Veritas InfoScale™ 7.4.2 Installation Guide - Linux

VERITAS™

Last updated: 2020-05-31

Legal Notice

Copyright © 2020 Veritas Technologies LLC. All rights reserved.

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

This product may contain third-party software for which Veritas is required to provide attribution to the third-party (“Third-Party Programs”). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:
https://www.veritas.com/about/legal/license-agreements

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

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

Technical Support

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

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

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

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

Documentation

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

Documentation feedback

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

Contents

Section 1 Planning and preparation 8

Chapter 1 Introducing Veritas InfoScale 9

About the Veritas InfoScale product suite 9
Components of the Veritas InfoScale product suite 9
About the co-existence of Veritas InfoScale products 11

Chapter 2 Licensing Veritas InfoScale 12

About Veritas InfoScale product licensing 12
About telemetry data collection in InfoScale 13
Licensing notes 15
Registering Veritas InfoScale using permanent license key file 17
Registering Veritas InfoScale using keyless license 18
Managing InfoScale licenses 20
About the `vxlicinstupgrade` utility 22
Generating license report with `vxlicrep` command 23

Chapter 3 System requirements 24

Important release information 24
Disk space requirements 25
Hardware requirements 25
SF and SFHA hardware requirements 26
SFCFS and SFCFSHA hardware requirements 26
SF Oracle RAC and SF Sybase CE hardware requirements 27
VCS hardware requirements 28
Supported operating systems and database versions 29
Number of nodes supported 29

Chapter 4 Preparing to install 30

Mounting the ISO image 30
Setting up ssh or rsh for inter-system communications 31
Obtaining installer patches 31
Disabling external network connection attempts 32

Verifying the systems before installation 32
Setting up the private network 33
Optimizing LLT media speed settings on private NICs 36
Guidelines for setting the media speed for LLT interconnects 36
Guidelines for setting the maximum transmission unit (MTU) for LLT interconnects in Flexible Storage Sharing (FSS) environments 36
Setting up shared storage 37
Setting up shared storage: SCSI 37
Setting up shared storage: Fibre Channel 38
Synchronizing time settings on cluster nodes 40
Setting the kernel.hung_task_panic tunable 40
Planning the installation setup for SF Oracle RAC and SF Sybase CE systems 40
Planning your network configuration 41
Planning the storage 44
Planning volume layout 49
Planning file system design 50
Setting the umask before installation 50
Setting the kernel.panic tunable 50
Configuring the I/O scheduler 51

Section 2 Installation of Veritas InfoScale 52

Chapter 5 Installing Veritas InfoScale using the installer 53

Installing Veritas InfoScale using the installer 53

Chapter 6 Installing Veritas InfoScale using response files 56

About response files 56
Syntax in the response file 57
Installing Veritas InfoScale using response files 57
Response file variables to install Veritas InfoScale 58
Sample response files for Veritas InfoScale installation 60

Chapter 7 Installing Veritas Infoscale using operating system-specific methods 62

Verifying Veritas InfoScale RPMs 62
About installing Veritas InfoScale using operating system-specific methods 64
Installing Veritas InfoScale using Kickstart 64
Sample Kickstart configuration file 66
Installing Veritas InfoScale using yum 68
Installing Veritas InfoScale using the Red Hat Satellite server 71
Using Red Hat Satellite server to install Veritas InfoScale products 72

Chapter 8 Completing the post installation tasks 74

Verifying product installation 74
Setting environment variables 75
Managing the Veritas telemetry collector on your server 76
Next steps after installation 78

Section 3 Uninstallation of Veritas InfoScale 80

Chapter 9 Uninstalling Veritas InfoScale using the installer 81

Removing VxFS file systems 81
Removing rootability 82
Moving volumes to disk partitions 83
Moving volumes onto disk partitions using VxVM 83
Removing the Replicated Data Set 85
Uninstalling Veritas InfoScale RPMs using the installer 87
Removing the Storage Foundation for Databases (SFDB) repository 88

Chapter 10 Uninstalling Veritas InfoScale using response files 90

Uninstalling Veritas InfoScale using response files 90
Response file variables to uninstall Veritas InfoScale 91
Sample response file for Veritas InfoScale uninstallation 92

Section 4 Installation reference 93

Appendix A Installation scripts 94

Installation script options 94

Appendix B Tunable files for installation 100

About setting tunable parameters using the installer or a response file 100
Setting tunables for an installation, configuration, or upgrade 101
Setting tunables with no other installer-related operations 102
Setting tunables with an un-integrated response file 103
Preparing the tunables file 104
Setting parameters for the tunables file 104
Tunables value parameter definitions 105

Appendix C Troubleshooting installation issues 113

Restarting the installer after a failed network connection 113
About the VRTSspt RPM troubleshooting tools 113
Incorrect permissions for root on remote system 114
Inaccessible system 115

Section 1

Planning and preparation

- Chapter 1. Introducing Veritas InfoScale
- Chapter 2. Licensing Veritas InfoScale
- Chapter 3. System requirements
- Chapter 4. Preparing to install

Chapter 1

Introducing Veritas InfoScale

This chapter includes the following topics:

- About the Veritas InfoScale product suite
- Components of the Veritas InfoScale product suite
- About the co-existence of Veritas InfoScale products

About the Veritas InfoScale product suite

The Veritas InfoScale product suite addresses enterprise IT service continuity needs. They provide resiliency and software defined storage for critical services across a data center in physical, virtual, and cloud environments. The clustering solution provides high availability and disaster recovery for applications across geographies.

The Veritas InfoScale product suite offers the following products:

- Veritas InfoScale Foundation
- Veritas InfoScale Storage
- Veritas InfoScale Availability
- Veritas InfoScale Enterprise

Components of the Veritas InfoScale product suite

Each new InfoScale product consists of one or more components. Each component within a product offers a unique capability that you can configure for use in your environment.

Table 1-1 lists the components of each Veritas InfoScale product.

Table 1-1 Veritas InfoScale product suite

Product	Description	Components
Veritas InfoScale™ Foundation	Veritas InfoScale™ Foundation delivers a comprehensive solution for heterogeneous online storage management while increasing storage utilization and enhancing storage I/O path availability.	Storage Foundation (SF) Standard (entry-level features)
Veritas InfoScale™ Storage	Veritas InfoScale™ Storage enables organizations to provision and manage storage independently of hardware types or locations while delivering predictable Quality-of-Service, higher performance, and better Return-on-Investment.	Storage Foundation (SF) Enterprise including Replication Storage Foundation Cluster File System (SFCFS)
Veritas InfoScale™ Availability	Veritas InfoScale™ Availability helps keep an organization's information and critical business services up and running on premise and across globally dispersed data centers.	Cluster Server (VCS) including HA/DR
Veritas InfoScale™ Enterprise	Veritas InfoScale™ Enterprise addresses enterprise IT service continuity needs. It provides resiliency and software defined storage for critical services across your datacenter infrastructure.	Cluster Server (VCS) including HA/DR Storage Foundation (SF) Enterprise including Replication Storage Foundation and High Availability (SFHA) Storage Foundation Cluster File System High Availability (SFCFSHA) Storage Foundation for Oracle RAC (SF Oracle RAC) Storage Foundation for Sybase ASE CE (SFSYBASECE)

About the co-existence of Veritas InfoScale products

You cannot install an InfoScale product on a system where another InfoScale product is already installed.

Chapter 2

Licensing Veritas InfoScale

This chapter includes the following topics:

- About Veritas InfoScale product licensing
- About telemetry data collection in InfoScale
- Licensing notes
- Registering Veritas InfoScale using permanent license key file
- Registering Veritas InfoScale using keyless license
- Managing InfoScale licenses
- Generating license report with vxlicrep command

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 product with a permanent license
  When you purchase a Veritas InfoScale product, you receive a License Key certificate. The certificate specifies the products and the number of product licenses purchased.
  See "Registering Veritas InfoScale using permanent license key file" on page 17.
- Install product without a permanent 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 using keyless license" on page 18.

- Veritas collects licensing and platform related information from InfoScale products as part of the Veritas Product Improvement Program. The information collected helps identify how customers deploy and use the product, and enables Veritas to manage customer licenses more efficiently. See "About telemetry data collection in InfoScale" on page 13.

Visit the Veritas licensing Support website, for more information about the licensing process.

www.veritas.com/licensing/process

About telemetry data collection in InfoScale

The Veritas Telemetry Collector is used to collect licensing and platform related information from InfoScale products as part of the Veritas Product Improvement Program. The Veritas Telemetry Collector sends this information to an edge server.

The information collected helps identify how customers deploy and use the product, and enables Veritas to manage customer licenses more efficiently. The edge server does not collect any private information and only uses information specific to product, licensing, and platform (which includes operating system and server hardware).

Table 2-1 Information sent by the collector

Category	Information attributes
Product	■ Telemetry data version ■ Cluster ID ■ Product version ■ Time stamp

Table 2-1 Information sent by the collector *(continued)*

Category	Information attributes
Licensing	■ Product ID ■ Serial number ■ Serial ID ■ License meter ■ Fulfillment ID ■ Platform ■ Version ■ SKU type ■ VXKEYLESS ■ License type ■ SKU
Operating system	■ Platform name ■ Version ■ TL number ■ Kernel/SRU
Server hardware	■ Architecture ■ CPU op-mode(s) ■ CPU(s) ■ Core(s) per socket ■ Thread(s) per core ■ Socket(s) ■ Vendor ID ■ CPU model name ■ CPU frequency ■ Hypervisor vendor ■ Memory

By default, the Veritas Telemetry Collector will collect telemetry data every Tuesday at 1:00 A.M. as per the local system time. The time and interval of data collection can be customized by the user if required.

You can configure the Veritas Telemetry Collector while installing or upgrading the product, See “Installing Veritas InfoScale using the installer” on page 53.. You can also manage the Veritas Telemetry Collector on each of your servers by using the `/opt/VRTSvlic/tele/bin/TelemetryCollector` command. For more information, See “Managing the Veritas telemetry collector on your server” on page 76.

Configure the firewall policy such that the ports required for telemetry data collection are not blocked. Refer to your respective firewall or OS vendor documents for the required configuration.

Note: Ensure that you reboot the server after uninstalling the product to ensure that all services related to the Veritas Telemetry Collector are stopped successfully.

Licensing notes

Review the following licensing notes before you install or upgrade the product.

- If you use a keyless license option, you must configure Veritas InfoScale Operations Manager within two months of product installation and add the node as a managed host to the Veritas InfoScale Operations Manager Management Server. Failing this, a warning message for non-compliance is displayed periodically.
  For more details, refer to Veritas InfoScale Operations Manager product documentation.
- Note the following limitation in case of InfoScale Availability and InfoScale Storage co-existence:
  If Keyless licensing type is selected during the product installation, checks performed to monitor the number of days of product installation are based on the InfoScale Storage component. As a result, if you do not enter a valid license key file or do not add the host as a managed host within 60 days of InfoScale Storage installation, a non-compliance error is logged every 4 hrs in the Event Viewer.
- The text-based license keys that are used in 7.3.1 and earlier versions are not supported when upgrading to later versions. If your current product is installed using a permanent license key and you do not have a permanent license key file for the newer InfoScale version, you can temporarily upgrade using the keyless licensing. Then you must procure a permanent license key file from the Veritas license certificate and portal within 60 days, and upgrade using the permanent license key file to continue using the product.
- The license key file must be present on the same node where you are trying to install the product.

  Note: The license key file must not be saved in the root directory (`/`) or the default license directory on the local host (`/etc/vx/licesnes/lic`). You can save the license key file inside any other directory on the local host.

- You can manage the license keys using the `vxlicinstupgrade` utility.
  See "Managing InfoScale licenses" on page 20.
- Before upgrading the product, review the licensing details and back up the older license key. If the upgrade fails for some reason, you can temporarily revert to the older product using the older license key to avoid any application downtime.
- You can use the license assigned for higher Stock Keeping Units (SKU) to install the lower SKUs.
  For example, if you have procured a license that is assigned for InfoScale Enterprise, you can use the license for installing any of the following products:
  - InfoScale Foundation
  - InfoScale Storage
  - InfoScale Availability

  The following table provides details about the license SKUs and the corresponding products that can be installed:

License SKU procured	**Products that can be installed**			
	InfoScale Foundation	InfoScale Storage	InfoScale Availability	InfoScale Enterprise
InfoScale Foundation	✓	X	X	X
InfoScale Storage	✓	✓	X	X
InfoScale Availability	X	X	✓	X
InfoScale Enterprise	✓	✓	✓	✓

Note: At any given point in time you can install only one product.

Registering Veritas InfoScale using permanent license key file

Slf license key files are required while registering Veritas InfoScale using a permanent license key file. Ensure that the license key file is downloaded on the local host, where you want to install or upgrade the product.

Note: The license key file must not be saved in the root directory (`/`) or the default license directory on the local host (`/etc/vx/licesnes/lic`). You can save the license key file inside any other directory on the local host.

You can register your permanent license key file in the following ways:

Using the `installer`

You can register your InfoScale product using a permanent license key file during the installation process.

- Run the following command:

  ```
  ./installer
  ```

- During the installation, the following interactive message appears:

  ```
  1) Enter a valid license key(key file path needed)
  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 the license key.
- Then provide the absolute path of the .slf license key file saved on the current node.
  Example:
  `/downloads/InfoScale_keys/XYZ.slf`

Alternatively, you can register your InfoScale product using the installer menu.

- Run the following command:

  ```
  ./installer
  ```

- Select the **L) License a Product** option in the installer menu.
- Then proceed to provide the licensing details as prompted.

To install InfoScale using the installer:

See "Installing Veritas InfoScale using the installer" on page 53.

Manual

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

```
# cd /opt/VRTS/bin
# ./vxlicinstupgrade -k <key file path>
```

or

```
# ./vxlicinst -k <key file path>
```

then,

```
# vxdctl license init
```

Note: It is recommended to use the `vxlicinstupgrade` utility to manage licenses. The `vxlicinst` utility is expected to be deprecated in near future.

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

```
# cd /opt/VRTS/bin
# ./vxlicinstupgrade -k <key file path>
```

For more information:

See "Managing InfoScale licenses" 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 using keyless license

You can enable keyless licensing for your product in the following ways:

Using the `installer`

You can enable keyless licensing for InfoScale during the installation process.

- Run the following command:

  ```
  ./installer
  ```

- During the installation, the following interactive message appears:

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

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

- Enter **2** to enable keyless licensing.

Alternatively, you can enable keyless licensing for your InfoScale product using the installer menu.

- Run the following command:

  ```
  ./installer
  ```

- Select the **L) License a Product** option in the installer menu.
- Then proceed to enable keyless licensing as prompted.

To install InfoScale using the installer:

See "Installing Veritas InfoScale using the installer" on page 53.

Manual

If you are performing a fresh installation or upgrade, perform the following steps:

1. Change your current working directory:

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

2. View the keyless product code for the product you want to install:

   ```
   # vxkeyless displayall
   ```

3. Enter the product code in the exact format as displayed in the previous step:

   ```
   # vxkeyless set <product code>
   ```

   Example:

   ```
   # vxkeyless set ENTERPRISE
   ```

For more information:

See "Managing InfoScale licenses" on page 20.

Warning: Within 60 days of choosing this option, you must install a valid license key file 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:

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

Managing InfoScale licenses

After you have installed a Veritas InfoScale product, you may need to manage the product license, for example, to switch from a keyless to a permanent license type.

You can manage your licenses by using the `vxlicinstupgrade` or `vxkeyless` utilities which are located in the product installation directory.

Using the `vxlicinstupgrade`

To add or update a permanent license, run the following commands:

```
# cd /opt/VRTS/bin
```

```
# ./vxlicinstupgrade -k <key file path>
```

Where, the *<key file path>* is the absolute path of the .slf license key file saved on the current node.

Example:

```
/downloads/InfoScale_keys/XYZ.slf
```

For more information on `vxlicinstupgrade` utility:

See "About the `vxlicinstupgrade` utility" on page 22.

For more information on permanent licensing:

See "Registering Veritas InfoScale using permanent license key file" on page 17.

Using the `vxkeyless`

To add or update a keyless license, perform the following steps:

1. Change your current working directory:

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

2. View the keyless product code for the product you want to install:

   ```
   # vxkeyless displayall
   ```

3. Enter the product code in the exact format as displayed in the previous step:

   ```
   # vxkeyless set <keyless license text-string>
   ```

   Example:

   ```
   # vxkeyless set ENTERPRISE
   ```

For more information on keyless licensing:

See "Registering Veritas InfoScale using keyless license" on page 18.

About the `vxlicinstupgrade` utility

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

- Upgrade to another Veritas InfoScale license
- Update a keyless 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 is keyless or permanent and if the user is trying to install the keyless or permanent license of the same product.
  Example: If the 7.4.2 Foundation Keyless license key is already installed on a system and the user tries to install another 7.4.2 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 file is a permanent license of the same product
  Example: If the 7.4.2 Foundation Keyless license key is already installed on a system and the user tries to install 7.4.2 Foundation permanent license key file, then the `vxlicinstupgrade` utility installs the new license at `/etc/vx/licenses/lic` and the 7.4.2 Foundation Keyless key is deleted.
- The `vxlicinstupgrade` utility in Veritas InfoScale does not support managing the text-based license keys used in versions before 7.4.
- If the current key is of a lower version and the user tries to install a higher version license key.
  Example: If 7.0 Storage license key is already installed on a system and the user tries to install 7.4.2 Storage license key file, then the `vxlicinstupgrade` utility installs the new license at `/etc/vx/licenses/lic` and the 7.0 Storage key is deleted.

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

Generating license report with `vxlicrep` command

The `vxlicrep` command generates a report of the product licenses in use on your system.

To display a license report:

- Enter the `# vxlicrep` command without any options to display the report of all the product licenses on your system, or
- Enter the `# vxlicrep` command with any of the following options to display the type of report required:

-g	default report
-k <key>	print report for input key
-v	print version
-h	display this help

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:
  https://www.veritas.com/support/en_US/article.000126340
- 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.000126344
- 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.000126342

Disk space requirements

Table 3-1 lists the minimum disk space requirements for RHEL and supported RHEL-compatible distributions for each product when the /opt, /root, /var, and /bin directories are created on the same disk..

Table 3-1 Disk space requirements for RHEL and supported RHEL-compatible distributions

Product name	RHEL 7 (MB)	RHEL 8 (MB)
Veritas InfoScale Foundation	980	1058
Veritas InfoScale Availability	756	868
Veritas InfoScale Storage	1408	1483
Veritas InfoScale Enterprise	1460	1533

Table 3-2 lists the minimum disk space requirements for SLES each product when the /opt, /root, /var, and /bin directories are created on the same disk..

Table 3-2 Disk space requirements

Product name	SLES 12 (MB)	SLES 15 (MB)
Veritas InfoScale Foundation	1192	897
Veritas InfoScale Availability	831	745
Veritas InfoScale Storage	1766	1386
Veritas InfoScale Enterprise	1823	1432

Hardware requirements

This section lists the hardware requirements for Veritas InfoScale.

Table 3-3 lists the hardware requirements for each component in Veritas InfoScale.

Table 3-3 Hardware requirements for components in Veritas InfoScale

Component	Requirement
Storage Foundation (SF) Storage Foundation for High Availability (SFHA)	See "SF and SFHA hardware requirements" on page 26.
Storage Foundation Cluster File System (SFCFS) and Storage Foundation Cluster File System for High Availability (SFCFSHA)	See "SFCFS and SFCFSHA hardware requirements" on page 26.
Storage Foundation for Oracle RAC (SF Oracle RAC) Storage Foundation for Sybase CE (SF Sybase CE)	See "SF Oracle RAC and SF Sybase CE hardware requirements" on page 27.
Cluster Server (VCS)	See "VCS hardware requirements" on page 28.

For additional information, see the hardware compatibility list (HCL) at:

https://www.veritas.com/support/en_US/article.000126344

SF and SFHA hardware requirements

Table 3-4 lists the hardware requirements for SF and SFHA.

Table 3-4 SF and SFHA hardware requirements

Item	Requirement
Memory	Each system requires at least 1 GB.

SFCFS and SFCFSHA hardware requirements

Table 3-5 lists the hardware requirements for SFCFSHA.

Table 3-5 Hardware requirements for SFCFSHA

Requirement	Description
Memory (Operating System)	2 GB of memory.

Table 3-5 Hardware requirements for SFCFSHA *(continued)*

Requirement	Description
CPU	A minimum of 2 CPUs.
Node	All nodes in a Cluster File System must have the same operating system version.
Shared storage	Shared storage can be one or more shared disks or a disk array connected either directly to the nodes of the cluster or through a Fibre Channel Switch. Nodes can also have non-shared or local devices on a local I/O channel. It is advisable to have `/`, `/usr`, `/var` and other system partitions on local devices. In a Flexible Storage Sharing (FSS) environment, shared storage may not be required.
Fibre Channel or iSCSI storage	Each node in the cluster must have a Fibre Channel I/O channel or iSCSI storage to access shared storage devices. The primary component of the Fibre Channel fabric is the Fibre Channel switch.
Cluster platforms	There are several hardware platforms that can function as nodes in a Veritas InfoScale cluster. See the *Veritas InfoScale 7.4.2 Release Notes*. For a cluster to work correctly, all nodes must have the same time. If you are not running the Network Time Protocol (NTP) daemon, make sure the time on all the systems comprising your cluster is synchronized.
SAS or FCoE	Each node in the cluster must have an SAS or FCoE I/O channel to access shared storage devices. The primary components of the SAS or Fibre Channel over Ethernet (FCoE) fabric are the switches and HBAs.

SF Oracle RAC and SF Sybase CE hardware requirements

Table 3-6 lists the hardware requirements for basic clusters.

Table 3-6 Hardware requirements for basic clusters

Item	Description
DVD drive	A DVD drive on one of the nodes in the cluster.

Table 3-6 Hardware requirements for basic clusters *(continued)*

Item	Description
Disks	All shared storage disks support SCSI-3 Persistent Reservations (PR). **Note:** The coordinator disk does not store data, so configure the disk as the smallest possible LUN on a disk array to avoid wasting space. The minimum size required for a coordinator disk is 128 MB.
RAM	Each system requires at least 2 GB.
Swap space	For SF Oracle RAC: See the Oracle Metalink document: 169706.1
Network	Two or more private links and one public link. Links must be 100BaseT or gigabit Ethernet directly linking each node to the other node to form a private network that handles direct inter-system communication. These links must be of the same type; you cannot mix 100BaseT and gigabit. Veritas recommends gigabit Ethernet using enterprise-class switches for the private links. Oracle RAC requires that all nodes use the IP addresses from the same subnet.
Fiber Channel or SCSI host bus adapters	At least one additional SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.

VCS hardware requirements

Table 3-7 lists the hardware requirements for a VCS cluster.

Table 3-7 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.

Table 3-7 Hardware requirements for a VCS cluster *(continued)*

Item	Description
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). **Note:** SFHA also supports non-SCSI3 server-based fencing configuration in virtual environments that do not support SCSI-3 PR-compliant storage.
Network Interface Cards (NICs)	In addition to the built-in public NIC, VCS requires at least one more NIC per system. Veritas recommends two additional NICs. You can also configure aggregated interfaces. Veritas recommends that you turn off the spanning tree on the LLT switches, and set port-fast on.
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 256 megabytes.

Supported operating systems and database versions

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

Number of nodes supported

Veritas InfoScale supports cluster configurations up to 128 nodes.

SFHA, SFCFSHA, SF Oracle RAC: Flexible Storage Sharing (FSS) only supports cluster configurations with up to 64 nodes.

SFHA, SFCFSHA: SmartIO writeback caching only supports cluster configurations with up to 2 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
- Setting up shared storage
- Synchronizing time settings on cluster nodes
- Setting the kernel.hung_task_panic tunable
- Planning the installation setup for SF Oracle RAC and SF Sybase CE systems

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. Mount the image:

```
# mount -o loop <ISO_image_path> /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 Veritas InfoScale version 7.0 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.

See “Disabling external network connection attempts” on page 32.

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 1.

3 Unzip the patch tar file. For example, run the following command:

```
# gunzip cpi-7.4.2P2-patches.tar.gz
```

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

```
# tar -xvf cpi-7.4.2P2-patches.tar
patches/
patches/CPI7.4.2P2.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/CPI7.4.2P2.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, RPMs, 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

This topic applies to VCS, SFHA, SFCFS, SFCFSHA, SF Oracle RAC, and SF Sybase CE.

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.

Refer to the *Cluster Server 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

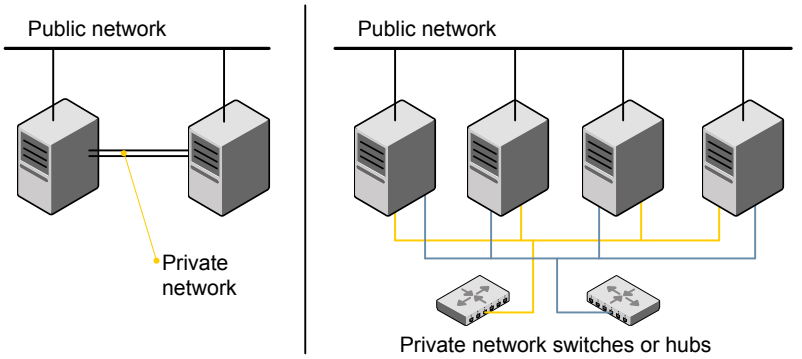

You need to configure at least two independent networks between the cluster nodes with a network switch for each network. You can also interconnect multiple layer 2 switches for advanced failure protection. Such connections for LLT are called cross-links.

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

Figure 4-2 Private network setup with crossed links

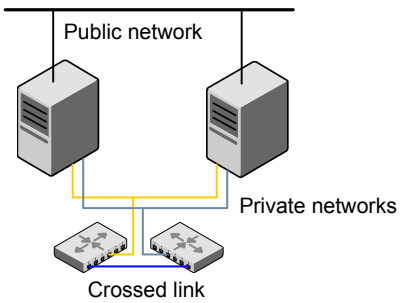

Veritas recommends one of the following two configurations:

- 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 Veritas InfoScale private NICs on each system.

3 Use crossover Ethernet cables, switches, or independent hubs for each Veritas InfoScale 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 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.

Guidelines for setting the maximum transmission unit (MTU) for LLT interconnects in Flexible Storage Sharing (FSS) environments

Review the following guidelines for setting the MTU for LLT interconnects in FSS environments:

- Set the maximum transmission unit (MTU) to the highest value (typically 9000) supported by the NICs when LLT (both high priority and low priority links) is configured over Ethernet or UDP. Ensure that the switch is also set to 9000 MTU.

  Note: MTU setting is not required for LLT over RDMA configurations.

- For virtual NICs, all the components—the virtual NIC, the corresponding physical NIC, and the virtual switch—must be set to 9000 MTU.

- If a higher MTU cannot be configured on the public link (because of restrictions on other components such as a public switch), do not configure the public link in LLT. LLT uses the lowest of the MTU that is configured among all high priority and low priority links.

Setting up shared storage

This topic applies to VCS, SFHA, SFCFSHA, SF Oracle RAC, and SF Sybase CE.

The sections describe how to set up the SCSI and the Fibre Channel devices that the cluster systems share.

Setting up shared storage: SCSI

Perform the following steps to set up shared storage.

To set up shared storage

1. Connect the disk to the first cluster system.
2. Power on the disk.
3. Connect a terminator to the other port of the disk.
4. Boot the system. The disk is detected while the system boots.
5. Press CTRL+A to bring up the SCSI BIOS settings for that disk.

   Set the following:

   - Set Host adapter SCSI ID = 7, or to an appropriate value for your configuration.
   - Set Host Adapter BIOS in Advanced Configuration Options to Disabled.
6. Format the shared disk and create required partitions on it.

   Perform the following:

   - Identify your shared disk name. If you have two internal SCSI hard disks, your shared disk is /dev/sdc.
     Identify whether the shared disk is sdc, sdb, and so on.
   - Type the following command:

     ```
     # fdisk /dev/shareddiskname
     ```

     For example, if your shared disk is sdc, type:

     ```
     # fdisk /dev/sdc
     ```

- Create disk groups and volumes using Volume Manager utilities.
- To apply a file system on the volumes, type:

  ```
  # mkfs -t fs-type /dev/vx/dsk/disk-group/volume
  ```

  For example, enter the following command:

  ```
  # mkfs -t vxfs /dev/vx/dsk/dg/vol01
  ```

  Where the name of the disk group is dg, the name of the volume is vol01, and the file system type is vxfs.

7. Power off the disk.
8. Remove the terminator from the disk and connect the disk to the other cluster system.
9. Power on the disk.
10. Boot the second system. The system can now detect the disk.
11. Press Ctrl+A to bring up the SCSI BIOS settings for the disk.

    Set the following:

    - Set Host adapter SCSI ID = 6, or to an appropriate value for your configuration. Note that the SCSI ID should be different from the one configured on the first cluster system.
    - Set Host Adapter BIOS in Advanced Configuration Options to Disabled.
12. Verify that you can view the shared disk using the `fdisk` command.

Setting up shared storage: Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up shared storage for Fibre Channel

1. Connect the Fibre Channel disk to a cluster system.
2. Boot the system and change the settings of the Fibre Channel. Perform the following tasks for all QLogic adapters in the system:
   - Press Alt+Q to bring up the QLogic adapter settings menu.
   - Choose **Configuration Settings**.
   - Click Enter.
   - Choose **Advanced Adapter Settings**.
   - Click Enter.

- Set the Enable Target Reset option to **Yes** (the default value).
- Save the configuration.
- Reboot the system.

3 Verify that the system detects the Fibre Channel disks properly.

4 Create volumes. Format the shared disk and create required partitions on it and perform the following:

- Identify your shared disk name. If you have two internal SCSI hard disks, your shared disk is /dev/sdc.
  Identify whether the shared disk is sdc, sdb, and so on.
- Type the following command:

  ```
  # fdisk /dev/shareddiskname
  ```

  For example, if your shared disk is sdc, type:

  ```
  # fdisk /dev/sdc
  ```

- Create disk groups and volumes using Volume Manager utilities.
- To apply a file system on the volumes, type:

  ```
  # mkfs -t fs-type /dev/vx/rdsk/disk-group/volume
  ```

  For example, enter the following command:

  ```
  # mkfs -t vxfs /dev/vx/rdsk/dg/vol01
  ```

  Where the name of the disk group is dg, the name of the volume is vol01, and the file system type is vxfs.

5 Repeat step 2 and step 3 for all nodes in the clusters that require connections with Fibre Channel.

6 Power off this cluster system.

7 Connect the same disks to the next cluster system.

8 Turn on the power for the second system.

9 Verify that the second system can see the disk names correctly—the disk names should be the same.

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.

Setting the kernel.hung_task_panic tunable

The topic applies to SFHA, SFCFSHA, and VCS.

By default, in the Linux kernel the `kernel.hung_task_panic` tunable is enabled and the `kernel.hung_task_timeout_secs` tunable is set to a default non-zero value.

To ensure that the node does not panic, the `kernel.hung_task_panic` tunable must be disabled. If `kernel.hung_task_panic` is enabled, then it causes the kernel to panic when any of the following kernel threads waits for more than the `kernel.hung_task_timeout_secs` value:

- The vxfenconfig thread in the vxfen configuration path waits for GAB to seed.
- The vxfenswap thread in the online coordinator disks replacement path waits for the snapshot of peer nodes of the new coordinator disks.

To disable the `kernel.hung_task_panic` tunable:

- Set the `kernel.hung_task_panic` tunable to zero (0) in the `/etc/sysctl.conf` file. This step ensures that the change is persistent across node restarts.
- Run the command on each node.
  `# sysctl -w kernel.hung_task_panic=0`

To verify the `kernel.hung_task_panic tunable` value, run the following command:

- `# sysctl -a | grep hung_task_panic`

Planning the installation setup for SF Oracle RAC and SF Sybase CE systems

This section provides guidelines and best practices for planning resilient, high-performant clusters. These best practices suggest optimal configurations for your core clustering infrastructure such as network and storage. Recommendations are also provided on planning for continuous data protection and disaster recovery.

Review the following planning guidelines before you install Veritas InfoScale:

- Planning your network configuration
  See "Planning your network configuration" on page 41.
- Planning the storage
  See "Planning the storage" on page 44.
- Planning volume layout
  See "Planning volume layout" on page 49.
- Planning file system design
  See "Planning file system design" on page 50.

Planning your network configuration

The following practices are recommended for a resilient network setup:

- Configure the private cluster interconnect over multiple dedicated gigabit Ethernet links. All single point of failures such as network interface cards (NIC), switches, and interconnects should be eliminated.
- The NICs used for the private cluster interconnect should have the same characteristics regarding speed, MTU, and full duplex on all nodes. Do not allow the NICs and switch ports to auto-negotiate speed.
- Configure non-routable IP addresses for private cluster interconnects.
- The default value for LLT peer inactivity timeout is 16 seconds.
  For SF Oracle RAC: The value should be set based on service availability requirements and the propagation delay between the cluster nodes in case of campus cluster setup. The LLT peer inactivity timeout value indicates the interval after which Veritas InfoScale on one node declares the other node in the cluster dead, if there is no network communication (heartbeat) from that node.
  The default value for the CSS miss-count in case of Veritas InfoScale is 600 seconds. The value of this parameter is much higher than the LLT peer inactivity timeout so that the two clusterwares, VCS and Oracle Clusterware, do not interfere with each other's decisions on which nodes should remain in the cluster in the event of network split-brain. Veritas I/O fencing is allowed to decide on the surviving nodes first, followed by Oracle Clusterware. The CSS miss-count value indicates the amount of time Oracle Clusterware waits before evicting another node from the cluster, when it fails to respond across the interconnect.
  For more information, see the Oracle Metalink document: 782148.1

Planning the public network configuration for Oracle RAC

Identify separate public virtual IP addresses for each node in the cluster. Oracle RAC requires one public virtual IP address for the Oracle RAC listener process on

each node. Public virtual IP addresses are used by client applications to connect to the Oracle RAC database and help mitigate TCP/IP timeout delays.

For SF Oracle RAC: For Oracle 11g Release 2 and later versions, additionally, you need a Single Client Access Name (SCAN) registered in Enterprise DNS that resolves to three IP addresses (recommended). Oracle Clusterware/Grid Infrastructure manages the virtual IP addresses.

Planning the private network configuration for Oracle RAC

Oracle RAC requires a minimum of one private IP address on each node for Oracle Clusterware heartbeat.

You must use UDP IPC for the database cache fusion traffic. The Oracle RAC UDP IPC protocol requires an IP address. Depending on your deployment needs, this IP address may be a dedicated IP address or one that is shared with Oracle Clusterware.

For Oracle and later versions, you must use UDP IPC for the database cache fusion traffic.

Note: The private IP addresses of all nodes that are on the same physical network must be in the same IP subnet.

The following practices provide a resilient private network setup:

- Configure Oracle Clusterware interconnects over LLT links to prevent data corruption.
  In an Veritas InfoScale cluster, the Oracle Clusterware heartbeat link MUST be configured as an LLT link. If Oracle Clusterware and LLT use different links for their communication, then the membership change between VCS and Oracle Clusterware is not coordinated correctly. For example, if only the Oracle Clusterware links are down, Oracle Clusterware kills one set of nodes after the expiry of the css-misscount interval and initiates the Oracle Clusterware and database recovery, even before CVM and CFS detect the node failures. This uncoordinated recovery may cause data corruption.
- Oracle Clusterware interconnects need to be protected against NIC failures and link failures. For Oracle RAC 11.2.0.1 versions, the PrivNIC or MultiPrivNIC agent can be used to protect against NIC failures and link failures, if multiple links are available. Even if link aggregation solutions in the form of bonded NICs are implemented, the PrivNIC or MultiPrivNIC agent can be used to provide additional protection against the failure of the aggregated link by failing over to available alternate links. These alternate links can be simple NIC interfaces or bonded NICs.

An alternative option is to configure the Oracle Clusterware interconnects over bonded NIC interfaces.
See "High availability solutions for Oracle RAC private network" on page 43.

Note: The PrivNIC and MultiPrivNIC agents are no longer supported in Oracle RAC 11.2.0.2 and later versions for managing cluster interconnects.

For 11.2.0.2 and later versions, Veritas recommends the use of alternative solutions such as bonded NIC interfaces or Oracle High Availability IP (HAIP).

- Configure Oracle Cache Fusion traffic to take place through the private network. Veritas also recommends that all UDP cache-fusion links be LLT links.
  Oracle database clients use the public network for database services. Whenever there is a node failure or network failure, the client fails over the connection, for both existing and new connections, to the surviving node in the cluster with which it is able to connect. Client failover occurs as a result of Oracle Fast Application Notification, VIP failover and client connection TCP timeout. It is strongly recommended not to send Oracle Cache Fusion traffic through the public network.
- Use NIC bonding to provide redundancy for public networks so that Oracle RAC can fail over virtual IP addresses if there is a public link failure.

High availability solutions for Oracle RAC private network

Table 4-1 lists the high availability solutions that you may adopt for your private network.

Table 4-1 High availability solutions for Oracle RAC private network

Options	Description
Using link aggregation/ NIC bonding for Oracle Clusterware	Use a native NIC bonding solution to provide redundancy, in case of NIC failures. Make sure that a link configured under a aggregated link or NIC bond is not configured as a separate LLT link. When LLT is configured over a bonded interface, do one of the following steps to prevent GAB from reporting jeopardy membership: ■ Configure an additional NIC under LLT in addition to the bonded NIC. ■ Add the following line in the `/etc/llttab` file: `set-dbg-minlinks 2`

Table 4-1 High availability solutions for Oracle RAC private network *(continued)*

Options	Description
Using HAIP	Starting with Oracle RAC 11.2.0.2, Oracle introduced the High Availability IP (HAIP) feature for supporting IP address failover. The purpose of HAIP is to perform load balancing across all active interconnect interfaces and fail over existing non-responsive interfaces to available interfaces. HAIP has the ability to activate a maximum of four private interconnect connections. These private network adapters can be configured during the installation of Oracle Grid Infrastructure or after the installation using the `oifcfg` utility.

Planning the public network configuration for Oracle RAC

Public interconnects are used by the clients to connect to Oracle RAC database. The public networks must be physically separated from the private networks.

See Oracle RAC documentation for more information on recommendations for public network configurations.

Planning the private network configuration for Oracle RAC

Private interconnect is an essential component of a shared disk cluster installation. It is a physical connection that allows inter-node communication. Veritas recommends that these interconnects and LLT links must be the same. You must have the IP addresses configured on these interconnects, persistent after reboot. You must use solutions specific to the operating System.

See Oracle RAC documentation for more information on recommendations for private network configurations.

Planning the storage

Veritas InfoScale provides the following options for shared storage:

- CVM
  CVM provides native naming (OSN) as well as enclosure-based naming (EBN). Use enclosure-based naming for easy administration of storage. Enclosure-based naming guarantees that the same name is given to a shared LUN on all the nodes, irrespective of the operating system name for the LUN.
- CFS
- **For SF Oracle RAC:** Local storage

With FSS, local storage can be used as shared storage. The local storage can be in the form of Direct Attached Storage (DAS) or internal disk drives.

- **For SF Oracle RAC:**Oracle ASM over CVM

The following recommendations ensure better performance and availability of storage.

- Use multiple storage arrays, if possible, to ensure protection against array failures. The minimum recommended configuration is to have two HBAs for each host and two switches.
- Design the storage layout keeping in mind performance and high availability requirements. Use technologies such as striping and mirroring.
- Use appropriate stripe width and depth to optimize I/O performance.
- Use SCSI-3 persistent reservations (PR) compliant storage.
- Provide multiple access paths to disks with HBA/switch combinations to allow DMP to provide high availability against storage link failures and to provide load balancing.

Planning the storage

Table 4-2 lists the type of storage required for SF Oracle RAC and SF Sybase CE.

Table 4-2 Type of storage required for SF Oracle RAC and SF Sybase CE

Files	Type of storage
SF Oracle RAC and SF Sybase CE binaries	Local
SF Oracle RAC and SF Sybase CE database storage management repository	Shared

Planning the storage for Oracle RAC

Review the storage options and guidelines for Oracle RAC:

- Storage options for OCR and voting disk
  See "Planning the storage for OCR and voting disk" on page 46.
- Storage options for the Oracle RAC installation directories (ORACLE_BASE, CRS_HOME or GRID_HOME (depending on Oracle RAC version), and ORACLE_HOME)
  See "Planning the storage for Oracle RAC binaries and data files" on page 48.

Planning the storage for OCR and voting disk

Review the following notes before you proceed:

- Set the disk detach policy setting to (`local`) with ioship off for OCR and voting disk.
- Configure OCR and voting disk on non-replicated shared storage when you configure global clusters.
- If you plan to use FSS, configure OCR and voting disk on SAN storage.

OCR and voting disk storage configuration for external redundancy

Figure 4-3 illustrates the OCR and voting disk storage options for external redundancy.

Figure 4-3 OCR and voting disk storage configuration for external redundancy

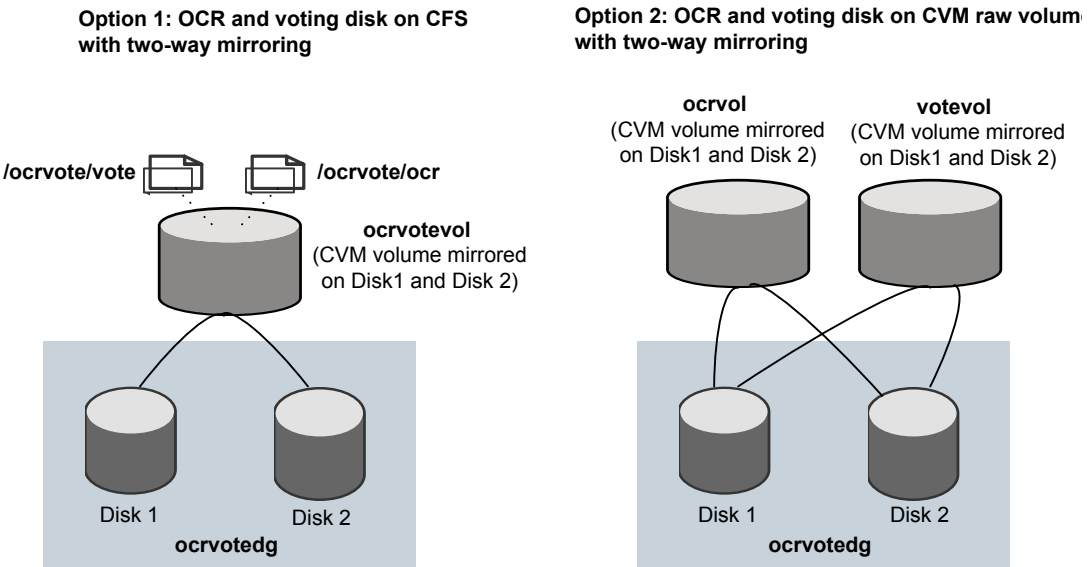

- If you want to place OCR and voting disk on a clustered file system (option 1), you need to have two separate files for OCR and voting information respectively on CFS mounted on a CVM mirrored volume.
- If you want to place OCR and voting disk on ASM disk groups that use CVM raw volumes (option 2), you need to use two CVM mirrored volumes for configuring OCR and voting disk on these volumes.

For both option 1 and option 2:

- The option **External Redundancy** must be selected at the time of installing Oracle Clusterware/Grid Infrastructure.
- The installer needs at least two LUNs for creating the OCR and voting disk storage.
  See the Oracle RAC documentation for Oracle RAC's recommendation on the required disk space for OCR and voting disk.

OCR and voting disk storage configuration for normal redundancy

Figure 4-4 illustrates the OCR and voting disk storage options for normal redundancy.

Figure 4-4 OCR and voting disk storage configuration for normal redundancy

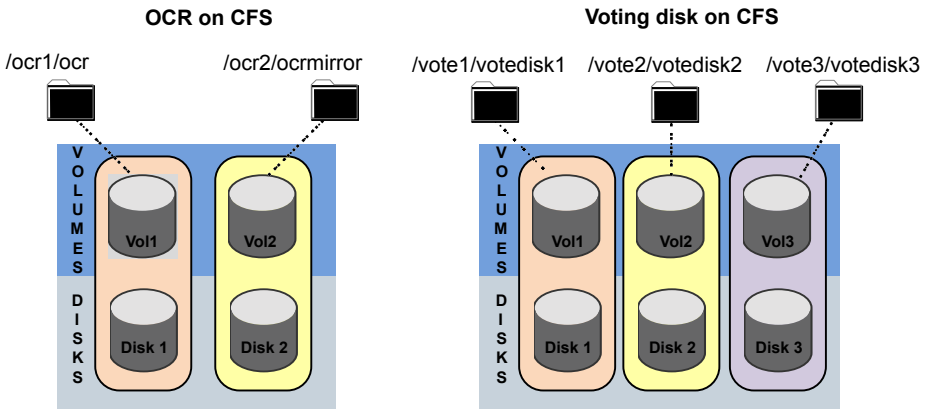

The OCR and voting disk files exist on separate cluster file systems.

Configure the storage as follows:

- Create separate filesystems for OCR and OCR mirror.
- Create separate filesystems for a minimum of 3 voting disks for redundancy.
- The option **Normal Redundancy** must be selected at the time of installing Oracle Clusterware/Grid Infrastructure.

Note: It is recommended that you configure atleast resource dependency for high availability of the OCR and voting disk resources.

Planning the storage for Oracle RAC binaries and data files

The Oracle RAC binaries can be stored on local storage or on shared storage, based on your high availability requirements.

Note: Veritas recommends that you install the Oracle Clusterware and Oracle RAC database binaries local to each node in the cluster.

Consider the following points while planning the installation:

- Local installations provide improved protection against a single point of failure and also allows for applying Oracle RAC patches in a rolling fashion.
- CFS installations provide a single Oracle installation to manage, regardless of the number of nodes. This scenario offers a reduction in storage requirements and easy addition of nodes.

Table 4-3 lists the type of storage for Oracle RAC binaries and data files.

Table 4-3 Type of storage for Oracle RAC binaries and data files

Oracle RAC files	Type of storage
Oracle base	Local
Oracle Clusterware/Grid Infrastructure binaries	Local Placing the Oracle Grid Infrastructure binaries on local disks enables rolling upgrade of the cluster.
Oracle RAC database binaries	Local Placing the Oracle RAC database binaries on local disks enables rolling upgrade of the cluster.
Database datafiles	Shared Store the Oracle RAC database files on CFS rather than on raw device or CVM raw device for easier management. Create separate clustered file systems for each Oracle RAC database. Keeping the Oracle RAC database datafiles on separate mount points enables you to unmount the database for maintenance purposes without affecting other databases. If you plan to store the Oracle RAC database on ASM, configure the ASM disk groups over CVM volumes to take advantage of dynamic multi-pathing.
Database recovery data (archive, flash recovery)	Shared Place archived logs on CFS rather than on local file systems.

Planning for Oracle RAC ASM over CVM

Review the following information on storage support provided by Oracle RAC ASM:

Supported by ASM	ASM provides storage for data files, control files, Oracle Cluster Registry devices (OCR), voting disk, online redo logs and archive log files, and backup files.
Not supported by ASM	ASM does not support Oracle binaries, trace files, alert logs, export files, tar files, core files, and application binaries.

The following practices offer high availability and better performance:

- Use CVM mirrored volumes with dynamic multi-pathing for creating ASM disk groups. Select external redundancy while creating ASM disk groups.
- The CVM raw volumes used for ASM must be used exclusively for ASM. Do not use these volumes for any other purpose, such as creation of file systems. Creating file systems on CVM raw volumes used with ASM may cause data corruption.
- Do not link the Veritas ODM library when databases are created on ASM. ODM is a disk management interface for data files that reside on the Veritas File System.
- Use a minimum of two Oracle RAC ASM disk groups. Store the data files, one set of redo logs, and one set of control files on one disk group. Store the Flash Recovery Area, archive logs, and a second set of redo logs and control files on the second disk group.
  For more information, see Oracle RAC's ASM best practices document.
- Do not configure DMP meta nodes as ASM disks for creating ASM disk groups. Access to DMP meta nodes must be configured to take place through CVM.
- Do not combine DMP with other multi-pathing software in the cluster.
- Do not use coordinator disks, which are configured for I/O fencing, as ASM disks. I/O fencing disks should not be imported or used for data.
- Volumes presented to a particular ASM disk group should be of the same speed and type.

Planning volume layout

The following recommendations ensure optimal layout of VxVM/CVM volumes:

- Mirror the volumes across two or more storage arrays, if using VxVM mirrors. Keep the Fast Mirror Resync regionsize equal to the database block size to reduce the copy-on-write (COW) overhead. Reducing the regionsize increases the amount of Cache Object allocations leading to performance overheads.

- Distribute the I/O load uniformly on all Cache Objects when you create multiple Cache Objects.
- Implement zoning on SAN switch to control access to shared storage. Be aware that physical disks may be shared by multiple servers or applications and must therefore be protected from accidental access.
- Choose DMP I/O policy based on the storage network topology and the application I/O pattern.
- Exploit thin provisioning for better return on investment.
- **For SF Oracle RAC:**
  Separate the Oracle recovery structures from the database files to ensure high availability when you design placement policies.
  Separate redo logs and place them on the fastest storage (for example, RAID 1+ 0) for better performance.
  Use "third-mirror break-off" snapshots for cloning the Oracle log volumes. Do not create Oracle log volumes on a Space-Optimized (SO) snapshot.
  Create as many Cache Objects (CO) as possible when you use Space-Optimized (SO) snapshots for Oracle data volumes.

Planning file system design

The following recommendations ensure an optimal file system design for databases:

- Create separate file systems for Oracle RAC binaries, data, redo logs, and archive logs. This ensures that recovery data is available if you encounter problems with database data files storage.
- Always place archived logs on CFS file systems rather then local file systems.
- **For SF Oracle RAC:** If using VxVM mirroring, use ODM with CFS for better performance. ODM with SmartSync enables faster recovery of mirrored volumes using Oracle resilvering.

Setting the umask before installation

The topic applies to SF Oracle RAC.

Set the umask to provide appropriate permissions for Veritas InfoScale binaries and files. This setting is valid only for the duration of the current session.

```
# umask 0022
```

Setting the kernel.panic tunable

The topic applies to SF Oracle RAC and SF Sybase CE.

By default, the kernel.panic tunable is set to zero. Therefore the kernel does not restart automatically if a node panics. To ensure that the node restarts automatically after it panics, this tunable must be set to a non-zero value.

To set the kernel.panic tunable

1. Set the kernel.panic tunable to a desired value in the /etc/sysctl.conf file.

   For example, kernel.panic = 10, will assign a value 10 seconds to the kernel.panic tunable. This step makes the change persistent across restarts.

2. Run the command:

   ```
   sysctl -w kernel.panic=10
   ```

   In case of a panic, the node will restart after 10 seconds.

Configuring the I/O scheduler

The topic applies to SF Oracle RAC and SF Sybase CE.

Veritas recommends using the Linux 'deadline' I/O scheduler for database workloads. Configure your system to boot with the 'elevator=deadline' argument to select the 'deadline' scheduler.

For information on configuring the 'deadline' scheduler for your Linux distribution, see the operating system documentation.

To determine whether a system uses the deadline scheduler, look for "elevator=deadline" in /proc/cmdline.

To configure a system to use the deadline scheduler

1. Include the elevator=deadline parameter in the boot arguments of the GRUB or ELILO configuration file. The location of the appropriate configuration file depends on the system's architecture and Linux distribution. For x86_64, the configuration file is /boot/grub/menu.lst

   - A setting for the elevator parameter is always included by SUSE in its ELILO and its GRUB configuration files. In this case, change the parameter from elevator=cfq to elevator=deadline.

2. Reboot the system once the appropriate file has been modified.

   See the operating system documentation for more information on I/O schedulers.

Section 2

Installation of Veritas InfoScale

- Chapter 5. Installing Veritas InfoScale using the installer
- Chapter 6. Installing Veritas InfoScale using response files
- Chapter 7. Installing Veritas Infoscale using operating system-specific methods
- Chapter 8. Completing the post installation tasks

Chapter 5

Installing Veritas InfoScale using the installer

This chapter includes the following topics:

- Installing Veritas InfoScale using the installer

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 list of available products is displayed. Select the product that you want to install on your system.

```
1)  Veritas InfoScale Foundation
2)  Veritas InfoScale Availability
3)  Veritas InfoScale Storage
4)  Veritas InfoScale Enterprise
b)  Back to previous menu
Select a product to install: [1-4,b,q]
```

6 The installer asks whether you want to configure the product.

```
Would you like to configure InfoScale Enterprise 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.

7 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.pdf file
present on media? [y,n,q,?] y
```

8 The installer performs the pre-checks. If it is a fresh system, the product is set as the user defined it. If the system already has a different product installed, the product is set as Veritas InfoScale Enterprise with a warning message after pre-check.

```
Veritas InfoScale Availability is installed. Installation of two
products is not supported, Veritas InfoScale Enterprise will be
installed to include Veritas InfoScale Storage and Veritas
InfoScale Availability on all the systems.
```

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

```
1)  Enter a valid license key(key file path needed)
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.

See "Registering Veritas InfoScale using permanent license key file" on page 17.

---

10 An edge server is used to collect telemetry data that will be used to help users monitor their InfoScale licenses more effectively. Enter the host name or IP address of the edge server that you want to use.

```
The Veritas Cloud Receiver (VCR) is a preconfigured, cloud-based
edge server deployed by Veritas. Enter telemetry.veritas.com to
use the Veritas Cloud Receiver as an edge server for your
environment.
Enter the edge server's hostname/ip: [q,?] ?
```

The Veritas Cloud Receiver (VCR) is a preconfigured, cloud-based edge server deployed by Veritas. Enter **telemetry.veritas.com** to use the Veritas Cloud Receiver as an edge server for your environment. For more information, See "About telemetry data collection in InfoScale" on page 13.

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

Chapter 6

Installing Veritas InfoScale using response files

This chapter includes the following topics:

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

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

To install Veritas InfoScale using response files

1. Make sure the systems where you want to install Veritas InfoScale 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 Veritas InfoScale.
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 Veritas InfoScale post-installation tasks.

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

Response file variables to install Veritas InfoScale

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

Table 6-1 Response file variables for installing Veritas InfoScale

Variable	Description
CFG{opt}{install}	Installs Veritas InfoScale RPMs. 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
CFG{keys}{vxkeyless} CFG{keys}{licensefile}	`CFG{keys}{vxkeyless}` gives the keyless key to be registered on the system. `CFG{keys}{licensefile}` gives the absolute file path to the permanent license key to be registered on the system. List 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

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

Variable	Description
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional
CFG{opt}{tmppath}	Defines the location where a working directory is created to store temporary files and the RPMs that are needed during the install. The default location is /opt/VRTStmp. 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
$CFG{edgeserver_host}	Use this parameter to configure the edge server. Enter **telemetry.veritas.com** to use the Veritas Cloud Receiver, which is a preconfigured, cloud-based edge server deployed by Veritas. Optional or required: required **Note:** An edge server is used to collect licensing and platform related information from InfoScale products as part of the Veritas Product Improvement Program. The information collected helps identify how customers deploy and use the product, and enables Veritas to manage customer licenses more efficiently.

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

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

Sample response files for Veritas InfoScale installation

The following example shows a response file for installing Veritas InfoScale using a keyless license.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{keys}{keyless}=[ qw(ENTERPRISE) ];
$CFG{opt}{gco}=1;
$CFG{opt}{install}=1;
$CFG{prod}="ENTERPRISE742";
$CFG{systems}=[ qw(system1 system2) ];
$CFG{edgeserver_host}="telemetry.veritas.com";
$CFG{edgeserver_port}=443;

1;
```

The following example shows a response file for installing Veritas InfoScale using a permanent license.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{keys}{licensefile}=["<path_to_license_key_file>"];
$CFG{opt}{gco}=1;
```

```
$CFG{opt}{install}=1;
$CFG{prod}="ENTERPRISE742";
$CFG{systems}=[ qw(system1 system2) ];
$CFG{edgeserver_host}="telemetry.veritas.com";
$CFG{edgeserver_port}=443;
1;
```

Chapter 7

Installing Veritas Infoscale using operating system-specific methods

This chapter includes the following topics:

- Verifying Veritas InfoScale RPMs
- About installing Veritas InfoScale using operating system-specific methods
- Installing Veritas InfoScale using Kickstart
- Installing Veritas InfoScale using yum
- Installing Veritas InfoScale using the Red Hat Satellite server

Verifying Veritas InfoScale RPMs

InfoScale RPMs include digital signatures in order to verify their authenticity. If you want to install the RPMs manually, you must import keys first. To import keys, perform the following steps:

1. Import the Veritas GPG key to verify InfoScale packages:

```
# rpm --import RPM-GPG-KEY-veritas-infoscale7
```

2. Display the list of Veritas keys installed for RPM verification:

```
# rpm -q gpg-pubkey --qf '%{name}-%{version}-%{release}
 -->%{summary}\n' | grep Veritas
```

3. Display the fingerprint of the Veritas key file:

```
# gpg --quiet --with-fingerprint ./RPM-GPG-KEY-veritas-infoscale7
```

For example:

```
Key fingerprint = C031 8CAB E668 4669 63DB  C8EA 0B0B C720 A17A 604B
```

To display details about the installed Veritas key file, use the `rpm -qi` command followed by the output from the previous command:

```
# rpm -qi <gpg-pubkey-file>
```

You can also use the following command to show information for the installed Veritas key file:

```
# rpm -qi `rpm -q gpg-pubkey --qf '%{name}-%{version}-%{release}
 --> %{summary}\n' | awk '/Veritas/ { print $1 }'`
```

To check the GnuPG signature of an RPM file after importing the builder's GnuPG key, use the following command:

```
# rpm -K <rpm-file>
```

Where <*rpm-file*> is the filename of the RPM package.

If the signature of the package is verified, and it is not corrupt, the following message is displayed:

```
md5 gpg OK
```

To verify the signature for all Veritas InfoScale RPMs:

```
# for i in *.rpm; do rpm -K $i; done
VRTSamf-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSaslapm-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTScavf-7.4.2.0000-GENERIC.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTScps-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSdbac-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSdbed-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSdocker-plugin-1.4-Linux.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSfsadv-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSfssdk-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSgab-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSglm-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSgms-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSllt-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSodm-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSperl-5.30.0.0-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
```

```
VRTSpython-3.7.4.1-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSsfcpi-7.4.2.0000-GENERIC.noarch.rpm: rsa sha1 (md5) pgp md5 OK
VRTSsfmh-7.4.2.0000_Linux.rpm: rsa sha1 (md5) pgp md5 OK
VRTSspt-7.4.2.0000-RHEL7.noarch.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvbs-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvcs-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvcsag-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvcsea-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvcsnr-7.4.2.0000-GENERIC.noarch.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvcswiz-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSveki-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvlic-4.01.74.004-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvxfen-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvxfs-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
VRTSvxvm-7.4.2.0000-RHEL7.x86_64.rpm: rsa sha1 (md5) pgp md5 OK
```

About installing Veritas InfoScale using operating system-specific methods

On Linux (RHEL and supported RHEL-compatible distributions), you can install Veritas InfoScale using the following methods:

- You can install Veritas InfoScale using Kickstart.
  See “Installing Veritas InfoScale using Kickstart” on page 64.
- You can install Veritas InfoScale using yum.
  See “Installing Veritas InfoScale using yum” on page 68.
- You can install Veritas InfoScale using the Red Hat Satellite server.
  See “Installing Veritas InfoScale using the Red Hat Satellite server” on page 71.

Installing Veritas InfoScale using Kickstart

You can install Veritas InfoScale using Kickstart. Kickstart is supported for Red Hat Enterprise Linux operating system.

To install Veritas InfoScale using Kickstart

1. Create a directory for the Kickstart configuration files.

   ```
   # mkdir /kickstart_files/
   ```

2. Generate the Kickstart configuration files. The configuration files have the extension `.ks.`:

   Enter the following command:

   ```
   # ./installer -kickstart /kickstart_files/
   ```

   The system lists the files.

   The output includes the following:

   ```
   The kickstart script for ENTERPRISE is generated at
   /kickstart_files/kickstart_enterprise.ks
   ```

3. Set up an NFS exported location which the Kickstart client can access. For example, if `/nfs_mount_kickstart` is the directory which has been NFS exported, the NFS exported location may look similar to the following:

   ```
   # cat /etc/exports
   /nfs_mount_kickstart  * (rw,sync,no_root_squash)
   ```

4. Copy the `rpms` directory from the installation media to the NFS location.
5. Verify the contents of the directory.

   ```
   # ls /nfs_mount_kickstart/
   ```

6. In the Veritas InfoScale Kickstart configuration file, modify the `BUILDSRC` variable to point to the actual NFS location. The variable has the following format:

   ```
   BUILDSRC="hostname_or_ip:/nfs_mount_kickstart"
   ```

7. Append the entire modified contents of the Kickstart configuration file to the operating system `ks.cfg` file.
8. Launch the Kickstart installation for the operating system.
9. After the operating system installation is complete, check the file `/opt/VRTStmp/kickstart.log` for any errors that are related to the installation of RPMs and product installer scripts.

10 Verify that all the product RPMs have been installed. Enter the following command:

```
# rpm -qa | grep -i vrts
```

11 If you do not find any installation issues or errors, configure the product stack. Enter the following command:

```
# /opt/VRTS/install/installer -configure sys1 sys2
```

Sample Kickstart configuration file

The following is a sample RedHat Enterprise Linux 7 (RHEL7) Kickstart configuration file.

```
# The packages below are required and will be installed from
OS installation media automatically during the automated installation
of products in the DVD, if they have not been installed yet.

%packages
libudev.i686
device-mapper
device-mapper-libs
parted
libgcc.i686
compat-libstdc++-33
ed
ksh
nss-softokn-freebl.i686
glibc.i686
libstdc++.i686
audit-libs.i686
cracklib.i686
db4.i686
libselinux.i686
pam.i686
libattr.i686
libacl.i686
%end

%post --nochroot
# Add necessary scripts or commands here to your need
# This generated kickstart file is only for the automated
```

```
installation of products in the DVD

PATH=$PATH:/sbin:/usr/sbin:/bin:/usr/bin
export PATH

#
# Notice:
# * Modify the BUILDSRC below according to your real environment
# * The location specified with BUILDSRC should be NFS accessible
#   to the Kickstart Server
# * Copy the whole directories of rpms from installation media
#   to the BUILDSRC
#

BUILDSRC="<hostname_or_ip>:/path/to/rpms"

#
# Notice:
# * You do not have to change the following scripts
#

# define path varibles
ROOT=/mnt/sysimage
BUILDDIR="${ROOT}/build"
RPMDIR="${BUILDDIR}/rpms"

# define log path
KSLOG="${ROOT}/opt/VRTStmp/kickstart.log"

echo "==== Executing kickstart post section: ====" >> ${KSLOG}

mkdir -p ${BUILDDIR}
mount -t nfs -o nolock,vers=3 ${BUILDSRC} ${BUILDDIR} >> ${KSLOG} 2>&1

# install rpms one by one
for RPM in VRTSperl VRTSvlic VRTSspt VRTSvxvm VRTSaslapm VRTSvxfs
VRTSfsadv VRTSllt VRTSgab VRTSvxfen VRTSamf VRTSvcs VRTScps VRTSvcsag
VRTSvcsea VRTSdbed VRTSglm VRTScavf VRTSgms VRTSodm VRTSdbac VRTSsfmh
VRTSvbs VRTSsfcpi VRTSvcswiz
do
    echo "Installing package  -- $RPM" >> ${KSLOG}
    rpm -U -v --root ${ROOT} ${RPMDIR}/${RPM}-* >> ${KSLOG} 2>&1
done
```

```
umount ${BUILDDIR}

CALLED_BY=KICKSTART ${ROOT}/opt/VRTS/install/bin/UXRT7.4.2/
add_install_scripts >> ${KSLOG} 2>&1

echo "==== Completed kickstart file ====" >> ${KSLOG}

exit 0
%end
```

Installing Veritas InfoScale using yum

You can install Veritas InfoScale using yum. yum is supported for Red Hat Enterprise operating system.

To install Veritas InfoScale using yum

1 Configure a yum repository on a client system.

- Create a `.repo` file under `/etc/yum.repos.d/`. An example of this `.repo` file for Veritas InfoScale is:

  ```
  # cat /etc/yum.repos.d/veritas_infoscale7.repo
  [repo-Veritas InfoScale]
  name=Repository for Veritas InfoScale
  baseurl=file:///path/to/repository/
  enabled=1
  gpgcheck=1
  gpgkey=file:///path/to/repository/RPM-GPG-KEY-veritas-infoscale7
  ```

  The values for the `baseurl` attribute can start with http://, ftp://, or file:///. The URL you choose needs to be able to access the `repodata` directory. It also needs to access all the Veritas InfoScale RPMs in the repository that you create or update.

- Run the following commands to get the yum repository updated:

  ```
  # yum repolist

  # yum updateinfo
  ```

- Check the yum group information:

  ```
  # yum grouplist | grep 7.4.2
     AVAILABILITY7.4.2
  ```

```
   ENTERPRISE7.4.2
   FOUNDATION7.4.2
   STORAGE7.4.2

# yum groupinfo AVAILABILITY7.4.2

# yum groupinfo FOUNDATION7.4.2

# yum groupinfo STORAGE7.4.2

# yum groupinfo ENTERPRISE7.4.2
```

- Check the yum configuration. List Veritas InfoScale RPMs.

```
# yum list 'VRTS*'
Available Packages
VRTSperl.x86_64        5.16.1.4-RHEL5.2
VRTSsfcpi.noarch       7.4.2.0.000-GENERIC
VRTSvlic.x86_64        3.02.7.4.2.010-0
...
```

The Veritas InfoScale RPMs may not be visible immediately if:

- The repository was visited before the Veritas InfoScale RPMs were added, and
- The local cache of its metadata has not expired.

To eliminate the local cache of the repositories' metadata and get the latest information from the specified `baseurl`, run the following commands:

```
# yum clean expire-cache
# yum list 'VRTS*'
```

Refer to the *Red Hat Enterpirse Linux Deployment Guide* for more information on yum repository configuration.

2 Install the RPMs on the target systems.

- **To install all the RPMs**

  1. Specify each RPM name as its yum equivalent. For example:

     ```
     # yum install VRTSvlic VRTSperl ... VRTSsfcpi
     ```

  2. Specify all of the Veritas InfoScale RPMs using its RPM glob. For example:

     ```
     # yum install 'VRTS*'
     ```

3. Specify the group name if a group is configured for Veritas InfoScale's RPMs. This name should keep consistency with the one in xml file. In this example, the group name is *ENTERPRISE7.4.2*:

```
# yum install @ENTERPRISE7.4.2
```

Or

```
# yum groupinstall -y ENTERPRISE7.4.2
```

- **To install one RPM at a time**

1. Run the `installer -allpkgs` command to determine RPM installation order.

```
# ./installer -allpkgs

InfoScale Foundation: PKGS: VRTSperl VRTSvlic VRTSspt
VRTSveki VRTSvxvm VRTSaslapm VRTSvxfs VRTSsfmh VRTSsfcpi

InfoScale Availability: PKGS: VRTSperl VRTSvlic VRTSspt
VRTSveki VRTSllt VRTSgab VRTSvxfen VRTSamf VRTSvcs VRTScps
VRTSvcsag VRTSvcsea VRTSsfmh VRTSvbs VRTSvcswiz VRTSsfcpi

InfoScale Storage: PKGS: VRTSperl VRTSvlic VRTSspt VRTSveki
VRTSvxvm VRTSaslapm VRTSvxfs VRTSfsadv VRTSllt VRTSgab VRTSvxfen
VRTSamf VRTSvcs VRTScps VRTSvcsag VRTSdbed VRTSglm VRTScavf
VRTSgms VRTSodm VRTSsfmh VRTSsfcpi

InfoScale Enterprise: PKGS: VRTSperl VRTSvlic VRTSspt VRTSveki
VRTSvxvm VRTSaslapm VRTSvxfs VRTSfsadv VRTSllt VRTSgab VRTSvxfen
VRTSamf VRTSvcs VRTScps VRTSvcsag VRTSvcsea VRTSdbed VRTSglm
VRTScavf VRTSgms VRTSodm VRTSdbac VRTSsfmh VRTSvbs VRTSvcswiz
VRTSsfcpi
```

2. Use the same order as the output from the `installer -allpkgs` command:

```
# yum install VRTSperl
# yum install VRTSvlic
  ...
# yum install VRTSsfcpi
```

3 After you install all the RPMs, use the `/opt/VRTS/install/installer` command to license, configure, and start the product.

If the `VRTSsfcpi` RPM is installed before you use yum to install Veritas InfoScale, the RPM is not upgraded or uninstalled. If the `/opt/VRTS/install/installer` script is not created properly, use the `/opt/VRTS/install/bin/UXRT7.4.2/add_install_scripts` script after all the other Veritas InfoScale RPMs are installed. For example, your output may be similar to the following, depending on the products you install:

```
# /opt/VRTS/install/bin/UXRT7.4.2/add_install_scripts
Creating install/uninstall scripts for installed products
Creating /opt/VRTS/install/installer for UXRT7.4.2
Creating /opt/VRTS/install/showversion for UXRT7.4.2
```

To uninstall Veritas InfoScale using yum

- You can uninstall Veritas InfoScale using yum. Use one of the following commands depending on the product that you have installed:

```
# yum groupremove -y AVAILABILITY7.4.2

# yum groupremove -y FOUNDATION7.4.2

# yum groupremove -y STORAGE7.4.2

# yum groupremove -y ENTERPRISE7.4.2
```

Installing Veritas InfoScale using the Red Hat Satellite server

You can install Veritas InfoScale using the Red Hat Satellite server. Red Hat Satellite is supported for Red Hat Enterprise Linux operating system. You can install RPMs and rolling patches on the systems which the Red Hat Satellite server manages.

Red Hat Satellite server is a systems management solution. It lets you:

- Inventory the hardware and the software information of your systems.
- Install and update software on systems.
- Collect and distribute custom software RPMs into manageable groups.
- Provision (Kickstart) systems.
- Manage and deploy configuration files to systems.
- Monitor your systems.
- Provision virtual guests.
- Start, stop, and configure virtual guests.

In a Red Hat Satellite server, you can manage the system by creating a channel. A Red Hat Satellite channel is a collection of software RPMs. Using channels, you can segregate the RPMs by defining some rules. For instance, a channel may contain RPMs only from a specific Red Hat distribution. You can define channels according to your own requirement. You can create a channel that contains Veritas InfoScale RPMs for custom usage in your organization's network.

Channels are of two types:

- Base channel
  A base channel consists of RPMs based on a specific architecture and Red Hat Enterprise Linux release.
- Child channel
  A child channel is a channel which is associated with a base channel that contains extra custom RPMs like Veritas InfoScale.

A system can subscribe to only one base channel and multiple child channels of its base channel. The subscribed system can only install or update the RPMs that are available through its satellite channels.

For more information, see the *Red Hat Satellite5.6 User Guide*.

Using Red Hat Satellite server to install Veritas InfoScale products

You can use the Red Hat Satellite server to install Veritas InfoScale products on your system.

To use Red Hat Satellite server to install Veritas InfoScale products

1. Set the base channel, child channel, and target system by following the Red Hat Satellite documentation. You need to ensure that:
   - The base channel consists of RPMs based on RHEL6.8, RHEL6.9, RHEL6.10, RHEL7.4, RHEL7.5, or the RHEL7.6 release

- The child channel consists of Veritas InfoScale RPMs or patches.
- The target system is registered to the Red Hat Satellite.

2. Log on to the Red Hat Satellite admin page. Select the **Systems** tab. Click on the **target system**.
3. Select **Alter Channel Subscriptions** to alter the channel subscription of the target system.
4. Select the channel which contains the repository of Veritas InfoScale.
5. Enter the following command to check the YUM repository on the target system.

```
# yum repolist
```

6. Enter the following command to install the Veritas InfoScale RPMs using YUM:

```
# yum install @ENTERPRISE7.4.2
```

7. Enter the following command to generate the script of the installer:

```
# /opt/VRTS/install/bin/UXRT7.4.2/add_install_scripts
```

8. Enter the following command to configure Veritas InfoScale using the installer:

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

Chapter 8

Completing the post installation tasks

This chapter includes the following topics:

- Verifying product installation
- Setting environment variables
- Managing the Veritas telemetry collector on your server
- Next steps after installation

Verifying product installation

To verify the version of the installed product, use the following command:

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

To find out about the installed RPMs and its versions, use the following command:

```
# /opt/VRTS/install/showversion
```

After every product installation, the installer creates an installation log file and a 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. Veritas recommends that you keep the files for auditing, debugging, and future use.

The installation log file contains all commands that are executed during the procedure, their output, and the errors generated by the commands.

The summary file contains the results of the installation by the installer or the product installation scripts. The summary includes the list of the RPMs, and the status (success or failure) of each RPM, and information about the processes that 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, Veritas InfoScale commands are in `/opt/VRTS/bin`. Veritas InfoScale manual pages are stored in `/opt/VRTS/man`.

Specify `/opt/VRTS/bin` in your `PATH` after the path to the standard Linux commands.

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.

To invoke the VxFS-specific `df`, `fsdb`, `ncheck`, or `umount` commands, type the full path name: `/opt/VRTS/bin/`*command*.

To set your `MANPATH` environment variable to include `/opt/VRTS/man` do the following:

- If you want to use a shell such as `sh` or `bash`, enter the following:

  ```
  $ MANPATH=$MANPATH:/opt/VRTS/man; export MANPATH
  ```

- If you want to use a shell such as `csh` or `tcsh`, enter the following:

  ```
  % setenv MANPATH $(MANPATH):/opt/VRTS/man
  ```

On a Red Hat system, also include the 1m manual page section in the list defined by your `MANSECT` environment variable.

- If you want to use a shell such as `sh` or `bash`, enter the following:

  ```
  $ MANSECT=$MANSECT:1m; export MANSECT
  ```

- If you want to use a shell such as `csh` or `tcsh`, enter the following:

  ```
  % setenv MANSECT $(MANSECT):1m
  ```

  If you use the `man`(1) command to access manual pages, set `LC_ALL=C` in your shell to ensure that they display correctly.

Managing the Veritas telemetry collector on your server

The Veritas telemetry collector on your server sends telemetry data to the edge server. You can manage the Veritas telemetry collector on each of your servers by using the `/opt/VRTSvlic/tele/bin/TelemetryCollector` command. See the following table for a list of operations that you can perform to manage the Veritas telemetry collector along with examples of each of the commands.

Table 8-1 Commands used to manage the collector

Operation	Description

Register an edge server and start the collector	Use the following command to register an edge server with the collector. Note that the collector is automatically started after running this command.
	`/opt/VRTSvlic/tele/bin/TelemetryCollector -start --hostname=<hostname_or_IP> --port=<port_number>`
	■ *<hostname_or_IP>* is the host name or IP address of the edge server you want to register.
	■ *<port_number>* is the port number of the edge server that is used for communication.
	Example:
	`/opt/VRTSvlic/tele/bin/TelemetryCollector -start --hostname=telemetry.veritas.com --port=443`

Start the collector (if the collector is not already running)	Use the following command if you want to start a collector that is not sending telemetry data to the edge server. `/opt/VRTSvlic/tele/bin/TelemetryCollector -start`
Restart the collector (if the collector is already running)	Use the following command to restart the collector that is sending telemetry data to the edge server. `/opt/VRTSvlic/tele/bin/TelemetryCollector -restart`
Check whether the collector is running or not	Use the following command to check the status of the collector on your server. `/opt/VRTSvlic/tele/bin/TelemetryCollector -status`
Check whether the collector is registered with an edge server	Use the following command to check whether the collector is registered with an edge server. `/opt/VRTSvlic/tele/bin/TelemetryCollector -registered`

Table 8-1 Commands used to manage the collector *(continued)*

Operation	Description
Update the host name, IP address, or port number of the edge server	Use the following command to update the host name or IP address and port number of the edge server. `/opt/VRTSvlic/tele/bin/TelemetryCollector -update --hostname=<hostname_or_IP> --port=<port_number>` ■ *<hostname_or_IP>* is the host name or IP address of the edge server you want to register. ■ *<port_number>* is the port number of the edge server that is used for communication. **Examples**: ■ `/opt/VRTSvlic/tele/bin/TelemetryCollector -update --hostname=telemetry.veritas.com --port=443` ■ `/opt/VRTSvlic/tele/bin/TelemetryCollector -update --port=443`

Table 8-1 Commands used to manage the collector *(continued)*

Operation	Description
Configure how often telemetry data is sent to the edge server	Use either of the following commands to define how often you want the collector to send telemetry data to the edge server. ■ To send telemetry data once every month: `/opt/VRTSvlic/tele/bin/TelemetryCollector -update --d=<day>` ■ To send telemetry data once every day: `/opt/VRTSvlic/tele/bin/TelemetryCollector -update --h=<hour>` ■ To send telemetry data once every week: `/opt/VRTSvlic/tele/bin/TelemetryCollector -update wday=<weekday>` Use the following variables in the above commands. ■ *<day>* is the day of the month that you want the collector to send telemetry data. ■ *<hour>* is the hour when you want the collector to send the telemetry data.. ■ *<weekday>* is the day of the week that you want the collector to send telemetry data. **Examples**: ■ `/opt/VRTSvlic/tele/bin/TelemetryCollector -update --wday=MON` ■ `/opt/VRTSvlic/tele/bin/TelemetryCollector -update --h=6` ■ `/opt/VRTSvlic/tele/bin/TelemetryCollector -update --d=14`

Next steps after installation

Once installation is complete, you can configure a component of your choice.

Table 8-2 lists the components and the respective Configuration and Upgrade guides that are available.

Table 8-2 Guides available for configuration

Component	Document name
Storage Foundation	See *Storage Foundation Configuration and Upgrade Guide* See *Storage Foundation Administrator's Guide*
Storage Foundation and High Availability	See *Storage Foundation and High Availability Configuration and Upgrade Guide*
Storage Foundation Cluster File System HA	See *Storage Foundation Cluster File System High Availability Configuration and Upgrade Guide* See *Storage Foundation Cluster File System High Availability Administrator's Guide*
Cluster Server	See *Cluster Server Configuration and Upgrade Guide* See *Cluster Server Administrator's Guide*
Storage Foundation for Oracle RAC	See *Storage Foundation for Oracle RAC Configuration and Upgrade Guide* See *Storage Foundation for Oracle RAC Administrator's Guide*
Storage Foundation for Sybase SE	See *Storage Foundation for Sybase ASE CE Configuration and Upgrade Guide* See *Storage Foundation for Sybase ASE CE Administrator's Guide*

Section 3

Uninstallation of Veritas InfoScale

- Chapter 9. Uninstalling Veritas InfoScale using the installer
- Chapter 10. Uninstalling Veritas InfoScale using response files

Chapter 9

Uninstalling Veritas InfoScale using the installer

This chapter includes the following topics:

- Removing VxFS file systems
- Removing rootability
- Moving volumes to disk partitions
- Removing the Replicated Data Set
- Uninstalling Veritas InfoScale RPMs using the installer
- Removing the Storage Foundation for Databases (SFDB) repository

Removing VxFS file systems

The VxFS RPM cannot be removed if there are any mounted VxFS file systems. Unmount all VxFS file systems before removing the RPM. After you remove the VxFS RPM, VxFS file systems are not mountable or accessible until another VxFS RPM is installed. It is advisable to back up VxFS file systems before installing a new VxFS RPM. If VxFS will not be installed again, all VxFS file systems must be converted to a new file system type.

To remove VxFS file systems

1. Check if any VxFS file systems or Storage Checkpoints are mounted:

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

2. Make backups of all data on the file systems that you wish to preserve, or recreate them as non-VxFS file systems on non-VxVM volumes or partitions.
3. Unmount all Storage Checkpoints and file systems:

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

4. Comment out or remove any VxFS file system entries from the `/etc/fstab` file.

Removing rootability

Perform this procedure if you configured rootability by encapsulating the root disk.

To remove rootability

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

   ```
   # df -v /
   ```

   The root disk is under VxVM control if `/dev/vx/dsk/rootdg/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

2. Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

   For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

   ```
   # vxplex -o rm dis mirrootvol-01 mirswapvol-01
   ```

   Warning: Do not remove the plexes that correspond to the original root disk partitions.

3. Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices:

   ```
   # /etc/vx/bin/vxunroot
   ```

   Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

Moving volumes to disk partitions

All volumes must be moved to disk partitions.

This can be done using one of the following procedures:

- Back up the system fully onto tape and then recover from it.
- Back up each file system individually and then recover them all after creating new file systems on disk partitions.
- Use VxVM to move volumes incrementally onto disk partitions as described in the following section.

Moving volumes onto disk partitions using VxVM

Use the following procedure to move volumes onto disk partitions.

To move volumes onto disk partitions

1. Evacuate disks using the `vxdiskadm` program or the `vxevac` script. You should consider the amount of target disk space required for this before you begin.

   Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.

2. Remove the evacuated disks from VxVM control using the following commands:

   ```
   # vxdg -g diskgroup rmdisk disk _media_name
   # vxdisk rm disk_access_name
   ```

3. Decide which volume to move first. If the volume to be moved is mounted, unmount it.

4. If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that data on the volume is synced.

5. Create a partition on free disk space of the same size as the volume. If there is not enough free space for the partition, a new disk must be added to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this volume.

6. `Copy` the data on the volume onto the newly created disk partition using a command similar to the following:

   ```
   # dd if=/dev/vx/dsk/diskgroup/volume-name of=/dev/sdb2
   ```

   where `sdb` is the disk outside of VxVM and `2` is the newly created partition on that disk.

7. Replace the entry for that volume (if present) in `/etc/fstab` with an entry for the newly created partition.

8. Mount the disk partition if the corresponding volume was previously mounted.

9. Stop the volume and remove it from VxVM using the following commands:

   ```
   # vxvol -g diskgroup -f stop volume_name
   # vxedit -g diskgroup -rf rm volume_name
   ```

10. Remove any disks that have become free (have no subdisks defined on them) by removing volumes from VxVM control. To check if there are still some subdisks remaining on a particular disk, use the following command:

    ```
    # vxprint -F "%sdnum" disk_media_name
    ```

11 If the output is not 0, there are still some subdisks on this disk that must be subsequently removed. If the output is 0, remove the disk from VxVM control using the following commands:

```
# vxdg -g diskgroup rmdisk disk_media_name
# vxdisk rm disk_access_name
```

12 The free space now created can be used for adding the data in the next volume to be removed.

13 After all volumes have been converted into disk partitions successfully, reboot the system. After the reboot, none of the volumes should be open. To verify that none of the volumes are open, use the following command:

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

14 If any volumes remain open, repeat the steps listed above.

Removing the Replicated Data Set

If you use VVR, you need to perform the following steps. This section gives the steps to remove a Replicated Data Set (RDS) when the application is either active or stopped.

Note: If you are upgrading Volume Replicator, do not remove the Replicated Data Set.

To remove the Replicated Data Set

1. Verify that all RLINKs are up-to-date:

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

   If the Secondary is not required to be up-to-date, proceed to 2 and stop replication using the `-f` option with the `vradmin stoprep` command.

2. Stop replication to the Secondary by issuing the following command on any host in the RDS:

   The `vradmin stoprep` command fails if the Primary and Secondary RLINKs are not up-to-date. Use the `-f` option to stop replication to a Secondary even when the RLINKs are not up-to-date.

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

   The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

   The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

3. Remove the Secondary from the RDS by issuing the following command on any host in the RDS:

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

   The argument local_rvgname is the name of the RVG on the local host and represents its RDS.

   The argument sec_hostname is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

4. Remove the Primary from the RDS by issuing the following command on the Primary:

   ```
   # vradmin -g diskgroup delpri local_rvgname
   ```

   When used with the `-f` option, the `vradmin delpri` command removes the Primary even when the application is running on the Primary.

   The RDS is removed.

5. If you want to delete the SRLs from the Primary and Secondary hosts in the RDS, issue the following command on the Primary and all Secondaries:

   ```
   # vxedit -r -g diskgroup rm srl_name
   ```

Uninstalling Veritas InfoScale RPMs using the installer

Use the following procedure to remove Veritas InfoScale products.

Not all RPMs 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 Veritas InfoScale 7.4.2 with a previous version of Veritas InfoScale.

To shut down and remove the installed Veritas InfoScale RPMs

1. Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/fstab`. Failing to remove these entries could result in system boot problems later.
2. Unmount all mount points for VxFS file systems.

   ```
   # umount /mount_point
   ```

3. If the VxVM RPM (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

   See "Removing rootability" on page 82.

4. If a cache area is online, you must take the cache area offline before uninstalling the VxVM RPM. Use the following command to take the cache area offline:

   ```
   # sfcache offline cachename
   ```

5. Make sure you have performed all of the prerequisite steps.
6. 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
   ```

7 Move to the `/opt/VRTS/install` directory and run the uninstall script.

```
# cd /opt/VRTS/install

# ./installer -uninstall
```

8 The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall Veritas InfoScale.

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

9 The uninstall script prompts you to stop the product processes. If you respond yes, the processes are stopped and the RPMs are uninstalled.

The uninstall script creates log files and displays the location of the log files.

10 Most RPMs have kernel components. In order to ensure complete removal, a system reboot is recommended after all RPMs have been removed.

11 In case the uninstallation fails to remove any of the VRTS RPMs, check the installer logs for the reason for failure or try to remove the RPMs manually using the following command:

```
# rpm -e VRTSvxvm
```

Removing the Storage Foundation for Databases (SFDB) repository

After removing the product, you can remove the SFDB repository file and any backups.

Removing the SFDB repository file disables the SFDB tools.

To remove the SFDB repository

1. Identify the SFDB repositories created on the host.

   Oracle:

   ```
   # cat /var/vx/vxdba/rep_loc
   ```

   ```
   {
      "sfae_rept_version" : 1,
      "oracle" : {
         "SFAEDB" : {
            "location" : "/data/sfaedb/.sfae",
            "old_location" : "",
            "alias" : [
               "sfaedb"
            ]
         }
      }
   }
   ```

2. Remove the directory identified by the `location` key.

   Oracle:

   ```
   # rm -rf /data/sfaedb/.sfae
   ```

   DB2 9.5 and 9.7:

   ```
   # rm -rf /db2data/db2inst1/NODE0000/SQL00001/.sfae
   ```

   DB2 10.1 and 10.5:

   ```
   # rm -rf /db2data/db2inst1/NODE0000/SQL00001/MEMBER0000/.sfae
   ```

3. Remove the repository location file.

   ```
   # rm -rf /var/vx/vxdba/rep_loc
   ```

   This completes the removal of the SFDB repository.

Chapter 10

Uninstalling Veritas InfoScale using response files

This chapter includes the following topics:

- Uninstalling Veritas InfoScale using response files
- Response file variables to uninstall Veritas InfoScale
- Sample response file for Veritas InfoScale uninstallation

Uninstalling Veritas InfoScale using response files

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

To perform an automated uninstallation

1. Make sure that you meet the prerequisites to uninstall Veritas InfoScale.
2. Copy the response file to the system where you want to uninstall Veritas InfoScale.
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

Table 10-1 lists the response file variables that you can define to configure Veritas InfoScale.

Table 10-1 Response file variables for uninstalling Veritas InfoScale

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 RPMs that are needed during the install. The default location is /opt/VRTStmp. List or scalar: scalar Optional or required: optional
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional

Table 10-1 Response file variables for uninstalling Veritas InfoScale *(continued)*

Variable	Description
CFG{opt}{uninstall}	Uninstalls Veritas InfoScale RPMs. List or scalar: scalar Optional or required: optional

Sample response file for Veritas InfoScale uninstallation

The following example shows a response file for uninstalling Veritas InfoScale

```
our %CFG;

$CFG{opt}{uninstall}=1;
$CFG{opt}{vr}=1;
$CFG{prod}="ENTERPRISE742";
$CFG{systems}=[ qw("system1", "system2") ];

1;
```

Section 4

Installation reference

- Appendix A. Installation scripts
- Appendix B. Tunable files for installation
- Appendix C. 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
-allpkgs	Displays all RPMs required for the specified product. The RPMs are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-comcleanup	The `-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.
-configure	Configures the product after installation.

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

Command Line Option	Function
-disable_dmp_native_support	Disables Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases RPM upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.
–hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.
-install	Used to install products on system
-online_upgrade	Used to perform online upgrade. Using this option, the installer upgrades the whole cluster and also supports customer's application zero down time during the upgrade procedure. Now this option is supported only in VCS.
-patch_path	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-patch2_path	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch3_path	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch4_path	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch5_path	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.

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

Command Line Option	Function
–keyfile *ssh_key_file*	Specifies a key file for secure shell (SSH) installs. This option passes `-I ssh_key_file` to every SSH invocation.
–kickstart *dir_path*	Produces a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of required RPMs in the correct order for installing, in a format that can be used for Kickstart installations. The *dir_path* indicates the path to the directory in which to create the file.
-license	Registers or updates product licenses on the specified systems.
–logpath *log_path*	Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved.
-noipc	Disables the installer from making outbound networking calls to Veritas Services and Operations Readiness Tool (SORT) in order to automatically obtain patch and release information updates.
-nolic	Allows installation of product RPMs without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
-pkgtable	Displays product's RPMs in correct installation order by group.
–postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Veritas recommends doing a precheck before installing a product.
-prod	Specifies the product for operations.
-component	Specifies the component for operations.

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

Command Line Option	Function
-redirect	Displays progress details without showing the progress bar.
-require	Specifies an installer patch file.
-requirements	The `-requirements` option displays required OS version, required RPMs 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.
-rolling_upgrade	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
-rollingupgrade_phase1	The `-rollingupgrade_phase1` option is used to perform rolling upgrade Phase-I. In the phase, the product kernel RPMs get upgraded to the latest version.
-rollingupgrade_phase2	The `-rollingupgrade_phase2` option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.
-rsh	Specify this option when you want to use rsh and RCP for communication between systems instead of the default ssh and SCP.
–serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.

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

Command Line Option	Function
-settunables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the `-tunablesfile` option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-timeout	The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the `-timeout` option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 prevents the script from timing out. The `-timeout` option does not work with the `-serial` option
–tmppath *tmp_path*	Specifies a directory other than `/opt/VRTStmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where RPMs are copied on remote systems before installation.
-tunables	Lists all supported tunables and create a tunables file template.
-tunablesfile *tunables_file*	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-uninstall	This option is used to uninstall the products from systems
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.

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

Command Line Option	Function
-version	Checks and reports the installed products and their versions. Identifies the installed and missing RPMs and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing RPMs and patches where applicable. Lists the installed patches, patches, and available updates for the installed product if an Internet connection is available.
-yumgroupxml	The `-yumgroupxml` option is used to generate a yum group definition XML file. The `createrepo` command can use the file on Redhat Linux to create a yum group for automated installation of all RPMs for a product. An available location to store the XML file should be specified as a complete path. The `-yumgroupxml` option is supported on RHEL and supported RHEL-compatible distributions only.

Appendix B

Tunable files for installation

This appendix includes the following topics:

- About setting tunable parameters using the installer or a response file
- Setting tunables for an installation, configuration, or upgrade
- Setting tunables with no other installer-related operations
- Setting tunables with an un-integrated response file
- Preparing the tunables file
- Setting parameters for the tunables file
- Tunables value parameter definitions

About setting tunable parameters using the installer or a response file

You can set non-default product and system tunable parameters using a tunables file. With the file, you can set tunables such as the I/O policy or toggle native multi-pathing. The tunables file passes arguments to the installer script to set tunables. With the file, you can set the tunables for the following operations:

- When you install, configure, or upgrade systems.

  ```
  # ./installer -tunablesfile tunables_file_name
  ```

  See "Setting tunables for an installation, configuration, or upgrade" on page 101.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -settunables [
sys1 sys2 ...]
```

See “Setting tunables with no other installer-related operations” on page 102.

- When you apply the tunables file with an un-integrated response file.

```
# ./installer -responsefile response_file_name -tunablesfile
tunables_file_name
```

See “Setting tunables with an un-integrated response file” on page 103.

See “About response files” on page 56.

You must select the tunables that you want to use from this guide.

See “Tunables value parameter definitions” on page 105.

Setting tunables for an installation, configuration, or upgrade

You can use a tunables file for installation procedures to set non-default tunables. You invoke the installation script with the `tunablesfile` option. The tunables file passes arguments to the script to set the selected tunables. You must select the tunables that you want to use from this guide.

See “Tunables value parameter definitions” on page 105.

Note: Certain tunables only take effect after a system reboot.

To set the non-default tunables for an installation, configuration, or upgrade

1. Prepare the tunables file.

   See “Preparing the tunables file” on page 104.

2. Make sure the systems where you want to install Veritas InfoScale meet the installation requirements.
3. Complete any preinstallation tasks.
4. Copy the tunables file to one of the systems where you want to install, configure, or upgrade the product.
5. Mount the product disc and navigate to the directory that contains the installation program.

6 Start the installer for the installation, configuration, or upgrade. For example:

```
# ./installer -tunablesfile /tmp/tunables_file
 -settunables [sys1 sys2 ...]
```

Where /tmp/*tunables_file* is the full path name for the tunables file.

7 Proceed with the operation. When prompted, accept the tunable parameters.

Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.

8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with no other installer-related operations

You can use the installer to set tunable parameters without any other installer-related operations. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See "Tunables value parameter definitions" on page 105.

Note: Certain tunables only take effect after a system reboot.

To set tunables with no other installer-related operations

1 Prepare the tunables file.

See "Preparing the tunables file" on page 104.

2 Make sure the systems where you want to install Veritas InfoScale meet the installation requirements.

3 Complete any preinstallation tasks.

4 Copy the tunables file to one of the systems that you want to tune.

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

6 Start the installer with the `-settunables` option.

```
# ./installer -tunablesfile tunables_file_name -settunables [
sys123 sys234 ...]
```

Where /tmp/*tunables_file* is the full path name for the tunables file.

7 Proceed with the operation. When prompted, accept the tunable parameters.

Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.

8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with an un-integrated response file

You can use the installer to set tunable parameters with an un-integrated response file. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See "Tunables value parameter definitions" on page 105.

Note: Certain tunables only take effect after a system reboot.

To set tunables with an un-integrated response file

1 Make sure the systems where you want to install Veritas InfoScale meet the installation requirements.

2 Complete any preinstallation tasks.

3 Prepare the tunables file.

See "Preparing the tunables file" on page 104.

4 Copy the tunables file to one of the systems that you want to tune.

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

6 Start the installer with the `-responsefile` and `-tunablesfile` options.

```
# ./installer -responsefile response_file_name -tunablesfile
tunables_file_name
```

Where *response_file_name* is the full path name for the response file and *tunables_file_name* is the full path name for the tunables file.

7 Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.

8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Preparing the tunables file

A tunables file is a Perl module and consists of an opening and closing statement, with the tunables defined between. Use the hash symbol at the beginning of the line to comment out the line. The tunables file opens with the line "our %TUN;" and ends with the return true "1;" line. The final return true line only needs to appear once at the end of the file. Define each tunable parameter on its own line.

You can use the installer to create a tunables file template, or manually format tunables files you create.

To create a tunables file template

- Start the installer with the `-tunables` option. Enter the following:

  ```
  # ./installer -tunables
  ```

  You see a list of all supported tunables, and the location of the tunables file template.

To manually format tunables files

- Format the tunable parameter as follows:

  ```
  $TUN{"tunable_name"}{"system_name"|"*"}=value_of_tunable;
  ```

For the *system_name*, use the name of the system, its IP address, or a wildcard symbol. The *value_of_tunable* depends on the type of tunable you are setting. End the line with a semicolon.

The following is an example of a tunables file.

```
#
# Tunable Parameter Values:
#
our %TUN;

$TUN{"tunable1"}{"*"}=1024;
$TUN{"tunable3"}{"sys123"}="SHA256";

1;
```

Setting parameters for the tunables file

Each tunables file defines different tunable parameters. The values that you can use are listed in the description of each parameter. Select the tunables that you want to add to the tunables file and then configure each parameter.

See "Tunables value parameter definitions" on page 105.

Each line for the parameter value starts with $TUN. The name of the tunable is in curly brackets and double-quotes. The system name is enclosed in curly brackets and double-quotes. Finally define the value and end the line with a semicolon, for example:

```
$TUN{"dmp_daemon_count"}{"node123"}=16;
```

In this example, you are changing the dmp_daemon_count value from its default of 10 to 16. You can use the wildcard symbol "*" for all systems. For example:

```
$TUN{"dmp_daemon_count"}{"*"}=16;
```

Tunables value parameter definitions

When you create a tunables file for the installer you can only use the parameters in the following list.

Prior to making any updates to the tunables, refer to the *Storage Foundation Cluster File System High Availability Administrator's Guide* for detailed information on product tunable ranges and recommendations.

Table B-1 describes the supported tunable parameters that can be specified in a tunables file.

Table B-1 Supported tunable parameters

Tunable	Description
autoreminor	(Veritas Volume Manager) Enable reminoring in case of conflicts during disk group import.
autostartvolumes	(Veritas Volume Manager) Enable the automatic recovery of volumes.
dmp_cache_open	(Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached.
dmp_daemon_count	(Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks.
dmp_delayq_interval	(Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy.

Table B-1 Supported tunable parameters *(continued)*

Tunable	Description
dmp_fast_recovery	(Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Dynamic Multi-Pathing is started.
dmp_health_time	(Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy.
dmp_log_level	(Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed.
dmp_low_impact_probe	(Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled.
dmp_lun_retry_timeout	(Dynamic Multi-Pathing) The retry period for handling transient errors.
dmp_monitor_fabric	(Dynamic Multi-Pathing) Whether the Event Source daemon (`vxesd`) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Dynamic Multi-Pathing is started.
dmp_monitor_ownership	(Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored.
dmp_native_support	(Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices.
dmp_path_age	(Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy.
dmp_pathswitch_blks_shift	(Dynamic Multi-Pathing) The default number of contiguous I/O blocks sent along a DMP path to an array before switching to the next available path.
dmp_probe_idle_lun	(Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs.
dmp_probe_threshold	(Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon.

Table B-1 Supported tunable parameters *(continued)*

Tunable	Description
dmp_restore_cycles	(Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic.
dmp_restore_interval	(Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths.
dmp_restore_policy	(Dynamic Multi-Pathing) The policy used by DMP path restoration thread.
dmp_restore_state	(Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started.
dmp_retry_count	(Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed.
dmp_scsi_timeout	(Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP.
dmp_sfg_threshold	(Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature.
dmp_stat_interval	(Dynamic Multi-Pathing) The time interval between gathering DMP statistics.
fssmartmovethreshold	(Veritas Volume Manager) The file system usage threshold for SmartMove (percent). This tunable must be set after Veritas Volume Manager is started.
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer can only set the system default value of max_diskq. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.

Table B-1 Supported tunable parameters *(continued)*

Tunable	Description
read_ahead	(Veritas File System) The 0 value disables read ahead functionality, the 1 value (default) retains traditional sequential read ahead behavior, and the 2 value enables enhanced read ahead for all reads. The installer can only set the system default value of read_ahead. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
read_nstream	(Veritas File System) The number of parallel read requests of size read_pref_io that can be outstanding at one time. The installer can only set the system default value of read_nstream. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
read_pref_io	(Veritas File System) The preferred read request size. The installer can only set the system default value of read_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
reclaim_on_delete_start_time	(Veritas Volume Manager) Time of day to start reclamation for deleted volumes. This tunable must be set after Veritas Volume Manager is started.
reclaim_on_delete_wait_period	(Veritas Volume Manager) Days to wait before starting reclamation for deleted volumes. This tunable must be set after Veritas Volume Manager is started.
same_key_for_alldgs	(Veritas Volume Manager) Use the same fencing key for all disk groups. This tunable must be set after Veritas Volume Manager is started.
sharedminorstart	(Veritas Volume Manager) Start of range to use for minor numbers for shared disk groups. This tunable must be set after Veritas Volume Manager is started.

Table B-1 Supported tunable parameters *(continued)*

Tunable	Description
storage_connectivity	(Veritas Volume Manager) The CVM storage connectivity type. This tunable must be set after Veritas Volume Manager is started.
usefssmartmove	(Veritas Volume Manager) Configure SmartMove feature (all, thinonly, none). This tunable must be set after Veritas Volume Manager is started.
vol_checkpt_default	(Veritas File System) Size of VxVM storage checkpoints (kBytes). This tunable requires a system reboot to take effect.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for Volume Replicator.
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires a system reboot to take effect.
vol_fmr_logsz	(Veritas Volume Manager) Maximum size of bitmap Fast Mirror Resync uses to track changed blocks (KBytes). This tunable requires a system reboot to take effect.
vol_max_adminio_poolsz	(Veritas Volume Manager) Maximum amount of memory used by VxVM admin I/O's (bytes). This tunable requires a system reboot to take effect.
vol_max_nmpool_sz	(Veritas Volume Manager) Maximum name pool size (bytes).
vol_max_rdback_sz	(Veritas Volume Manager) Storage Record readback pool maximum (bytes).
vol_max_wrspool_sz	(Veritas Volume Manager) Maximum memory used in clustered version of Volume Replicator .

Table B-1 Supported tunable parameters *(continued)*

Tunable	Description
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (kBytes). This tunable requires a system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires a system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires a system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (kBytes). This tunable requires a system reboot to take effect.
vol_min_lowmem_sz	(Veritas Volume Manager) Low water mark for memory (bytes).
vol_nm_hb_timeout	(Veritas Volume Manager) Volume Replicator timeout value (ticks).
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by Volume Replicator (bytes).
vol_stats_enable	(Veritas Volume Manager) Enable VxVM I/O stat collection.
vol_subdisk_num	(Veritas Volume Manager) Maximum number of subdisks attached to a single VxVM plex. This tunable requires a system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions. This tunable requires a system reboot to take effect.
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires a system reboot to take effect.

Table B-1 Supported tunable parameters *(continued)*

Tunable	Description
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (kBytes). This tunable requires a system reboot to take effect.
voldrl_volumemax_drtregs	(Veritas Volume Manager) Max per volume dirty regions in log-plex DRL.
voldrl_volumemax_drtregs_20	(Veritas Volume Manager) Max per volume dirty regions in DCO version 20.
voldrl_dirty_regions	(Veritas Volume Manager) Number of regions cached for DCO version 30.
voliomem_chunk_size	(Veritas Volume Manager) Size of VxVM memory allocation requests (bytes). This tunable requires a system reboot to take effect.
voliomem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires a system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires a system reboot to take effect.
volpagemod_max_memsz	(Veritas Volume Manager) Maximum paging module memory used by Instant Snapshots (Kbytes).

Table B-1 Supported tunable parameters *(continued)*

Tunable	Description
volraid_rsrtransmax	(Veritas Volume Manager) Maximum number of VxVM RAID-5 transient reconstruct operations in parallel. This tunable requires a system reboot to take effect.
vxfs_mbuf	(Veritas File System) Maximum memory used for VxFS buffer cache. This tunable requires a system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires a system reboot to take effect.
write_nstream	(Veritas File System) The number of parallel write requests of size write_pref_io that can be outstanding at one time. The installer can only set the system default value of write_nstream. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
write_pref_io	(Veritas File System) The preferred write request size. The installer can only set the system default value of write_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.

Appendix C

Troubleshooting installation issues

This appendix includes the following topics:

- Restarting the installer after a failed network connection
- About the VRTSspt RPM troubleshooting tools
- Incorrect permissions for root on remote system
- Inaccessible system

Restarting the installer after a failed network connection

If an installation is aborted because of a failed network connection, restarting the installer will detect the previous installation. The installer prompts to resume the installation. If you choose to resume the installation, the installer proceeds from the point where the installation aborted. If you choose not to resume, the installation starts from the beginning.

About the VRTSspt RPM troubleshooting tools

The VRTSspt RPM provides a group of tools for troubleshooting a system and collecting information on its configuration. If you install and use the VRTSspt RPM, 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 RPM, 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 Veritas InfoScale 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 Linux 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.