP
2001:454e:205a:110:203:baff:feee:10: [b,q,?] 64
```

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:

      NIC: bge0
      IP: 2001:454e:205a:110:203:baff:feee:10
      Prefix: 64

Is this information correct? [y,n,q] (y)
```

If you want to set up trust relationships for your secure cluster, refer to the following topics:

See "Setting up trust relationships for your InfoScale Availability cluster" on page 92.

See "Configuring a secure cluster node by node" on page 94.

Configuring InfoScale Availability in secure mode

Configuring InfoScale Availability in secure mode ensures that all the communication between the systems is encrypted and users are verified against security credentials. InfoScale Availability user names and passwords are not used when a cluster is running in secure mode.

To configure InfoScale Availability in secure mode

1. To install and configure InfoScale Availability in secure mode, run the command:

   ```
   # ./installer -security
   ```

2. The installer displays the following question before the installer stops the product processes:

   - Do you want to grant read access to everyone? [y,n,q,?]
     - To grant read access to all authenticated users, type **y**.
     - To grant usergroup specific permissions, type **n**.
   - Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
     - To specify usergroups and grant them read access, type **y**
     - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.

3. To verify the cluster is in secure mode after configuration, run the command:

   ```
   # haclus -value SecureClus
   ```

   The command returns 1 if cluster is in secure mode, else returns 0.

Setting up trust relationships for your InfoScale Availability cluster

If you need to use an external authentication broker for authenticating InfoScale Availability users, you must set up a trust relationship between InfoScale Availability

and the broker. For example, if Veritas InfoScale Operations Manager is your external authentication broker, the trust relationship ensures that InfoScale Availability accepts the credentials that VOM issues.

Perform the following steps to set up a trust relationship between your InfoScale Availability cluster and a broker.

To set up a trust relationship

1. Ensure that you are logged in as superuser on one of the nodes in the cluster.
2. Enter the following command:

   ```
   # /opt/VRTS/install/installer  -securitytrust
   ```

   The installer specifies the location of the log files. It then lists the cluster information such as cluster name, cluster ID, node names, and service groups.
3. When the installer prompts you for the broker information, specify the IP address, port number, and the data directory for which you want to establish trust relationship with the broker.

   ```
   Input the broker name of IP address: 15.193.97.204

   Input the broker port: (14545)
   ```

   Specify a port number on which broker is running or press Enter to accept the default port.

   ```
   Input the data directory to setup trust with: (/var/VRTSvcs/
   vcsauth/data/HAD)
   ```

   Specify a valid data directory or press Enter to accept the default directory.
4. The installer performs one of the following actions:
   - If you specified a valid directory, the installer prompts for a confirmation.

     ```
     Are you sure that you want to setup trust for the InfoScale Availabili
     with the broker 15.193.97.204 and port 14545? [y,n,q] y
     ```

     The installer sets up trust relationship with the broker for all nodes in the cluster and displays a confirmation.

     ```
     Setup trust with broker 15.193.97.204 on cluster node1
     ........Done

     Setup trust with broker 15.193.97.204 on cluster node2
     ........Done
     ```

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

- If you entered incorrect details for broker IP address, port number, or directory name, the installer displays an error. It specifies the location of the log files, summary file, and response file and exits.

Configuring a secure cluster node by node

For environments that do not support passwordless ssh or passwordless rsh, you cannot use the `-security` option to enable secure mode for your cluster. Instead, you can use the `-securityonenode` option to configure a secure cluster node by node. Moreover, to enable security in fips mode, use the `-fips` option together with `-securityonenode`.

Table 14-2 lists the tasks that you must perform to configure a secure cluster.

Table 14-2 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See "Configuring the first node" on page 94.
Configure security on the remaining nodes	See "Configuring the remaining nodes" on page 95.
Complete the manual configuration steps	See "Completing the secure cluster configuration" on page 96.

Configuring the first node

Perform the following steps on one node in your cluster.

To configure security on the first node

1. Ensure that you are logged in as superuser.
2. Enter the following command:

```
# /opt/VRTS/install/installer -securityonenode
```

The installer lists information about the cluster, nodes, and service groups. If InfoScale Availability is not configured or if InfoScale Availability is not running on all nodes of the cluster, the installer prompts whether you want to continue configuring security. It then prompts you for the node that you want to configure.

```
InfoScale Availability is not running on all systems in this cluster. All
must be in RUNNING state. Do you want to continue? [y,n,q] (n) y

1) Perform security configuration on first node and export
security configuration files.

2) Perform security configuration on remaining nodes with
security configuration files.

Select the option you would like to perform [1-2,q.?] 1
```

Warning: All InfoScale Availability configurations about cluster users are deleted when you configure the first node. You can use the `/opt/VRTSvcs/bin/hauser` command to create cluster users manually.

3. The installer completes the secure configuration on the node. It specifies the location of the security configuration files and prompts you to copy these files to the other nodes in the cluster. The installer also specifies the location of log files, summary file, and response file.
4. Copy the security configuration files from the location specified by the installer to temporary directories on the other nodes in the cluster.

Configuring the remaining nodes

On each of the remaining nodes in the cluster, perform the following steps.

To configure security on each remaining node

1. Ensure that you are logged in as superuser.
2. Enter the following command:

```
# /opt/VRTS/install/install -securityonenode
```

The installer lists information about the cluster, nodes, and service groups. If InfoScale Availability is not configured or if InfoScale Availability is not running on all nodes of the cluster, the installer prompts whether you want to continue configuring security. It then prompts you for the node that you want to configure. Enter **2**.

```
InfoScale Availability is not running on all systems in this cluster. All
must be in RUNNING state. Do you want to continue? [y,n,q] (n) y

1) Perform security configuration on first node and export
security configuration files.

2) Perform security configuration on remaining nodes with
security configuration files.

Select the option you would like to perform [1-2,q.?]  2
Enter the security conf file directory: [b]
```

The installer completes the secure configuration on the node. It specifies the location of log files, summary file, and response file.

Completing the secure cluster configuration

Perform the following manual steps to complete the configuration.

To complete the secure cluster configuration

1. On the first node, freeze all service groups except the ClusterService service group.

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

   # /opt/VRTSvcs/bin/hagrp -list Frozen=0

   # /opt/VRTSvcs/bin/hagrp -freeze groupname -persistent


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

2. On the first node, stop the InfoScale Availability engine.

   ```
   # /opt/VRTSvcs/bin/hastop -all -force
   ```

3. On all nodes, stop the CmdServer.

   ```
   # /opt/VRTSvcs/bin/CmdServer -stop
   ```

4 To grant access to all users, add or modify `SecureClus=1` and `DefaultGuestAccess=1` in the cluster definition.

For example:

To grant read access to everyone:

```
Cluster clus1 (
SecureClus=1
DefaultGuestAccess=1
)
```

Or

To grant access to only root:

```
Cluster clus1 (
SecureClus=1
)
```

Or

To grant read access to specific user groups, add or modify SecureClus=1 and GuestGroups={} to the cluster definition.

For example:

```
cluster clus1 (
SecureClus=1
GuestGroups={staff, guest}
```

5 Modify `/etc/VRTSvcs/conf/config/main.cf` file on the first node, and add `-secure` to the WAC application definition if GCO is configured.

For example:

```
Application wac (
                StartProgram = "/opt/VRTSvcs/bin/wacstart -secure"
                StopProgram = "/opt/VRTSvcs/bin/wacstop"
                MonitorProcesses = {"/opt/VRTSvcs/bin/wac -secure"}
                RestartLimit = 3
                )
```

6 On all nodes, create the `/etc/VRTSvcs/conf/config/.secure` file.

```
# touch /etc/VRTSvcs/conf/config/.secure
```

7 On the first node, start InfoScale Availability. Then start InfoScale Availability on the remaining nodes.

```
# /opt/VRTSvcs/bin/hastart
```

8 On all nodes, start CmdServer.

```
# /opt/VRTSvcs/bin/CmdServer
```

9 On the first node, unfreeze the service groups.

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

# /opt/VRTSvcs/bin/hagrp -list Frozen=1

# /opt/VRTSvcs/bin/hagrp -unfreeze groupname -persistent

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

Adding InfoScale Availability users

If you have enabled a secure InfoScale Availability cluster, you do not need to add InfoScale Availability users now. Otherwise, on systems operating under an English locale, you can add InfoScale Availability users at this time.

To add InfoScale Availability users

1 Review the required information to add InfoScale Availability users.

2 Reset the password for the Admin user, if necessary.

```
Do you wish to accept the default cluster credentials of
'admin/password'? [y,n,q] (y) n
Enter the user name: [b,q,?] (admin)
Enter the password:
Enter again:
```

3 To add a user, enter **y** at the prompt.

```
Do you want to add another user to the cluster? [y,n,q] (y)
```

4 Enter the user's name, password, and level of privileges.

```
Enter the user name: [b,q,?] smith
Enter New Password:*******

Enter Again:*******
Enter the privilege for user smith (A=Administrator, O=Operator,
G=Guest): [b,q,?] a
```

5 Enter `n` at the prompt if you have finished adding users.

```
Would you like to add another user? [y,n,q] (n)
```

6 Review the summary of the newly added users and confirm the information.

Configuring SMTP email notification

You can choose to configure InfoScale Availability to send event notifications to SMTP email services. You need to provide the SMTP server name and email addresses of people to be notified. Note that you can also configure the notification after installation.

Refer to the *Cluster Server Administrator's Guide* for more information.

To configure SMTP email notification

1 Review the required information to configure the SMTP email notification.

2 Specify whether you want to configure the SMTP notification.

If you do not want to configure the SMTP notification, you can skip to the next configuration option.

See "Configuring SNMP trap notification" on page 101.

3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the SMTP server's host name.

```
Enter the domain-based hostname of the SMTP server
(example: smtp.yourcompany.com): [b,q,?] smtp.example.com
```

- Enter the email address of each recipient.

```
Enter the full email address of the SMTP recipient
(example: user@yourcompany.com): [b,q,?] ozzie@example.com
```

- Enter the minimum security level of messages to be sent to each recipient.

```
Enter the minimum severity of events for which mail should be
sent to ozzie@example.com  [I=Information, W=Warning,
E=Error, S=SevereError]: [b,q,?] w
```

4 Add more SMTP recipients, if necessary.

- If you want to add another SMTP recipient, enter `y` and provide the required information at the prompt.

```
Would you like to add another SMTP recipient? [y,n,q,b] (n) y

Enter the full email address of the SMTP recipient
(example: user@yourcompany.com): [b,q,?] harriet@example.com

Enter the minimum severity of events for which mail should be
sent to harriet@example.com  [I=Information, W=Warning,
E=Error, S=SevereError]: [b,q,?] E
```

- If you do not want to add, answer **n**.

```
Would you like to add another SMTP recipient? [y,n,q,b] (n)
```

5 Verify and confirm the SMTP notification information.

```
SMTP Address: smtp.example.com
Recipient: ozzie@example.com receives email for Warning or
higher events
Recipient: harriet@example.com receives email for Error or
higher events

Is this information correct? [y,n,q] (y)
```

Configuring SNMP trap notification

You can choose to configure InfoScale Availability to send event notifications to SNMP management consoles. You need to provide the SNMP management console name to be notified and message severity levels.

Note that you can also configure the notification after installation.

Refer to the *Cluster Server Administrator's Guide* for more information.

To configure the SNMP trap notification

1. Review the required information to configure the SNMP notification feature of InfoScale Availability.
2. Specify whether you want to configure the SNMP notification.

   See "Configuring global clusters" on page 103.

3. Provide information to configure SNMP trap notification.

   Provide the following information:

   - Enter the SNMP trap daemon port.

     ```
     Enter the SNMP trap daemon port: [b,q,?] (162)
     ```

   - Enter the SNMP console system name.

     ```
     Enter the SNMP console system name: [b,q,?] sys5
     ```

   - Enter the minimum security level of messages to be sent to each console.

     ```
     Enter the minimum severity of events for which SNMP traps
     should be sent to sys5 [I=Information, W=Warning, E=Error,
     S=SevereError]: [b,q,?] E
     ```

4. Add more SNMP consoles, if necessary.

   - If you want to add another SNMP console, enter `y` and provide the required information at the prompt.

     ```
     Would you like to add another SNMP console? [y,n,q,b] (n) y
     Enter the SNMP console system name: [b,q,?] sys4
     Enter the minimum severity of events for which SNMP traps
     should be sent to sys4 [I=Information, W=Warning,
     E=Error, S=SevereError]: [b,q,?] S
     ```

   - If you do not want to add, answer `n`.

```
Would you like to add another SNMP console? [y,n,q,b] (n)
```

5 Verify and confirm the SNMP notification information.

```
SNMP Port: 162
Console: sys5 receives SNMP traps for Error or
higher events
Console: sys4 receives SNMP traps for SevereError or
higher events

Is this information correct? [y,n,q] (y)
```

Configuring global clusters

You can configure global clusters to link clusters at separate locations and enable wide-area failover and disaster recovery. The installer adds basic global cluster information to the InfoScale Availability configuration file. You must perform additional configuration tasks to set up a global cluster.

See the *Cluster Server Administrator's Guide* for instructions to set up SFHA global clusters.

To configure the global cluster option

1 Review the required information to configure the global cluster option.

2 Specify whether you want to configure the global cluster option.

If you skip this option, the installer proceeds to configure InfoScale Availability based on the configuration details you provided.

3 Provide information to configure this cluster as global cluster.

The installer prompts you for a NIC, a virtual IP address, and value for the netmask.

Completing the InfoScale Availability configuration

After you enter the InfoScale Availability configuration information, the installer prompts to stop the InfoScale Availability processes to complete the configuration process. The installer continues to create configuration files and copies them to each system. The installer also configures a cluster UUID value for the cluster at the end of the configuration. After the installer successfully configures InfoScale Availability, it restarts InfoScale Availability and its related processes.

To complete the InfoScale Availability configuration

1. If prompted, press Enter at the following prompt.

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

2. Review the output as the installer stops various processes and performs the configuration. The installer then restarts InfoScale Availability and its related processes.

3. Enter y at the prompt to send the installation information to Veritas.

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

4. After the installer configures VCS successfully, note the location of summary, log, and response files that installer creates.

   The files provide the useful information that can assist you with the configuration and can also assist future configurations.

summary file	Describes the cluster and its configured resources.
log file	Details the entire configuration.
response file	Contains the configuration information that can be used to perform secure or unattended installations on other systems.

Chapter 15

Configuring InfoScale Availability clusters for data integrity

This chapter includes the following topics:

- Setting up disk-based I/O fencing using installer
- Setting up non-SCSI-3 I/O fencing in virtual environments using installer
- Setting up majority-based I/O fencing using installer
- Enabling or disabling the preferred fencing policy

Setting up disk-based I/O fencing using installer

You can configure I/O fencing using the `-fencing` option of the installer.

Note: In InfoScale Availability 7.0.1, disk-based fencing on Solaris 10 x64 is for the co-existence of SF and Availability only.

Configuring disk-based I/O fencing using installer

Note: The installer stops and starts VCS to complete I/O fencing configuration. Make sure to unfreeze any frozen InfoScale Availability service groups in the cluster for the installer to successfully stop VCS.

To set up disk-based I/O fencing using the installer

1. Start the installer with `-fencing` option.

   ```
   # /opt/VRTS/install/installer  -fencing
   ```

   The installer starts with a copyright message and verifies the cluster information.

   Note the location of log files which you can access in the event of any problem with the configuration process.

2. Enter the host name of one of the systems in the cluster.

3. Confirm that you want to proceed with the I/O fencing configuration at the prompt.

   The program checks that the local node running the script can communicate with remote nodes and checks whether SFHA 7.0.1 is configured properly.

4. Review the output as the configuration program checks whether VxVM is already started and is running.

   - If the check fails, configure and enable VxVM before you repeat this procedure.
   - If the check passes, then the program prompts you for the coordinator disk group information.

5. Choose whether to use an existing disk group or create a new disk group to configure as the coordinator disk group.

   The program lists the available disk group names and provides an option to create a new disk group. Perform one of the following:

   - To use an existing disk group, enter the number corresponding to the disk group at the prompt.
     The program verifies whether the disk group you chose has an odd number of disks and that the disk group has a minimum of three disks.
   - To create a new disk group, perform the following steps:
     - Enter the number corresponding to the **Create a new disk group** option.
       The program lists the available disks that are in the CDS disk format in the cluster and asks you to choose an odd number of disks with at least three disks to be used as coordinator disks.
       Veritas recommends that you use three disks as coordination points for disk-based I/O fencing.
     - If the available VxVM CDS disks are less than the required, installer asks whether you want to initialize more disks as VxVM disks. Choose the disks you want to initialize as VxVM disks and then use them to create new disk group.

- Enter the numbers corresponding to the disks that you want to use as coordinator disks.
- Enter the disk group name.

6 Verify that the coordinator disks you chose meet the I/O fencing requirements.

You must verify that the disks are SCSI-3 PR compatible using the vxfentsthdw utility and then return to this configuration program.

See "Checking shared disks for I/O fencing" on page 110.

7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.

8 Verify and confirm the I/O fencing configuration information that the installer summarizes.

9 Review the output as the configuration program does the following:

- Stops InfoScale Availability and I/O fencing on each node.
- Configures disk-based I/O fencing and starts the I/O fencing process.
- Updates the InfoScale Availability configuration file main.cf if necessary.
- Copies the /etc/vxfenmode file to a date and time suffixed file /etc/vxfenmode-*date-time*. This backup file is useful if any future fencing configuration fails.
- Updates the I/O fencing configuration file /etc/vxfenmode.
- Starts InfoScale Availability on each node to make sure that the SFHA is cleanly configured to use the I/O fencing feature.

10 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.

11 Configure the Coordination Point Agent.

```
Do you want to configure Coordination Point Agent on
the client cluster? [y,n,q] (y)
```

12 Enter a name for the service group for the Coordination Point Agent.

```
Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen) vxfen
```

13 Set the level two monitor frequency.

```
Do you want to set LevelTwoMonitorFreq? [y,n,q] (y)
```

14 Decide the value of the level two monitor frequency.

```
Enter the value of the LevelTwoMonitorFreq attribute: [b,q,?] (5)
```

Installer adds Coordination Point Agent and updates the main configuration file.

Refreshing keys or registrations on the existing coordination points for disk-based fencing using the installer

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies InfoScale Availability about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points for disk-based I/O fencing using the installer

1 Start the installer with the `-fencing` option.

```
# /opt/VRTS/install/installer  -fencing
```

The installer starts with a copyright message and verifies the cluster information.

Note down the location of log files that you can access if there is a problem with the configuration process.

2 Confirm that you want to proceed with the I/O fencing configuration at the prompt.

The program checks that the local node running the script can communicate with the remote nodes and checks whether SFHA 7.0.1 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type the number corresponding to refresh registrations or keys on the existing coordination points.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-6,q]
```

4 Ensure that the disk group constitution that is used by the fencing module contains the same disks that are currently used as coordination disks.

5 Verify the coordination points.

```
For example,
Disk Group: fendg
Fencing disk policy: dmp
Fencing disks:
 emc_clariion0_62
 emc_clariion0_65
 emc_clariion0_66
```

Is this information correct? [y,n,q] **(y)**.

```
Successfully completed the vxfenswap operation
```

The keys on the coordination disks are refreshed.

6 Do you want to send the information about this installation to us to help improve installation in the future? [y,n,q,?] **(y)**.

7 Do you want to view the summary file? [y,n,q] **(n)**.

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

1 List the new external disks or the LUNs as recognized by the operating system. On each node, enter:

```
# vxdisk list
```

2 To initialize the disks as VxVM disks, use one of the following methods:

- Use the interactive vxdiskadm utility to initialize the disks as VxVM disks. For more information, see the *Storage Foundation Administrator's Guide*.
- Use the `vxdisksetup` command to initialize a disk as a VxVM disk.

```
# vxdisksetup -i device_name
```

The example specifies the CDS format:

```
# vxdisksetup -i c2t13d0
```

Repeat this command for each disk you intend to use as a coordinator disk.

Checking shared disks for I/O fencing

Make sure that the shared storage you set up while preparing to configure InfoScale Availability meets the I/O fencing requirements. You can test the shared disks using the vxfentsthdw utility. The two nodes must have `ssh` (default) or `rsh` communication. To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxfenadm` command with the `-i` option. This command option verifies that the same serial number for the LUN is returned on all paths to the LUN.

Make sure to test the disks that serve as coordinator disks.

The vxfentsthdw utility has additional options suitable for testing many disks. Review the options for testing the disk groups (`-g`) and the disks that are listed in a file (`-f`). You can also test disks without destroying data using the `-r` option.

See the *InfoScale Availability Administrator's Guide*.

Checking that disks support SCSI-3 involves the following tasks:

- Verifying the Array Support Library (ASL)
  See "Verifying Array Support Library (ASL)" on page 110.
- Verifying that nodes have access to the same disk
  See "Verifying that the nodes have access to the same disk" on page 111.
- Testing the shared disks for SCSI-3
  See "Testing the disks using vxfentsthdw utility" on page 112.

Verifying Array Support Library (ASL)

Make sure that the Array Support Library (ASL) for the array that you add is installed.

To verify Array Support Library (ASL)

1. If the Array Support Library (ASL) for the array that you add is not installed, obtain and install it on each node before proceeding.

   The ASL for the supported storage device that you add is available from the disk array vendor or Veritas technical support.

2. Verify that the ASL for the disk array is installed on each of the nodes. Run the following command on each node and examine the output to verify the installation of ASL.

   The following output is a sample:

   ```
   # vxddladm listsupport all


   LIBNAME             VID                 PID
   ============================================================
   libvx3par.so        3PARdata            VV
   libvxCLARiiON.so    DGC                 All
   libvxFJTSYe6k.so    FUJITSU             E6000
   libvxFJTSYe8k.so    FUJITSU             All
   libvxap.so          Oracle                 All
   libvxatf.so         VERITAS             ATFNODES
   libvxcompellent.so  COMPELNT            Compellent Vol
   libvxcopan.so       COPANSYS            8814, 8818
   ```

3. Scan all disk drives and their attributes, update the VxVM device list, and reconfigure DMP with the new devices. Type:

   ```
   # vxdisk scandisks
   ```

   See the Veritas Volume Manager documentation for details on how to add and configure disks.

Verifying that the nodes have access to the same disk

Before you test the disks that you plan to use as shared data storage or as coordinator disks using the vxfentsthdw utility, you must verify that the systems see the same disk.

To verify that the nodes have access to the same disk

1. Verify the connection of the shared storage for data to two of the nodes on which you installed Veritas InfoScale Enterprise.

2. Ensure that both nodes are connected to the same disk during the testing. Use the vxfenadm command to verify the disk serial number.

```
# vxfenadm -i diskpath
```

Refer to the `vxfenadm` (1M) manual page.

For example, an EMC disk is accessible by the /dev/rdsk/c1t1d0s2 path on node A and the /dev/rdsk/c2t1d0s2 path on node B.

From node A, enter:

```
# vxfenadm -i /dev/rdsk/c1t1d0s2

Vendor id : EMC
Product id : SYMMETRIX
Revision : 5567
Serial Number : 42031000a
```

The same serial number information should appear when you enter the equivalent command on node B using the /dev/rdsk/c2t1d0s2 path.

On a disk from another manufacturer, Hitachi Data Systems, the output is different and may resemble:

```
Vendor id       : HITACHI
Product id      : OPEN-3
Revision        : 0117
Serial Number   : 0401EB6F0002
```

Testing the disks using vxfentsthdw utility

This procedure uses the /dev/rdsk/c1t1d0s2 disk in the steps.

If the utility does not show a message that states a disk is ready, the verification has failed. Failure of verification can be the result of an improperly configured disk array. The failure can also be due to a bad disk.

If the failure is due to a bad disk, remove and replace it. The vxfentsthdw utility indicates a disk can be used for I/O fencing with a message resembling:

```
The disk /dev/rdsk/c1t1d0s2 is ready to be configured for I/O Fencing on
node sys1
```

For more information on how to replace coordinator disks, refer to the *InfoScale Availability Administrator's Guide*.

To test the disks using vxfentsthdw utility

1. Make sure system-to-system communication functions properly.

2. From one node, start the utility.

3 The script warns that the tests overwrite data on the disks. After you review the overview and the warning, confirm to continue the process and enter the node names.

Warning: The tests overwrite and destroy data on the disks unless you use the `-r` option.

```
******** WARNING!!!!!!!! ********
THIS UTILITY WILL DESTROY THE DATA ON THE DISK!!

Do you still want to continue : [y/n] (default: n) y
Enter the first node of the cluster: sys1
Enter the second node of the cluster: sys2
```

4 Review the output as the utility performs the checks and reports its activities.

5 If a disk is ready for I/O fencing on each node, the utility reports success for each node. For example, the utility displays the following message for the node sys1.

```
The disk is now ready to be configured for I/O Fencing on node
sys1

ALL tests on the disk /dev/rdsk/c1t1d0s2 have PASSED
The disk is now ready to be configured for I/O fencing on node
sys1
```

6 Run the vxfentsthdw utility for each disk you intend to verify.

Note: Only dmp disk devices can be used as coordinator disks.

Setting up non-SCSI-3 I/O fencing in virtual environments using installer

If you have installed Veritas InfoScale Enterprise in virtual environments that do not support SCSI-3 PR-compliant storage, you can configure non-SCSI-3 fencing.

To configure I/O fencing using the installer in a non-SCSI-3 PR-compliant setup

1. Start the installer with `-fencing` option.

   ```
   # /opt/VRTS/install/install  -fencing
   ```

   The installer starts with a copyright message and verifies the cluster information.

2. Confirm that you want to proceed with the I/O fencing configuration at the prompt.

   The program checks that the local node running the script can communicate with remote nodes and checks whether SFHA 7.0.1 is configured properly.

3. For server-based fencing, review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

   ```
   Select the fencing mechanism to be configured in this
   Application Cluster
   [1-7,q] 1
   ```

4. Enter **n** to confirm that your storage environment does not support SCSI-3 PR.

   ```
   Does your storage environment support SCSI3 PR?
   [y,n,q] (y) n
   ```

5. Confirm that you want to proceed with the non-SCSI-3 I/O fencing configuration at the prompt.

6. For server-based fencing, enter the number of CP server coordination points you want to use in your setup.

7. For server-based fencing, enter the following details for each CP server:

   - Enter the virtual IP address or the fully qualified host name.
   - Enter the port address on which the CP server listens for connections. The default value is 443. You can enter a different port address. Valid values are between 49152 and 65535.

   The installer assumes that these values are identical from the view of the SFHA cluster nodes that host the applications for high availability.

8. For server-based fencing, verify and confirm the CP server information that you provided.

9. Verify and confirm the SFHA cluster configuration information.

   Review the output as the installer performs the following tasks:

- Updates the CP server configuration files on each CP server with the following details for only server-based fencing, :
  - Registers each node of the SFHA cluster with the CP server.
  - Adds CP server user to the CP server.
  - Adds SFHA cluster to the CP server user.
- Updates the following configuration files on each node of the SFHA cluster
  - `/etc/vxfenmode` file
  - `/etc/default/vxfen` file
  - `/etc/vxenviron` file
  - `/etc/llttab` file
  - `/etc/vxfentab` (only for server-based fencing)

10 Review the output as the installer stops SFHA on each node, starts I/O fencing on each node, updates the InfoScale Availability configuration file main.cf, and restarts SFHA with non-SCSI-3 fencing.

For server-based fencing, confirm to configure the CP agent on the SFHA cluster.

11 Confirm whether you want to send the installation information to us.

12 After the installer configures I/O fencing successfully, note the location of summary, log, and response files that installer creates.

The files provide useful information which can assist you with the configuration, and can also assist future configurations.

Setting up majority-based I/O fencing using installer

You can configure majority-based fencing for the cluster using the installer .

Perform the following steps to confgure majority-based I/O fencing

1 Start the installer with the -fencing option.

```
# /opt/VRTS/install/installer  -fencing
```

Where *version* is the specific release version. The installer starts with a copyright message and verifies the cluster information.

Note: Make a note of the log file location which you can access in the event of any issues with the configuration process.

---

2 Confirm that you want to proceed with the I/O fencing configuration at the prompt. The program checks that the local node running the script can communicate with remote nodes and checks whether InfoScale Availability is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type **3** to configure majority-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,b,q] 3
```

Note: The installer will ask the following question. Does your storage environment support SCSI3 PR? [y,n,q,?] Input 'y' if your storage environment supports SCSI3 PR. Other alternative will result in installer configuring non-SCSI3 fencing(NSF).

---

4 The installer then populates the /etc/vxfenmode file with the appropriate details in each of the application cluster nodes.

```
Updating /etc/vxfenmode file on sys1 ..... Done
Updating /etc/vxfenmode file on sys2 ..... Done
```

5 Review the output as the installer stops and restarts the InfoScale Availability and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

6 Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

7 Verify the fencing configuration.

```
# vxfenadm -d
```

Enabling or disabling the preferred fencing policy

You can enable or disable the preferred fencing feature for your I/O fencing configuration.

You can enable preferred fencing to use system-based race policy, group-based race policy, or site-based policy. If you disable preferred fencing, the I/O fencing configuration uses the default count-based race policy.

Preferred fencing is not applicable to majority-based I/O fencing.

To enable preferred fencing for the I/O fencing configuration

1. Make sure that the cluster is running with I/O fencing set up.

   ```
   # vxfenadm -d
   ```

2. Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

   ```
   # haclus -value UseFence
   ```

3. To enable system-based race policy, perform the following steps:

   - Make the InfoScale Availability configuration writable.

     ```
     # haconf -makerw
     ```

   - Set the value of the cluster-level attribute PreferredFencingPolicy as System.

     ```
     # haclus -modify PreferredFencingPolicy System
     ```

   - Set the value of the system-level attribute FencingWeight for each node in the cluster.
     For example, in a two-node cluster, where you want to assign sys1 five times more weight compared to sys2, run the following commands:

     ```
     # hasys -modify sys1 FencingWeight 50
     # hasys -modify sys2 FencingWeight 10
     ```

   - Save the InfoScale Availability configuration.

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

   - Verify fencing node weights using:

     ```
     # vxfenconfig -a
     ```

4 To enable group-based race policy, perform the following steps:

- Make the InfoScale Availability configuration writable.

  ```
  # haconf -makerw
  ```

- Set the value of the cluster-level attribute PreferredFencingPolicy as Group.

  ```
  # haclus -modify PreferredFencingPolicy Group
  ```

- Set the value of the group-level attribute Priority for each service group. For example, run the following command:

  ```
  # hagrp -modify service_group Priority 1
  ```

  Make sure that you assign a parent service group an equal or lower priority than its child service group. In case the parent and the child service groups are hosted in different subclusters, then the subcluster that hosts the child service group gets higher preference.

- Save the InfoScale Availability configuration.

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

5 To enable site-based race policy, perform the following steps:

- Make the InfoScale Availability configuration writable.

  ```
  # haconf -makerw
  ```

- Set the value of the cluster-level attribute PreferredFencingPolicy as Site.

  ```
  # haclus -modify PreferredFencingPolicy Site
  ```

- Set the value of the site-level attribute Preference for each site.

  ```
  For example,
  # hasite -modify Pune Preference 2
  ```

- Save the InfoScale Availability configuration.

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

6 To view the fencing node weights that are currently set in the fencing driver, run the following command:

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

1. Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

2. Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

3. To disable preferred fencing and use the default race policy, set the value of the cluster-level attribute PreferredFencingPolicy as Disabled.

```
# haconf -makerw
# haclus -modify PreferredFencingPolicy Disabled
# haconf -dump -makero
```

Chapter 16

Manually configuring InfoScale Availability

This chapter includes the following topics:

- About configuring InfoScale Availability manually
- Configuring LLT manually
- Configuring GAB manually
- Configuring InfoScale Availability manually
- Configuring InfoScale Availability in single node mode
- Starting LLT, GAB, and InfoScale Availability after manual configuration
- About configuring cluster using InfoScale Availability Cluster Configuration wizard
- Before configuring a InfoScale Availability cluster using the InfoScale Availability Cluster Configuration wizard
- Launching the InfoScale Availability Cluster Configuration wizard
- Configuring a cluster by using the InfoScale Availability cluster configuration wizard
- Adding a system to a InfoScale Availability cluster
- Modifying the InfoScale Availability configuration

About configuring InfoScale Availability manually

This section describes the procedures to manually configure InfoScale Availability.

Note: For manually configuring InfoScale Availability in single node mode, you can skip steps about configuring LLT manually and configuring GAB manually.

Configuring LLT manually

InfoScale Availability uses the Low Latency Transport (LLT) protocol for all cluster communications as a high-performance, low-latency replacement for the IP stack. LLT has two major functions.

It handles the following tasks:

- Traffic distribution
- Heartbeat traffic

To configure LLT over Ethernet, perform the following steps on each node in the cluster:

- Set up the file /etc/llthosts.
  See "Setting up /etc/llthosts for a manual installation" on page 121.
- Set up the file /etc/lltab.
  See "Setting up /etc/lltab for a manual installation" on page 122.
- Edit the following file on each node in the cluster to change the values of the LLT_START and the LLT_STOP environment variables to 1:
  /etc/default/llt

You can also configure LLT over UDP.

Setting up /etc/llthosts for a manual installation

The file llthosts(4) is a database. It contains one entry per system that links the LLT system ID (in the first column) with the LLT host name. You must ensure that contents of this file are identical on all the nodes in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.

Use vi or another editor, to create the file /etc/llthosts that contains the entries that resemble:

```
0 sys1
1 sys2
```

Setting up /etc/llttab for a manual installation

The /etc/llttab file must specify the system's ID number (or its node name), its cluster ID, and the network links that correspond to the system. In addition, the file can contain other directives. Refer also to the sample llttab file in /opt/VRTSllt.

See "About LLT directives in /etc/llttab file" on page 123.

Run the `dladm show-dev` command to query all NICs.

Use vi or another editor to create the file /etc/lltab that contains the entries that resemble the following:

- For x64:

```
set-node sys1
set-cluster 2
link e1000g0 /dev/e1000g:0 - ether - -
link e1000g1 /dev/e1000g:1 - ether - -
```

The first line must identify the system where the file exists. In the example, the value for `set-node` can be: sys1, 0, or the file name /etc/nodename. The file needs to contain the name of the system (sys1 in this example). The next line, beginning with the `set-cluster` command, identifies the cluster number, which must be a unique number when more than one cluster is configured on the same physical network connection.The next two lines, beginning with the `link` command, identify the two private network cards that the LLT protocol uses. The order of directives must be the same as in the sample `lltab` file in `/opt/VRTSllt`.

If you use different media speed for the private NICs, Veritas recommends that you configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example:

Use vi or another editor to create the file /etc/lltab that contains the entries that resemble the following:

- For x64:

```
set-node sys1
set-cluster 2
link e1000g0 /dev/e1000g:0 - ether - -
link e1000g1 /dev/e1000g:1 - ether - -
link-lowpri e1000g2 /dev/e1000g:2 - ether - -
```

LLT directives for a manual installation

Table 16-1 contains the LLT directives for a manual installation.

Table 16-1 LLT directives

Directive	Description
`set-node`	Assigns the system ID or symbolic name. The system ID number must be unique for each system in the cluster, and must be in the range 0-31. The symbolic name corresponds to the system ID, which is in /etc/llthosts file. Note that LLT fails to operate if any systems share the same ID.
`link`	Attaches LLT to a network interface. At least one link is required, and up to eight are supported. The first argument to link is a user-defined tag shown in the `lltstat(1M)` output to identify the link. It may also be used in `llttab` to set optional static MAC addresses. The second argument to link is the device name of the network interface. Its format is device_name:device_instance_number. The remaining four arguments to link are defaults; these arguments should be modified only in advanced configurations. There should be one link directive for each network interface. LLT uses an unregistered Ethernet SAP of 0xCAFE. If the SAP is unacceptable, refer to the `llttab(4)` manual page for information on how to customize SAP. Note that IP addresses do not need to be assigned to the network device; LLT does not use IP addresses.
`set-cluster`	Assigns a unique cluster number. Use this directive when more than one cluster is configured on the same physical network connection. LLT uses a default cluster number of zero.
`link-lowpri`	Use this directive in place of `link` for public network interfaces. This directive prevents InfoScale Availability communication on the public network until the network is the last link, and reduces the rate of heartbeat broadcasts. Note that LLT distributes network traffic evenly across all available network connections. In addition to enabling InfoScale Availability communication, it broadcasts heartbeats to monitor each network connection.

For more information about LLT directives, refer to the `llttab`(4) manual page.

About LLT directives in /etc/llttab file

Table 16-2 lists the LLT directives in /etc/llttab file for LLT over Ethernet.

Table 16-2 LLT directives

Directive	Description
`set-node`	Assigns the system ID or symbolic name. The system ID number must be unique for each system in the cluster, and must be in the range 0-63. The symbolic name corresponds to the system ID, which is in /etc/llthosts file. Note that LLT fails to operate if any systems share the same ID.
`set-cluster`	Assigns a unique cluster number. Use this directive when more than one cluster is configured on the same physical network connection. LLT uses a default cluster number of zero.
`link`	Attaches LLT to a network interface. At least one link is required, and up to eight are supported. LLT distributes network traffic evenly across all available network connections unless you mark the link as low-priority using the link-lowpri directive or you configured LLT to use destination-based load balancing. The first argument to link is a user-defined tag shown in the `lltstat(1M)` output to identify the link. It may also be used in `lltab` to set optional static MAC addresses. The second argument to link is the device name of the network interface. Its format is device_name:device_instance_number. The remaining four arguments to link are defaults; these arguments should be modified only in advanced configurations. There should be one link directive for each network interface. LLT uses an unregistered Ethernet SAP of 0xcafe. If the SAP is unacceptable, refer to the `lltab(4)` manual page for information on how to customize SAP. Note that IP addresses do not need to be assigned to the network device; LLT does not use IP addresses in LLT over Ethernet mode.
`link-lowpri`	Use this directive in place of `link` for public network interfaces. This directive prevents InfoScale Availability communication on the public network until the network is the last link, and reduces the rate of heartbeat broadcasts. If you use private NICs with different speed, use "link-lowpri" directive in place of "link" for all links with lower speed. Use the "link" directive only for the private NIC with higher speed to enhance LLT performance. LLT uses low-priority network links for InfoScale Availability communication only when other links fail.

For more information about the LLT directives, refer to the `lltab`(4) manual page.

Additional considerations for LLT for a manual installation

You must attach each network interface that is configured for LLT to a separate and distinct physical network.

By default, Oracle systems assign the same MAC address to all interfaces. Thus, connecting two or more interfaces to a network switch can cause problems. Consider the following example. You configure an IP on one public interface and LLT on another. Both interfaces are connected to a switch. The duplicate MAC address on the two switch ports can cause the switch to incorrectly redirect IP traffic to the LLT interface and vice versa. To avoid this issue, configure the system to assign unique MAC addresses by setting the eeprom(1M) parameter local-mac-address? to true.

Configuring GAB manually

InfoScale Availability uses the Group Membership Services/Atomic Broadcast (GAB) protocol for cluster membership and reliable cluster communications. GAB has two major functions.

It handles the following tasks:

- Cluster membership
- Cluster communications

To configure GAB

1 Set up an /etc/gabtab configuration file on each node in the cluster using vi or another editor. The following example shows an /etc/gabtab file:

```
/sbin/gabconfig -c -nN
```

Where the `-c` option configures the driver for use. The `-n`N option specifies that the cluster is not formed until at least N systems are ready to form the cluster. Veritas recommends that you set N to be the total number of systems in the cluster.

Warning: Veritas does not recommend the use of the `-c -x` option or `-x` option for `/sbin/gabconfig`. Using `-c -x` or `-x` can lead to a split-brain condition.

---

2 Edit the following file on each node in the cluster to change the values of the GAB_START and the GAB_STOP environment variables to 1:

/etc/default/gab

Configuring InfoScale Availability manually

InfoScale Availability configuration requires the types.cf and main.cf files on each system in the cluster. Both of the files are in the /etc/VRTSvcs/conf/config directory.

main.cf file — The main.cf configuration file requires the following minimum essential elements:

- An "include" statement that specifies the file, types.cf, which defines the InfoScale Availability bundled agent resource type definitions.
- The name of the cluster.
- The name of the systems that make up the cluster.

types.cf file — Note that the "include" statement in main.cf refers to the types.cf file. This text file describes the InfoScale Availability bundled agent resource type definitions. During new installations, the types.cf file is automatically copied in to the /etc/VRTSvcs/conf/config directory.

When you manually install InfoScale Availability, the file /etc/VRTSvcs/conf/config/main.cf contains only the line:

```
include "types.cf"
```

For a full description of the main.cf file, and how to edit and verify it, refer to the *Cluster Server Administrator's Guide*.

To configure InfoScale Availability manually

1 Log on as superuser, and move to the directory that contains the configuration file:

```
# cd /etc/VRTSvcs/conf/config
```

2 Use vi or another text editor to edit the main.cf file, defining your cluster name and system names. Refer to the following example.

An example main.cf for a two-node cluster:

```
include "types.cf"
cluster VCSCluster2 ( )
system sys1 ( )
system sys2 ( )
```

An example main.cf for a single-node cluster:

```
include "types.cf"
cluster VCSCluster1 ( )
system sn1 ( )
```

3 Save and close the main.cf file.

Configuring the cluster UUID when creating a cluster manually

You need to configure the cluster UUID when you manually create a cluster.

To configure the cluster UUID when you create a cluster manually

- On one node in the cluster, perform the following command to populate the cluster UUID on each node in the cluster.

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -configure nodeA
        nodeB ... nodeN
```

Where *nodeA*, *nodeB*, through *nodeN* are the names of the cluster nodes.

Configuring InfoScale Availability in single node mode

In addition to the steps mentioned in the manual configuration section, complete the following steps to configure InfoScale Availability in single node mode.

See “Configuring InfoScale Availability manually” on page 126.

To configure InfoScale Availability in single node mode

1. Disable the InfoScale Availability SMF service imported by VRTSvcs package.

   ```
   # svcadm disable -s system/vcs:default
   ```

2. Delete the InfoScale Availability SMF service configuration.

   ```
   # svccfg delete -f  system/vcs:default
   ```

3. Edit the following file to change the value of the ONENODE environment variable to **yes**.

   ```
   /etc/default/vcs
   ```

4. Import the SMF service for vcs-onenode.

   ```
   # svccfg import /etc/VRTSvcs/conf/vcs-onenode.xml
   ```

5. If the single node is intended only to manage applications, you can disable LLT, GAB, I/O fencing kernel modules.

   Note: Disabling InfoScale Availability kernel modules means that you cannot make the applications highly available across multiple nodes.

Starting LLT, GAB, and InfoScale Availability after manual configuration

After you have configured LLT, GAB, and InfoScale Availability, use the following procedures to start LLT, GAB, and InfoScale Availability.

To start LLT

1 On each node, run the following command to start LLT:

```
# svcadm enable llt
```

If LLT is configured correctly on each node, the console output resembles:

```
Jun 26 19:04:24 sys1 kernel: [1571667.550527] LLT INFO V-14-1-10009 LLT 6
```

2 On each node, run the following command to verify that LLT is running:

```
# /sbin/lltconfig
LLT is running
```

To start GAB

1 On each node, run the following command to start GAB:

```
# svcadm enable gab
```

If GAB is configured correctly on each node, the console output resembles:

```
Jun 26 19:10:34 sys1 kernel: [1572037.501731] GAB INFO
V-15-1-20021 GAB 6.0.100.000-SBLD available
```

2 On each node, run the following command to verify that GAB is running:

```
# /sbin/gabconfig -a
GAB Port Memberships
===================================
Port a gen a36e0003 membership 01
```

To start InfoScale Availability

- On each node, type:

```
# svcadm enable vcs
```

If InfoScale Availability is configured correctly on each node, the console output resembles:

```
Apr  5 14:52:02 sys1 gab: GAB:20036: Port h gen 3972a201
      membership 01
```

To start InfoScale Availability as single node

- Run the following command:

```
# svcadm enable vcs-onenode
```

About configuring cluster using InfoScale Availability Cluster Configuration wizard

Consider the following before configuring a cluster using InfoScale Availability Cluster Configuration wizard

- The InfoScale Availability Cluster Configuration wizard allows you to configure a InfoScale Availability cluster and add a node to the cluster.
  See "Configuring a cluster by using the InfoScale Availability cluster configuration wizard" on page 133.
- Veritas recommends that you first configure application monitoring using the wizard before using InfoScale Availability commands to add additional components or modify the existing configuration. Apart from configuring application availability, the wizard also sets up the other components required for successful application monitoring.

Before configuring a InfoScale Availability cluster using the InfoScale Availability Cluster Configuration wizard

Ensure that you complete the following tasks before launching the InfoScale Availability Cluster Configuration wizard to configure a InfoScale Availability cluster:

- Install InfoScale Availability or InfoScale Enterprise on the system on which you want to configure the InfoScale Availability cluster.
- You must have the following user privileges when you attempt to configure the InfoScale Availability cluster:
  - Configure Application Monitoring (Admin) privileges when you launch the wizard from the vSphere client.
  - Admin role privileges if you launch the wizard through VOM
- Install the application and the associated components that you want to monitor on the system.

- If you have configured a firewall, ensure that your firewall settings allow access to ports used by Cluster Server installer, wizards, and services.
  Verify that the following ports are not blocked by the firewall:

VMware environment	443, 5634, 14152, and 14153
Physical environment	5634, 14161, 14162, 14163, and 14164 At least one port from 14161, 14162, 14163, and 14164 must be open.

- You must not select bonded interfaces for cluster communication. A bonded interface is a logical NIC, formed by grouping several physical NICs together. All NICs in a bond have an identical MAC address, due to which you may experience the following issues:
  - Single Sign On (SSO) configuration failure.
  - The wizard may fail to discover the specified network adapters.
  - The wizard may fail to discover or validate the specified system name.
- The host name of the system must be resolvable through the DNS server or locally, using /etc/hosts file entries.

Launching the InfoScale Availability Cluster Configuration wizard

You must launch the InfoScale Availability Cluster Configuration wizard from the system where the disk residing on the shared datastore is attached.

You can launch the InfoScale Availability Cluster Configuration wizard from:

- VMware vSphere Client
  See Launching the InfoScale Availability Cluster Configuration wizard from VMware vSphere Client.
- A browser window
  See Launching the InfoScale Availability Cluster Configuration wizard from a browser window.

Launching the InfoScale Availability Cluster Configuration wizard from VMware vSphere Client

To launch the wizard from the VMware vSphere Client:

1. Launch the VMware vSphere Client and connect to the VMware vCenter Server that hosts the virtual machine.
2. From the vSphere Client's Inventory view in the left pane, select the virtual machine where you want to configure the InfoScale Availability cluster.
3. Select the **Symantec High Availability** tab.

   The tab displays various menus based on the what is configured on the system. The menu options launch the appropriate wizard panel based on the tasks that you choose to perform.

Launching the InfoScale Availability Cluster Configuration wizard from a browser window

You can launch the InfoScale Availability Cluster Configuration wizard from the Symantec High Availability view.

1. Open a browser window and enter the following URL:

   https://<IP_or_HostName>:5634/vcs/admin/application_health.html

   where <IP_or_HostName> is the IP address or host name of the system on which you want to configure the cluster.
2. Click the **Configure cluster** link on the Symantec High Availability view page to launch the wizard.

Note: At various stages of cluster configuration, the Symantec High Availability view offers different configuration options. These options launch appropriate wizard panels based on the tasks that you choose to perform.

See "Configuring a cluster by using the InfoScale Availability cluster configuration wizard" on page 133.

See "Adding a system to a InfoScale Availability cluster" on page 136.

Refer to the *Administering application monitoring from the Symantec High Availability view* section in *Cluster Server Administrator's Guide* for more information on the configurations possible from the Symantec High Availability view.

Configuring a cluster by using the InfoScale Availability cluster configuration wizard

Perform the following steps to configure a InfoScale Availability cluster by using the InfoScale Availability Cluster Configuration wizard.

To configure a InfoScale Availability cluster

1. Access the Symantec High Availability view (for any system belonging the required cluster).

   See "Launching the InfoScale Availability Cluster Configuration wizard" on page 131.

2. Review the information on the Welcome panel and click **Next**.

   The Configuration Inputs panel appears.

   The local system is by default selected as a cluster system.

3. If you do not want to add more systems to the cluster, skip this step. You can add systems later using the same wizard.

   To add a system to the cluster, click **Add System**.

   In the Add System dialog box, specify the following details for the system that you want to add to the InfoScale Availability cluster and click **OK**.

System Name or IP address	Specify the name or IP address of the system that you want to add to the InfoScale Availability cluster.
User name	Specify the user account for the system. Typically, this is the root user. The root user should have the necessary privileges.
Password	Specify the password for the user account you specified.
Use the specified user account on all systems	Select this check box to use the specified user account on all the cluster systems that have the same user name and password.

4. On the Configuration Inputs panel, do one of the following actions:
   - To add another system to the cluster, click Add System and repeat step 3.
   - To modify the specified User name or Password for a cluster system, use the edit icon.

- Click **Next**

5 If you do not want to modify the security settings for the cluster, click **Next**, and proceed to step 7.

By default, the wizard configures single sign-on for secure cluster communication. If you want to modify the security settings for the cluster, click **Advanced Settings**.

6 In the Advanced settings dialog box, specify the following details and click **OK**.

Use Single Sign-on	Select to configure single sign-on using InfoScale Availability Authentication Service for cluster communication. This option is enabled by default.
Use InfoScale Availability user privileges	Select to configure a user with administrative privileges to the cluster. Specify the username and password and click OK.

7 On the Network Details panel, select the type of network protocol to configure the InfoScale Availability cluster network links (Low Latency Transport or LLT module), and then specify the adapters for network communication.

The wizard configures the InfoScale Availability cluster communication links using these adapters. You must select a minimum of two adapters per cluster system.

Note: By default, the LLT links are configured over Ethernet.

Select **Use MAC address for cluster communication (LLT over Ethernet)** or select **Use IP address for cluster communication (LLT over UDP)**, and specify the following details for each cluster system.

- To configure LLT over Ethernet, select the adapter for each network communication link. You must select a different network adapter for each communication link.
- To configure LLT over UDP, select the type of IP protocol (IPv4 or IPv6), and then specify the required details for each communication link.

Network Adapter	Select a network adapter for the communication links. You must select a different network adapter for each communication link.
IP Address	Displays the IP address.
Port	Specify a unique port number for each link. For IPv4 and IPv6, the port range is from 49152 to 65535. A specified port for a link is used for all the cluster systems on that link.
Subnet mask (IPv4)	Displays the subnet mask details.
Prefix (IPv6)	Displays the prefix details.

By default, one of the links is configured as a low-priority link on a public network interface. The second link is configured as a high-priority link.

To change a high-priority link to a low-priority link, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

Note: Veritas recommends that you configure one of the links on a public network interface. You can assign the link on the public network interface as a low-priority link for minimal InfoScale Availability cluster communication over the link.

---

8 On the Configuration Summary panel, specify a cluster name and unique cluster ID and then click **Validate**.

Note: If multiple clusters exist in your network, the wizard validates if the specified cluster ID is a unique cluster ID among all clusters accessible from the current system. Among clusters that are not accessible from the current system, you must ensure that the cluster ID you specified is unique

---

9 Review the InfoScale Availability Cluster Configuration Details and then click **Next** to proceed with the configuration

10 On the Implementation panel, the wizard creates the InfoScale Availability cluster.

The wizard displays the status of the configuration task. After the configuration is complete, click **Next**.

If the configuration task fails, click **Diagnostic information** to check the details of the failure. Rectify the cause of the failure and run the wizard again to configure the InfoScale Availability cluster.

11 On the Finish panel, click **Finish** to complete the wizard workflow.

This completes the InfoScale Availability cluster configuration.

Adding a system to a InfoScale Availability cluster

Perform the following steps to add a system to a InfoScale Availability cluster by using the InfoScale Availability Cluster Configuration wizard.

The system from where you launch the wizard must be part of the cluster to which you want to add a new system.

To add a system to a InfoScale Availability cluster

1 Access the Symantec High Availability view (for any system belonging to the required cluster).

See "Launching the InfoScale Availability Cluster Configuration wizard" on page 131.

2 Click **Actions** > **Add System to InfoScale Availability Cluster**.

The InfoScale Availability Cluster Configuration Wizard is launched.

3 Review the information on the Welcome panel and click **Next**.

The Configuration Inputs panel appears, along with the cluster name, and a table of existing cluster systems.

4 To add a system to the cluster, click **Add System**.

5 In the Add System dialog box, specify the following details for the system that you want to add to the InfoScale Availability cluster and click **OK**.

System Name or IP address	Specify the name or IP address of the system that you want to add to the InfoScale Availability cluster.
User name	Specify the user account for the system. Typically, this is the root user. The root user should have the necessary privileges.
Password	Specify the password for the user account you specified.
Use the specified user account on all systems	Select this check box to use the specified user account on all the cluster systems that have the same user name and password.

6 On the Configuration Inputs panel, do one of the following actions:

- To add another system to the cluster, click **Add System** and repeat step 4.
- To modify the User name or Password for a cluster system, use the edit icon.
- Click **Next**

7 On the Network Details panel, specify the adapters for network communication (Low Latency Transport or LLT module of InfoScale Availability) for the system. The wizard configures the InfoScale Availability cluster communication links using these adapters. You must select a minimum of two adapters.

Note: You cannot modify the existing type of cluster communication (LLT over Ethernet or LLT over UDP).

- If the existing cluster uses LLT over Ethernet, select the adapter for each network communication link. You must select a different network adapter for each communication link.
- If the existing cluster uses LLT over UDP, select the type of IP protocol (IPv4 or IPv6), and then specify the required details for each communication link.

Network Adapter	Select a network adapter for the communication links. You must select a different network adapter for each communication link.
IP Address	Displays the IP address.
Port	Specify a unique port number for each link. For IPv4 and IPv6, the port range is from 49152 to 65535. A specified port for a link is used for all the cluster systems on that link.
Subnet mask (IPv4)	Displays the subnet mask details.
Prefix (IPv6)	Displays the prefix details.

By default, one of the links is configured as a low-priority link on a public network interface. The other link is configured as a high-priority link.

To change a high-priority link to a low-priority link, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

Note: Veritas recommends that you configure one of the links on a public network interface. You can assign the link on the public network interface as a low-priority link for minimal InfoScale Availability cluster communication over the link.

---

8 On the Configuration Summary panel, review the InfoScale Availability Cluster Configuration Details.

9 On the Implementation panel, the wizard creates the InfoScale Availability cluster.

The wizard displays the status of the configuration task. After the configuration is complete, click **Next**.

If the configuration task fails, click **Diagnostic information** to check the details of the failure. Rectify the cause of the failure and run the wizard again to add the required system to the InfoScale Availability cluster.

10 On the Finish panel, click **Finish** to complete the wizard workflow.

Modifying the InfoScale Availability configuration

After the successful installation of InfoScale Availability, you can modify the configuration of InfoScale Availability using several methods. You can dynamically modify the configuration from the command line, Veritas Operations Manager, or the Cluster Manager (Java Console). For information on management tools, refer to the *Cluster Server Administrator's Guide*.

You can also edit the main.cf file directly. For information on the structure of the main.cf file, refer to the *Cluster Server Administrator's Guide*.

Configuring the ClusterService group

When you have installed InfoScale Availability, and verified that LLT, GAB, and InfoScale Availability work, you can create a service group to include the optional features. These features include the InfoScale Availability notification components and the Global Cluster option. If you manually added InfoScale Availability to your cluster systems, you must manually create the ClusterService group. You can refer to the configuration examples of a system with a ClusterService group. See the *InfoScale Availability Administrator's Guide* for more information.

Chapter 17

Manually configuring the clusters for data integrity

This chapter includes the following topics:

- Setting up disk-based I/O fencing manually
- Setting up non-SCSI-3 fencing in virtual environments manually
- Setting up majority-based I/O fencing manually

Setting up disk-based I/O fencing manually

Table 17-1 lists the tasks that are involved in setting up I/O fencing.

Table 17-1

Task	Reference
Initializing disks as VxVM disks	See "Initializing disks as VxVM disks" on page 109.
Identifying disks to use as coordinator disks	See "Identifying disks to use as coordinator disks" on page 145.
Checking shared disks for I/O fencing	See "Checking shared disks for I/O fencing" on page 110.
Setting up coordinator disk groups	See "Setting up coordinator disk groups" on page 141.
Creating I/O fencing configuration files	See "Creating I/O fencing configuration files" on page 142.
Modifying SFHA configuration to use I/O fencing	See "Modifying InfoScale Availability configuration to use I/O fencing" on page 143.

Table 17-1 *(continued)*

Task	Reference
Configuring CoordPoint agent to monitor coordination points	
Verifying I/O fencing configuration	See "Verifying I/O fencing configuration" on page 144.

Setting up coordinator disk groups

From one node, create a disk group named vxfencoorddg. This group must contain three disks or LUNs. You must also set the coordinator attribute for the coordinator disk group. VxVM uses this attribute to prevent the reassignment of coordinator disks to other disk groups.

Note that if you create a coordinator disk group as a regular disk group, you can turn on the coordinator attribute in Volume Manager.

Refer to the *Storage Foundation Administrator's Guide* for details on how to create disk groups.

The following example procedure assumes that the disks have the device names c1t1d0s2, c2t1d0s2, and c3t1d0s2.

To create the vxfencoorddg disk group

1 On any node, create the disk group by specifying the device names:

```
# vxdg init vxfencoorddg c1t1d0s2 c2t1d0s2 c3t1d0s2
```

2 Set the coordinator attribute value as "on" for the coordinator disk group.

```
# vxdg -g vxfencoorddg set coordinator=on
```

3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

4 Import the disk group with the `-t` option to avoid automatically importing it when the nodes restart:

```
# vxdg -t import vxfencoorddg
```

5 Deport the disk group. Deporting the disk group prevents the coordinator disks from serving other purposes:

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

After you set up the coordinator disk group, you must do the following to configure I/O fencing:

- Create the I/O fencing configuration file /etc/vxfendg
- Update the I/O fencing configuration file /etc/vxfenmode

To update the I/O fencing files and start I/O fencing

1 On each nodes, type:

```
# echo "vxfencoorddg" > /etc/vxfendg
```

Do not use spaces between the quotes in the "vxfencoorddg" text.

This command creates the /etc/vxfendg file, which includes the name of the coordinator disk group.

2 On all cluster nodes specify the use of DMP disk policy in the `/etc/vxfenmode` file.

- `# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode`

3 To check the updated /etc/vxfenmode configuration, enter the following command on one of the nodes. For example:

```
# more /etc/vxfenmode
```

4 Ensure that you edit the following file on each node in the cluster to change the values of the VXFEN_START and the VXFEN_STOP environment variables to 1:

/etc/default/vxfen

Modifying InfoScale Availability configuration to use I/O fencing

After you add coordination points and configure I/O fencing, add the UseFence = SCSI3 cluster attribute to the InfoScale Availability configuration file /etc/VRTSvcs/conf/config/main.cf.

To modify InfoScale Availability configuration to enable I/O fencing

1. Save the existing configuration:

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

2. Stop InfoScale Availability on all nodes:

   ```
   # hastop -all
   ```

3. To ensure High Availability has stopped cleanly, run `gabconfig -a`.

   In the output of the commans, check that Port h is not present.

4. If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

   ```
   # svcadm disable -t vxfen
   ```

5. Make a backup of the main.cf file on all the nodes:

   ```
   # cd /etc/VRTSvcs/conf/config
   # cp main.cf main.orig
   ```

6. On one node, use vi or another text editor to edit the main.cf file. To modify the list of cluster attributes, add the UseFence attribute and assign its value as SCSI3.

   ```
   cluster clus1(
   UserNames = { admin = "cDRpdxPmHpzS." }
   Administrators = { admin }
   HacliUserLevel = COMMANDROOT
   CounterInterval = 5
   UseFence = SCSI3
   )
   ```

   Regardless of whether the fencing configuration is disk-based or server-based, the value of the cluster-level attribute UseFence is set to SCSI3.

7. Save and close the file.

8 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

9 Using rcp or another utility, copy the InfoScale Availability configuration file from a node (for example, sys1) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp sys1:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

Verifying I/O fencing configuration

Verify from the vxfenadm output that the SCSI-3 disk policy reflects the configuration in the /etc/vxfenmode file.

To verify I/O fencing configuration

1 On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the fencing mode is SCSI3 and the SCSI3 disk policy is dmp:

```
I/O Fencing Cluster Information:
================================

Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:

   * 0 (sys1)
   1 (sys2)

RFSM State Information:
   node 0 in state 8 (running)
   node 1 in state 8 (running)
```

2 Verify that the disk-based I/O fencing is using the specified disks.

```
# vxfenconfig -l
```

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See "Initializing disks as VxVM disks" on page 109.

Review the following procedure to identify disks to use as coordinator disks.

To identify the coordinator disks

1 List the disks on each node.

For example, execute the following commands to list the disks:

```
# vxdisk -o alldgs list
```

2 Pick three SCSI-3 PR compliant shared disks as coordinator disks.

See "Checking shared disks for I/O fencing" on page 110.

Setting up non-SCSI-3 fencing in virtual environments manually

To manually set up I/O fencing in a non-SCSI-3 PR compliant setup

1 Configure I/O fencing either in majority-based fencing mode with no coordination points or in server-based fencing mode only with CP servers as coordination points.

See "Setting up majority-based I/O fencing manually " on page 151.

2 Make sure that the SFHA cluster is online and check that the fencing mode is customized mode or majority mode.

```
# vxfenadm -d
```

3 Make sure that the cluster attribute UseFence is set to SCSI-3.

```
# haclus -value UseFence
```

4 On each node, edit the /etc/vxenviron file as follows:

```
data_disk_fencing=off
```

5 On each node, edit the /kernel/drv/vxfen.conf file as follows:

```
vxfen_vxfnd_tmt=25
```

6 On each node, edit the /etc/vxfenmode file as follows:

```
loser_exit_delay=55
vxfen_script_timeout=25
```

Refer to the sample /etc/vxfenmode file.

7 On each node, set the value of the LLT sendhbcap timer parameter value as follows:

- Run the following command:

```
lltconfig -T sendhbcap:3000
```

- Add the following line to the /etc/llttab file so that the changes remain persistent after any reboot:

```
set-timer senhbcap:3000
```

8 On any one node, edit the InfoScale Availability configuration file as follows:

- Make the InfoScale Availability configuration file writable:

```
# haconf -makerw
```

- For each resource of the type DiskGroup, set the value of the MonitorReservation attribute to 0 and the value of the Reservation attribute to NONE.

```
# hares -modify <dg_resource> MonitorReservation 0

# hares -modify <dg_resource> Reservation "NONE"
```

- Run the following command to verify the value:

```
# hares -list Type=DiskGroup MonitorReservation!=0

# hares -list Type=DiskGroup Reservation!="NONE"
```

The command should not list any resources.

- Modify the default value of the Reservation attribute at type-level.

```
# haattr -default DiskGroup Reservation "NONE"
```

- Make the InfoScale Availability configuration file read-only

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

9 Make sure that the UseFence attribute in the InfoScale Availability configuration file main.cf is set to SCSI-3.

10 To make these VxFEN changes take effect, stop and restart VxFEN and the dependent modules

- On each node, run the following command to stop InfoScale Availability:

  ```
  # svcadm disable -t vcs
  ```

- After InfoScale Availability takes all services offline, run the following command to stop VxFEN:

  ```
  # svcadm disable -t vxfen
  ```

- On each node, run the following commands to restart VxFEN and InfoScale Availability:

  ```
  # svcadm enable vxfen
  ```

Sample /etc/vxfenmode file for non-SCSI-3 fencing

```
#
# vxfen_mode determines in what mode InfoScale Availability I/O Fencing shoul
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=customized

# vxfen_mechanism determines the mechanism for customized I/O
# fencing that should be used.
#
# available options:
# cps      - use a coordination point server with optional script
#            controlled scsi3 disks
#
vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O fencing
```

```
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multi-pathing
#
scsi3_disk_policy=dmp

#
# Seconds for which the winning sub cluster waits to allow for the
# losing subcluster to panic & drain I/Os. Useful in the absence of
# SCSI3 based data disk fencing loser_exit_delay=55
#
# Seconds for which vxfend process wait for a customized fencing
# script to complete. Only used with vxfen_mode=customized
# vxfen_script_timeout=25

# security parameter is deprecated release 6.1 onwards since
# communication with CP server will always happen over HTTPS
# which is inherently secure. In pre-6.1 releases, it was used
# to configure secure communication to the cp server using
# VxAT (Veritas Authentication Service) available options:
# 0 - don't use Veritas Authentication Service for cp server
#   communication
# 1  - use Veritas Authentication Service for cp server
#   communication
security=1

#
# vxfen_honor_cp_order determines the order in which vxfen
# should use the coordination points specified in this file.
#
# available options:
# 0 - vxfen uses a sorted list of coordination points specified
# in this file, the order in which coordination points are specified
# does not matter.
#   (default)
# 1 - vxfen uses the coordination points in the same order they are
#    specified in this file

# Specify 3 or more odd number of coordination points in this file,
# each one in its own line. They can be all-CP servers, all-SCSI-3
# compliant coordinator disks, or a combination of CP servers and
```

```
# SCSI-3 compliant coordinator disks.
# Please ensure that the CP server coordination points are
# numbered sequentially and in the same order on all the cluster
# nodes.
#
# Coordination Point Server(CPS) is specified as follows:
#
#  cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames
# over different subnets, all of the IPs/names can be specified
# in a comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,
# ...,[<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
#  is the serial number of the CPS as a coordination point; must
#  start with 1.
# <vip>
#  is the virtual IP address of the CPS, must be specified in
#  square brackets ("[]").
# <vhn>
#  is the virtual hostname of the CPS, must be specified in square
#  brackets ("[]").
# <port>
#  is the port number bound to a particular <vip/vhn> of the CPS.
#  It is optional to specify a <port>. However, if specified, it
#  must follow a colon (":") after <vip/vhn>. If not specified, the
#  colon (":") must not exist after <vip/vhn>.
#
# For all the <vip/vhn>s which do not have a specified <port>,
# a default port can be specified as follows:
#
# port=<default_port>
#
#  Where <default_port> is applicable to all the <vip/vhn>s for which a
#  <port> is not specified. In other words, specifying <port> with a
#  <vip/vhn> overrides the <default_port> for that <vip/vhn>.
# If the <default_port> is not specified, and there are <vip/vhn>s for
# which <port> is not specified, then port number 14250 will be used
# for such <vip/vhn>s.
```

```
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[cps1.company.com]
#  cps2=[192.168.0.25]
# cps3=[cps2.company.com]:59999
#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [cps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
# - if default port 57777 were not specified, port 14250 would be
#  used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
#  vxfendg=<coordinator disk group name>
# Example:
#  vxfendg=vxfencoorddg
#
# Examples of different configurations:
#  1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
# cps1=[cps1.company.com]
# cps2=[cps2.company.com]
```

```
# cps3=[cps3.company.com]
# port=443
```

Setting up majority-based I/O fencing manually

Table 17-2 lists the tasks that are involved in setting up I/O fencing.

Task	Reference
Creating I/O fencing configuration files	Creating I/O fencing configuration files
Modifying InfoScale Availability configuration to use I/O fencing	Modifying InfoScale Availability configuration to use I/O fencing
Verifying I/O fencing configuration	Verifying I/O fencing configuration

Creating I/O fencing configuration files

To update the I/O fencing files and start I/O fencing

1. On all cluster nodes, run the following command

   ```
   # cp /etc/vxfen.d/vxfenmode_majority /etc/vxfenmode
   ```

2. To check the updated /etc/vxfenmode configuration, enter the following command on one of the nodes.

   ```
   # cat /etc/vxfenmode
   ```

3. Ensure that you edit the following file on each node in the cluster to change the values of the `VXFEN_START` and the `VXFEN_STOP` environment variables to 1.

   ```
   /etc/sysconfig/vxfen
   ```

Modifying InfoScale Availability configuration to use I/O fencing

After you configure I/O fencing, add the UseFence = SCSI3 cluster attribute to the InfoScale Availability configuration file /etc/VRTSvcs/conf/config/main.cf.

To modify InfoScale Availability configuration to enable I/O fencing

1. Save the existing configuration:

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

2. Stop InfoScale Availability on all nodes:

   ```
   # hastop -all
   ```

3. To ensure High Availability has stopped cleanly, run `gabconfig -a`.

   In the output of the commans, check that Port h is not present.

4. If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

   ```
   # svcadm disable -t vxfen
   ```

5. Make a backup of the main.cf file on all the nodes:

   ```
   # cd /etc/VRTSvcs/conf/config
   # cp main.cf main.orig
   ```

6. On one node, use vi or another text editor to edit the main.cf file. To modify the list of cluster attributes, add the UseFence attribute and assign its value as SCSI3.

   ```
   cluster clus1(
   UserNames = { admin = "cDRpdxPmHpzS." }
   Administrators = { admin }
   HacliUserLevel = COMMANDROOT
   CounterInterval = 5
   UseFence = SCSI3
   )
   ```

   For fencing configuration in any mode except the disabled mode, the value of the cluster-level attribute UseFence is set to SCSI3.

7. Save and close the file.

8 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

9 Using rcp or another utility, copy the InfoScale Availability configuration file from a node (for example, sys1) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp sys1:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

Verifying I/O fencing configuration

Verify from the vxfenadm output that the fencing mode reflects the configuration in the /etc/vxfenmode file.

To verify I/O fencing configuration

- On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the fencing mode is majority:

```
I/O Fencing Cluster Information:
================================

 Fencing Protocol Version: 201
 Fencing Mode: MAJORITY
 Cluster Members:

        * 0 (sys1)
          1 (sys2)

 RFSM State Information:
        node   0 in state  8 (running)
        node   1 in state  8 (running)
```

Sample /etc/vxfenmode file for majority-based fencing

```
#
# vxfen_mode determines in what mode InfoScale Availability I/O Fencing shoul
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# majority   - use majority based fencing
# disabled   - run the driver but don't do any actual fencing
#
# vxfen_mode=majority
```

Section 6

Uninstallation of Veritas InfoScale Availability

- Chapter 18. Uninstalling Veritas InfoScale Availability using the installer
- Chapter 19. Uninstalling Veritas InfoScale Availability using response files

Chapter 18

Uninstalling Veritas InfoScale Availability using the installer

This chapter includes the following topics:

- About removing Veritas InfoScale Availability
- Uninstalling InfoScale Availability packages using the product installer
- Uninstalling Veritas InfoScale Availability using the pkgrm command

About removing Veritas InfoScale Availability

This section covers uninstallation requirements and steps to uninstall the Veritas software.

Only users with superuser privileges can uninstall Veritas InfoScale Availability.

Warning: Failure to follow the instructions in the following sections may result in unexpected behavior.

Preparing to uninstall

Review the following before removing the Veritas software.

Remote uninstallation

You must configure remote communication to uninstall InfoScale Availability on remote systems. In a High Availability environment, you must meet the prerequisites to uninstall on all nodes in the cluster at one time.

The following prerequisites are required for remote uninstallation:

- Communication protocols must exist between systems. By default, the uninstall scripts use ssh.
- You must be able to execute ssh or rsh commands as superuser on all systems.
- The ssh or rsh must be configured to operate without requests for passwords or pass phrases

Uninstalling InfoScale Availability packages using the product installer

Use the following procedure to remove InfoScale Availability products.

Not all packages may be installed on your system depending on the choices that you made when you installed the software.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout version in InfoScale Availability 7.0.1 with a previous version of InfoScale Availability.

Language packages are uninstalled when you uninstall the English language packages.

To shut down and remove the installed InfoScale Availability packages

1. In an HA configuration, stop VCS processes on either the local system or all systems.

   To stop VCS processes on the local system:

   ```
   # hastop -local
   ```

   To stop VCS processes on all systems:

   ```
   # hastop -all
   ```

2. Move to the `/opt/VRTS/install` directory and run the uninstall script.

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

   ```
   # ./installer -uninstall
   ```

3. The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall InfoScale Availability.

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

4. The uninstall script prompts you to stop the product processes. If you respond yes, the processes are stopped and the packages are uninstalled.

   The uninstall script creates log files and displays the location of the log files.

5. Most packages have kernel components. In order to ensure complete removal, a system reboot is recommended after all packages have been removed.

6. To verify the removal of the packages, use the following commands:

   Solaris 10:

   ```
   # pkginfo | grep VRTS
   ```

7. In case the uninstallation fails to remove any of the VRTS packages, check the installer logs for the reason for failure or try to remove the packages manually using the following command:

   ```
   # pkgrm VRTSvcs
   ```

Uninstalling Veritas InfoScale Availability using the `pkgrm` command

Use the following procedure to uninstall Veritas InfoScale Availability using the `pkgrm` command.

If you want to uninstall Veritas InfoScale Availability using the `pkgrm` command, the packages must be removed in a specific order, or else the uninstallation fails. Removing the packages out of order results in some errors, including possible core dumps, although the packages are still removed.

To uninstall Veritas InfoScale Availability

1. Stop any Veritas daemons that are running.
2. Remove the packages in the following order:
   - For Veritas InfoScale Availability (Solaris 10):

```
# pkgrm VRTSvcswiz
# pkgrm VRTSvbs
# pkgrm VRTSsfmh
# pkgrm VRTSvcsea
# pkgrm VRTSat (if it exists)
# pkgrm VRTSvcsag
# pkgrm VRTScps
# pkgrm VRTSvcs
# pkgrm VRTSamf
# pkgrm VRTSvxfen
# pkgrm VRTSgab
# pkgrm VRTSllt
# pkgrm VRTSspt
# pkgrm VRTSsfcpi
# pkgrm VRTSperl
# pkgrm VRTSvlic
```

Uninstalling the language packages using the `pkgrm` command

If you want to remove only the language packages, you can do so with the `pkgrm` command.

If you use the product installer menu or the uninstallation script, you can remove the language packages along with the English packages.

To remove the language packages

- Use the `pkgrm` command to remove the appropriate packages.

  ```
  # pkgrm package_name package_name ...
  ```

  Because the packages do not contain any dependencies, you can remove them in any order.

Chapter 19

Uninstalling Veritas InfoScale Availability using response files

This chapter includes the following topics:

- Uninstalling InfoScale Availability using response files
- Response file variables to uninstall Veritas InfoScale Availability
- Sample response file for Veritas InfoScale Availability uninstallation

Uninstalling InfoScale Availability using response files

Typically, you can use the response file that the installer generates after you perform InfoScale Availability uninstallation on one system to uninstall InfoScale Availability on other systems.

To perform an automated uninstallation

1. Make sure that you meet the prerequisites to uninstall InfoScale Availability.
2. Copy the response file to the system where you want to uninstall InfoScale Availability.
3. Edit the values of the response file variables as necessary.

4 Start the uninstallation from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installer -responsefile
/tmp/response_file
```

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

Response file variables to uninstall Veritas InfoScale Availability

Table 19-1 lists the response file variables that you can define to configure InfoScale Availability.

Table 19-1 Response file variables for uninstalling InfoScale Availability

Variable	Description
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required
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

Table 19-1 Response file variables for uninstalling InfoScale Availability *(continued)*

Variable	Description
CFG{opt}{uninstall}	Uninstalls InfoScale Availability packages. List or scalar: scalar Optional or required: optional

Sample response file for Veritas InfoScale Availability uninstallation

The following example shows a response file for uninstalling Veritas InfoScale Availability

```
our %CFG;

$CFG{opt}{uninstall}=1;
$CFG{opt}{vr}=1;
$CFG{prod}="AVAILABILITY701";
$CFG{systems}=[ qw(system1 system2) ];

1;
```

Section 7

Installation reference

- Appendix A. Installation scripts
- Appendix B. Troubleshooting installation issues

Appendix A

Installation scripts

This appendix includes the following topics:

- Installation script options

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
-addnode	Adds a node to a high availability cluster.
-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.
-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.
-online_upgrade	Used to perform online upgrade. Using this option, the installer upgrades the whole cluster and also supports customer's application zero down time during the upgrade procedure. Now this option only supports VCS and ApplicationHA.
-patch_path	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-patch2_path	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.

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

Command Line Option	Function
-patch3_path	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch4_path	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch5_path	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
–jumpstart *dir_path*	Produces a sample finish file for Solaris JumpStart installation. The *dir_path* indicates the path to the directory in which to create the finish file.
–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 Services Operations Readiness Tool (SORT) in order to automatically obtain patch and release information updates.
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.

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

Command Line Option	Function
–pkginfo	Displays a list of packages and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkginfo option with the installer script to display VCS packages.
–pkgset	Discovers and displays the package group (minimum, recommended, all) and packages that are installed on the specified systems.
-pkgtable	Displays product's packages in correct installation order by group.
–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.
-requirements	The `-requirements` option displays required OS version, required packages and patches, file system space, and other system requirements in order to install the product.
–responsefile *response_file*	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.

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

Command Line Option	Function
–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.
-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.

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

Command Line Option	Function
-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.
-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.
-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.

Appendix B

Troubleshooting installation issues

This appendix includes the following topics:

- Restarting the installer after a failed connection
- About the VRTSspt package troubleshooting tools
- Incorrect permissions for root on remote system
- Inaccessible system

Restarting the installer after a failed connection

If an installation is killed because of a failed connection, you can restart the installer to resume the installation. The installer detects the existing installation. The installer prompts you whether you want to resume the installation. If you resume the installation, the installation proceeds from the point where the installation failed.

About the VRTSspt package troubleshooting tools

The VRTSspt package provides a group of tools for troubleshooting a system and collecting information on its configuration. If you install and use the VRTSspt package, it will be easier for Veritas Support to diagnose any issues you may have.

The tools can gather Veritas File System and Veritas Volume Manager metadata information and establish various benchmarks to measure file system and volume manager performance. Although the tools are not required for the operation of any Veritas InfoScale product, Veritas recommends installing them should a support case be needed to be opened with Veritas Support. Use caution when you use the VRTSspt package, and always use it in concert with Veritas Support.

Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Failed to setup rsh communication on 10.198.89.241:
'rsh 10.198.89.241 <command>' failed
Trying to setup ssh communication on 10.198.89.241.
Failed to setup ssh communication on 10.198.89.241:
Login denied

Failed to login to remote system(s) 10.198.89.241.
Please make sure the password(s) are correct and superuser(root)
can login to the remote system(s) with the password(s).
If you want to setup rsh on remote system(s), please make sure
rsh with command argument ('rsh <host> <command>') is not
denied by remote system(s).

Either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication

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

System verification did not complete successfully

The following errors were discovered on the systems:

The ssh permission denied on 10.198.89.241
rsh exited 1 on 10.198.89.241
either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

Suggested solution: You need to set up the systems to allow remote access using `ssh` or `rsh`.

Note: Remove remote shell permissions after completing the InfoScale Availability installation and configuration.

Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
 Verifying systems: 12% ..................................
 Estimated time remaining: 0:10 1 of 8
 Checking system communication ............................. Done
System verification did not complete successfully
The following errors were discovered on the systems:
cannot resolve hostname host1
Enter the  system names separated by spaces: q,? (host1)
```

Suggested solution: Verify that you entered the system name correctly; use the `ping`(1M) command to verify the accessibility of the host.

Index

Symbols

/etc/llttab
    LLT directives 123
2048 bit key and SHA256 signature certificates 60

A

about
    installation using operating system-specific methods 45
    response files 40
    Veritas InfoScale Availability 15
    Veritas InfoScale product licensing 16
    VRTSvlic package 21
    vxlicinstupgrade utility 20
adding
    ClusterService group 139
    system to InfoScale Availability cluster 136
    users 99
attributes
    UseFence 143, 151

B

bundled agents
    types.cf file 126

C

checking product versions 56
cluster configuration wizard
    about 130
    considerations 130
    launching 131
    launching from a browser window 132
    launching from vSphere Client 132
ClusterService group
    adding manually 139
commands
    gabconfig 125
    vxdisksetup (initializing disks) 109
configuration files
    types.cf 126
configuring
    GAB 125
    hardware 25
    LLT
        manual 121
    private network 30
    rsh 28
    ssh 28
    switches 30
configuring InfoScale Availability
    adding users 99
    event notification 100–101
    global clusters 103
    product installer 85
    starting 85
controllers
    private Ethernet 30
coordinator disks
    setting up 141
creating
    Flash archive 50
    post-deployment scripts 49
creating root user 34

D

delete
    non-root users 63
directives
    LLT 122–123
disabling
    external network connection attempts 29
disk space
    directories 25
    language pack 25
    required 25
disks
    adding and initializing 109
    coordinator 141
    testing with vxfentsthdw 110
    verifying node access 111
downloading maintenance releases and patches 56

E

eeprom
  parameters 30
Ethernet controllers 30

F

fibre channel 25
flarcreate 50
Flash archive 50
  post-deployment scripts 49

G

GAB
  manual configuration 125
  starting 129
gabconfig command 125
gabtab file
  creating 125
global clusters
  configuration 103

H

hardware
  configuring network and storage 25
Hardware requirements
  Veritas InfoScale 25
hubs 30

I

I/O fencing
  checking disks 110
  setting up 140
  shared storage 110
InfoScale Availability
  configuring 85
  coordinator disks 141
  starting 128–129
InfoScale Availability installation
  pre-installation tasks
    synchronizing time settings 34
    verifying systems 30
installation
  response file variables 42
  sample response file 44
  Veritas InfoScale 38
installation script options 165
installer patches
  obtaining either manually or automatically 28
installing
  InfoScale Availability using operating system-specific methods 45
  JumpStart 46
  language packages 39
  post 103
  required disk space 25
  using Flash archive 50
  using response files 41
  using the system command 51
ISO image
  mounting 27

J

JumpStart
  installing 46
Jumpstart
  Generating the finish scripts 46
  overview 46
  Preparing installation resources 47

K

keyless licensing
  Veritas InfoScale 18

L

language packages
  disk space 25
  removal 159
licensing
  registering Veritas InfoScale product license keys 17
limitatoins
  online upgrade 69
Live Upgrade
  administering Solaris 10 boot environments 78
  completing Solaris 10 upgrade 76
  preparing 74
  reverting to primary boot environment 78
  Solaris 10 systems 73
  supported upgrade paths 72
  upgrading Solaris 10 on alternate boot disk 75
  upgrading Solaris 10 using the installer 75
  verifying Solaris 10 upgrade 78
LLT
  directives 122–123

LLT *(continued)*
interconnects 34
manual configuration 121
starting 129
LLT directives
link 122–123
link-lowpri 122–123
set-cluster 122–123
set-node 122–123

M

MAC addresses 30
media speed 34
optimizing 33

N

network switches 30
non-SCSI-3 fencing
manual configuration 145
setting up 145
non-SCSI3 fencing
setting up 113
using installer 113

O

obtaining
installer patches either automatically or manually 28
optimizing
media speed 33

P

parameters
eeprom 30
post-deployment scripts 49
post-upgrade tasks 63, 71, 77
preparing
Live Upgrade 74
private network
configuring 30
product installer
InfoScale Availability configuration overview 85

R

RAM
installation requirement 25
re-establish
SFMH
VBS 66
trust with Steward 65
WAC communication 64
release information 24
requirements
Ethernet controllers 25
fibre channel 25
hardware 25
RAM Ethernet controllers 25
SCSI host bus adapter 25
response file variables
installation 42
uninstall 162
response files
about 40
installation 41
syntax 41
uninstalling 161
upgrading 79
rsh 86
configuration 28

S

sample response file
installation 44
uninstall 163
script-based installer
online upgrade 69
SCSI host bus adapter 25
SCSI-3 persistent reservations
verifying 140
setting
environment variables 55
SMTP email notification 100
SNMP trap notification 101
solaris10 brand zones 53
ssh 86
configuration 28
starting configuration
installInfoScale Availability program 86
product installer 86
starting InfoScale Availability after manual upgrade 128
supported operating systems 26
supported upgrade paths
Live Upgrade 72
switches 30

synchronizing time settings, before installing 34

T

types.cf 126
    bundled agents 126
types.cf file 126

U

uninstall
    response file variables 162
    sample response file 163
    using the installer 157
uninstalling
    about removing Veritas InfoScale Availability 156
    language packages 159
    preparing to uninstall 156
    remote 157
    using pkg uninstall command 159
    using pkgrm command 159
    using response files 161
updating licenses
    Veritas InfoScale 20
upgrade Steward 67
upgrading
    using response files 79
upgrading online
    using installer 69

V

verifying
    product installation 54
Veritas InfoScale
    Hardware requirements 25
    keyless licensing 18
    mounting ISO image 27
    product installer 38
    registering Veritas InfoScale product license keys 17
    updating licenses 20
Veritas InfoScale Availability
    about 15
vxdisksetup command 109