

Veritas InfoScale™ 7.3.1 Installation Guide - AIX

Last updated: 2017-12-04

Legal Notice

Copyright © 2017 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
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

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

<https://www.veritas.com/support>

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

<https://sort.veritas.com/documents>

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Section 1	Introduction to Veritas InfoScale	8
Chapter 1	Introducing Veritas InfoScale	9
	About the Veritas InfoScale product suite	9
	About Veritas InfoScale Foundation	10
	About Veritas InfoScale Storage	11
	About Veritas InfoScale Availability	11
	About Veritas InfoScale Enterprise	11
	Components of the Veritas InfoScale product suite	11
Chapter 2	Licensing Veritas InfoScale	13
	About Veritas InfoScale product licensing	13
	Registering Veritas InfoScale using product license keys	14
	Registering Veritas InfoScale product using keyless licensing	15
	Updating your product licenses	17
	Using the <code>vxlicinstupgrade</code> utility	17
	About the <code>VRTSvlic</code> package	19
Section 2	Planning and preparation	20
Chapter 3	System requirements	21
	Important release information	21
	Disk space requirements	22
	Hardware requirements	22
	SF and SFHA hardware requirements	23
	SFCFS and SFCFSA hardware requirements	23
	SF Oracle RAC hardware requirements	24
	VCS hardware requirements	25
	Virtual I/O Server (VIOS) requirements	26
	Supported operating systems and database versions	27
	Number of nodes supported	27

Chapter 4	Preparing to install	28
	Mounting the ISO image	28
	Setting up ssh or rsh for inter-system communications	29
	Obtaining installer patches	29
	Disabling external network connection attempts	30
	Verifying the systems before installation	31
	Setting up the private network	31
	Optimizing LLT media speed settings on private NICs	34
	Guidelines for setting the media speed for LLT interconnects	34
	Guidelines for setting the maximum transmission unit (MTU) for LLT interconnects in Flexible Storage Sharing (FSS) environments	35
	Setting up shared storage	35
	Setting the SCSI identifier value	35
	Setting up Fibre Channel	37
	Synchronizing time settings on cluster nodes	37
	Configuring LLT interconnects to use Jumbo Frames	37
	Planning the installation setup for SF Oracle RAC systems	38
	Planning your network configuration	39
	Planning the storage	42
	Planning volume layout	48
	Planning file system design	49
	Setting the umask before installation	49
	Updating the <code>SCSI reserve ODM</code> attribute settings for VIOS	49
Section 3	Installation of Veritas InfoScale	51
Chapter 5	Installing Veritas InfoScale using the installer	52
	Installing Veritas InfoScale using the installer	52
Chapter 6	Installing Veritas InfoScale using response files	55
	About response files	55
	Syntax in the response file	56
	Installing Veritas InfoScale using response files	56
	Response file variables to install Veritas InfoScale	57
	Sample response file for Veritas InfoScale installation	58

Chapter 7	Installing Veritas Infoscale using operating system-specific methods	59
	About installing Veritas InfoScale using operating system-specific methods	59
	Installing Veritas InfoScale using NIM and the installer	59
	Preparing the installation bundle on the NIM server	60
	Installing Veritas InfoScale on the NIM client using SMIT on the NIM server	60
	Installing Veritas InfoScale and the operating system on the NIM client using SMIT	61
Chapter 8	Completing the post installation tasks	63
	Verifying product installation	63
	Setting environment variables	64
	Next steps after installation	64
Section 4	Uninstallation of Veritas InfoScale	66
Chapter 9	Uninstalling Veritas InfoScale using the installer	67
	Preparing to uninstall a Veritas InfoScale product	67
	Moving volumes to physical disks	68
	Removing the Replicated Data Set	70
	Uninstalling Veritas InfoScale filesets using the installer	72
	Removing Storage Foundation products using SMIT	73
	Removing the Storage Foundation for Databases (SFDB) repository	75
Chapter 10	Uninstalling Veritas InfoScale using response files	77
	Uninstalling Veritas InfoScale using response files	77
	Response file variables to uninstall Veritas InfoScale	78
	Sample response file for Veritas InfoScale uninstallation	79
Section 5	Installation reference	80
Appendix A	Installation scripts	81
	Installation script options	81

Appendix B	Tunable files for installation	87
	About setting tunable parameters using the installer or a response file	87
	Setting tunables for an installation, configuration, or upgrade	88
	Setting tunables with no other installer-related operations	89
	Setting tunables with an un-integrated response file	90
	Preparing the tunables file	91
	Setting parameters for the tunables file	91
	Tunables value parameter definitions	92
Appendix C	Troubleshooting installation issues	100
	Restarting the installer after a failed network connection	100
	Troubleshooting an installation on AIX	100
	Incorrect permissions for root on remote system	101
	Resource temporarily unavailable	102
	Inaccessible system	102
Index	104

Introduction to Veritas InfoScale

- [Chapter 1. Introducing Veritas InfoScale](#)
- [Chapter 2. Licensing Veritas InfoScale](#)

Introducing Veritas InfoScale

This chapter includes the following topics:

- [About the Veritas InfoScale product suite](#)
- [About Veritas InfoScale Foundation](#)
- [About Veritas InfoScale Storage](#)
- [About Veritas InfoScale Availability](#)
- [About Veritas InfoScale Enterprise](#)
- [Components of the Veritas InfoScale product suite](#)

About the Veritas InfoScale product suite

The Veritas InfoScale product suite addresses enterprise IT service continuity needs. It draws on Veritas' long heritage of world-class availability and storage management solutions to help IT teams in realizing ever more reliable operations and better protected information across their physical, virtual, and cloud infrastructures. It provides resiliency and software defined storage for critical services across the datacenter infrastructure. It realizes better Return on Investment (ROI) and unlocks high performance by integrating next-generation storage technologies. The solution provides high availability and disaster recovery for complex multi-tiered applications across any distance. Management operations for Veritas InfoScale are enabled through a single, easy-to-use, web-based graphical interface, Veritas InfoScale Operations Manager.

The Veritas InfoScale product suite offers the following products:

- Veritas InfoScale Foundation

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

About Veritas InfoScale Foundation

Veritas InfoScale™ Foundation is specifically designed for enterprise edge-tier, departmental, and test/development systems. InfoScale Foundation combines the industry-leading File System and Volume Manager technology, and delivers a complete solution for heterogeneous online storage management while increasing storage utilization and enhancing storage I/O path availability.

Storage features included in InfoScale Foundation products are listed below:

- Veritas InfoScale Operations Manager Support
- Supports file systems upto 256 TB
- Device names using Array Volume IDs
- Dirty region logging
- Dynamic LUN expansion
- Dynamic Multi-pathing
- Enclosure based naming
- iSCSI device support
- Keyless licensing
- Online file system defragmentation
- Online file system grow & shrink
- Online relayout
- Online volume grow & shrink

Storage features included in InfoScale Storage and Enterprise products, but not included in the InfoScale Foundation product are listed below:

- Hot-relocation
- Remote mirrors for campus clusters
- SCSI-3 based I/O Fencing
- SmartMove
- Split-mirror snapshot

- Thin storage reclamation
- Flexible Storage Sharing

About Veritas InfoScale Storage

Veritas InfoScale™ Storage enables organizations to provision and manage storage independently of hardware types or locations. InfoScale Storage delivers predictable Quality-of-Service by identifying and optimizing critical workloads. InfoScale Storage increases storage agility enabling you to work with and manage multiple types of storage to achieve better ROI without compromising on performance and flexibility.

About Veritas InfoScale Availability

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

About 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. Realize better ROI and unlock high performance by integrating next-generation storage technologies. The solution provides high availability and disaster recovery for complex multi-tiered applications across any distance in physical and virtual environments.

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)

Licensing Veritas InfoScale

This chapter includes the following topics:

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

About Veritas InfoScale product licensing

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

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

- **Install with a license key for the product**
When you purchase a Veritas InfoScale product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased.
See [“Registering Veritas InfoScale using product license keys”](#) on page 14.
- **Install without a license key (keyless licensing)**
Installation without a license does not eliminate the need to obtain a license. The administrator and company representatives must ensure that a server or cluster is entitled to the license level for the products installed. Veritas reserves the right to ensure entitlement and compliance through auditing.

See “[Registering Veritas InfoScale product using keyless licensing](#)” on page 15.

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

www.veritas.com/licensing/process

Registering Veritas InfoScale using product license keys

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

Using the
 installer

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

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

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

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

Enter **1** to register your license key.

See “[Installing Veritas InfoScale using the installer](#)” on page 52.

- You can also register your license keys using the installer menu. Run the following command:

```
./installer
```

Select the **L) License a Product** option in the installer menu.

Manual

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

```
# cd /opt/VRTS/bin  
# ./vxlicinst -k license key  
# vxdctl license init
```

or

```
# vxlicinstupgrade -k
```

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

```
# cd /opt/VRTS/bin  
# ./vxlicinstupgrade -k license key
```

For more information:

See [“Using the vxlicinstupgrade utility”](#) on page 17.

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

Registering Veritas InfoScale product using keyless licensing

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

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

Using the installer

- Run the following command:

```
./installer
```

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

During the installation, you will get the following prompt:

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

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

Enter **2** for keyless licensing.

See [“Installing Veritas InfoScale using the installer”](#) on page 52.

- You can also register your license keys using the installer menu.

Run the following command:

```
./installer
```

Select the **L) License a Product** option in the installer menu.

Manual

Perform the following steps after installation or upgrade:

- 1** Change your current working directory:

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

- 2** View the possible settings for the product level:

```
# vxkeyless displayall
```

- 3** Register the desired product:

```
# vxkeyless set prod_levels
```

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

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

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

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

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

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

Updating your product licenses

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

Move from one product to another

Perform the following steps:

```
# export PATH=$PATH:/opt/VRTSvlic/bin
# vxkeyless set prod_levels
```

Move from keyless licensing to key-based licensing

You will need to remove the keyless licenses by using the NONE keyword.

Note: Clearing the keys disables the Veritas InfoScale products until you install a new key or set a new product level.

```
# vxkeyless [-q] set NONE
```

Register a Veritas InfoScale product using a license key:

See "[Registering Veritas InfoScale using product license keys](#)" on page 14.

Using the `vxlicinstupgrade` utility

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

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

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

- If the current license key is keyless or user-defined and if the user is trying to install the keyless or user defined key of the same product.

Example: If the 7.3.1 Foundation Keyless license key is already installed on a system and the user tries to install another 7.3.1 Foundation Keyless license key, then `vxlicinstupgrade` utility shows an error message:

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

- If the current key is keyless and the newly entered license key is user-defined of the same product

Example: If the 7.3.1 Foundation Keyless license key is already installed on a system and the user tries to install 7.3.1 Foundation user-defined license, then the `vxlicinstupgrade` utility installs the new licenses at `/etc/vx/licenses/lic` and all the 7.3.1 Foundation Keyless keys are deleted and backed up at `/var/vx/licenses/lic<date-timestamp>`.

- If the current key is of higher version and the user tries to install a lower version license key.

Example: If the 7.3.1 Enterprise license key is already installed on a system and the user tries to install the 6.0 SFSTD license key, then the `vxlicinstupgrade` utility shows an error message:

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

- If the current key is of a lower version and the user tries to install a higher version license key.

Example: If 6.0 SFSTD license key is already installed on a system and the user tries to install 7.3.1 Storage license key, then the `vxlicinstupgrade` utility installs the new licenses at `/etc/vx/licenses/lic` and all the 6.0 SFSTD keys are deleted and backed up at `/var/vx/licenses/lic<date-timestamp>`.

Supported Co-existence scenarios:

- InfoScale Foundation and InfoScale Availability
- InfoScale Storage and InfoScale Availability

Example: If the 7.3.1 Foundation or 7.3.1 Storage license key is already installed and the user tries to install 7.3.1 Availability license key or vice -versa, then the `vxlicinstupgrade` utility installs the new licenses and both the keys are preserved at `/etc/vx/licenses/lic`.

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

About the VRTSvlic package

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

<code>vxlicinstupgrade</code>	Installs or upgrades your license key when you have a product or older license already present on the system. See the <code>vxlicinstupgrade(1m)</code> manual page
<code>vxlicrep</code>	Displays the currently installed licenses
<code>vxlictest</code>	Retrieves the features and their descriptions that are encoded in a license key

Planning and preparation

- [Chapter 3. System requirements](#)
- [Chapter 4. Preparing to install](#)

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 disk space requirements for each product.

Table 3-1 Disk space requirements

Product name	Requirement
Veritas InfoScale Foundation	639 MB
Veritas InfoScale Availability	997 MB
Veritas InfoScale Storage	1719 MB
Veritas InfoScale Enterprise	1812 MB

Hardware requirements

This section lists the hardware requirements for Veritas InfoScale.

Table 3-2 lists the hardware requirements for each component in Veritas InfoScale.

Table 3-2 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 23.
Storage Foundation Cluster File System (SFCFS) and Storage Foundation Cluster File System for High Availability (SFCFSHA)	See "SFCFS and SFCFSHA hardware requirements" on page 23.
Storage Foundation for Oracle RAC (SF Oracle RAC)	See "SF Oracle RAC hardware requirements" on page 24.
Cluster Server (VCS)	See "VCS hardware requirements" on page 25.

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

Table 3-3 SF and SFHA hardware requirements

Item	Requirement
Memory	Each system requires at least 1 GB.
For DMP: Virtual I/O Server (VIOS) requirements	2.2.2.1 or later

SFCFS and SFCFSHA hardware requirements

Table 3-4 lists the hardware requirements for SFCFSHA.

Table 3-4 Hardware requirements for SFCFSHA

Requirement	Description
Memory (Operating System)	2 GB of memory.
CPU	A minimum of 2 CPUs.
Node	Storage Foundation Cluster File System High Availability supports mixed cluster environments with AIX 7.1 and 7.2 operating systems.
Shared storage	<p>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 <code>/</code>, <code>/usr</code>, <code>/var</code> and other system partitions on local devices.</p> <p>In a Flexible Storage Sharing (FSS) environment, shared storage may not be required.</p>
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.

Table 3-4 Hardware requirements for SFCFSHA (*continued*)

Requirement	Description
Cluster platforms	<p>There are several hardware platforms that can function as nodes in a Veritas InfoScale cluster.</p> <p>See the <i>Veritas InfoScale 7.3.1 Release Notes</i>.</p> <p>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.</p>
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 hardware requirements

[Table 3-5](#) lists the hardware requirements for basic clusters.

Table 3-5 Hardware requirements for basic clusters

Item	Description
DVD drive	A DVD drive on one of the nodes in the cluster.
Disks	<p>All shared storage disks support SCSI-3 Persistent Reservations (PR).</p> <p>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.</p>
RAM	Each system requires at least 2 GB.
Swap space	For SF Oracle RAC: See the Oracle Metalink document: 169706.1

Table 3-5 Hardware requirements for basic clusters (*continued*)

Item	Description
Network	<p>Two or more private links and one public link.</p> <p>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.</p> <p>Veritas recommends gigabit Ethernet using enterprise-class switches for the private links.</p> <p>Oracle RAC requires that all nodes use the IP addresses from the same subnet.</p> <p>You can also configure aggregated interfaces.</p>
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 3-6](#) lists the hardware requirements for a VCS cluster.

Table 3-6 Hardware requirements for a VCS cluster

Item	Description
DVD drive	One drive in a system that can communicate to all the nodes in the cluster.
Disks	<p>Typical configurations require that the applications are configured to use shared disks/storage to enable migration of applications between systems in the cluster.</p> <p>The SFHA I/O fencing feature requires that all data and coordinator disks support SCSI-3 Persistent Reservations (PR).</p> <p>Note: SFHA also supports non-SCSI3 server-based fencing configuration in virtual environments that do not support SCSI-3 PR-compliant storage.</p>

Table 3-6 Hardware requirements for a VCS cluster (*continued*)

Item	Description
Ethernet controllers	<p>In addition to the built-in public Ethernet controller, VCS requires at least one more Ethernet interface per system. Veritas recommends two additional interfaces.</p> <p>You can also configure aggregated interfaces.</p> <p>Veritas recommends that you turn off the spanning tree on the LLT switches, and set port-fast on.</p>
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.

Virtual I/O Server (VIOS) requirements

To run DMP in VIOS, the minimum VIOS level that is required is 2.2.2.1 or later.

Before installing DMP on VIOS, confirm the following:

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

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

```
# lsattr -El hdisk557 | grep res
reserve_policy single_path
Reserve Policy True
# chdev -l hdisk557 -a reserve_policy=no_reserve -P
hdisk557 changed
```
- If a path has the `reserve_lock` attribute set, change the `reserve_lock` attribute to `no`.

```
# lsattr -El hdisk558 | grep reserve_lock
reserve_lock yes
Reserve Device on open True
# chdev -l hdisk558 -a reserve_lock=no -P
hdisk558 changed
```

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 64 nodes.

SFHA, SFCFSHA, SF Oracle RAC: Flexible Storage Sharing (FSS) only supports cluster configurations with up to 8 nodes.

SFHA, SFCFSHA: SmartIO writeback caching only supports cluster configurations with up to 2 nodes.

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](#)
- [Configuring LLT interconnects to use Jumbo Frames](#)
- [Planning the installation setup for SF Oracle RAC systems](#)
- [Updating the SCSI reserve ODM attribute settings for VIOS](#)

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 Create a loopback device to which you can bind the ISO image file:

```
# mkdev -c loopback -s node -t loopback
loop0 Available
```

- 3 Bind the ISO image to the loopback device and mount the device:

```
# loopmount -i <ISO_image_path> -l loop0 \
-o "-V cdrfs -o ro" -m /mnt
```

Where *<ISO_image_path>* is the complete path to the ISO image

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 30.

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-7.3.1P2-patches.tar.gz
```

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

```
# tar -xvf cpi-7.3.1P2-patches.tar
patches/
patches/CPI731P2.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/CPI731P2.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, filesets, 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 and SF Oracle RAC

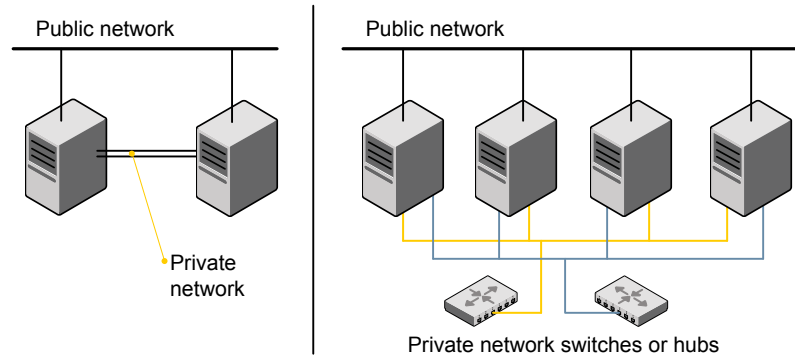
VCS requires you to set up a private network between the systems that form a cluster. You can use either NICs or aggregated interfaces to set up private network.

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.

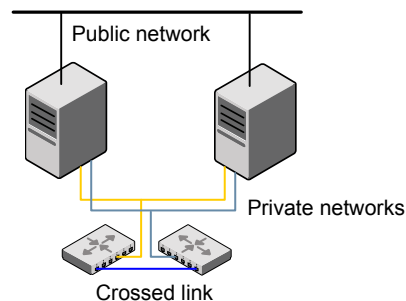
Figure 4-1 Private network setups: two-node and four-node clusters



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 Ethernet controllers 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, SFCFS, SFCFSHA and SF Oracle RAC


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

Setting the SCSI identifier value

SCSI adapters are typically set with a default identifier value of 7. Each device on a SCSI bus must have a unique SCSI identifier value. When more than one system is connected to a SCSI bus, you must change the SCSI identifier to a unique number. You must make this change to one or more systems, usually the unique number is 5 or 6.

Perform the procedure if you want to connect to shared storage with shared SCSI devices.

Figure 4-3 Cabling the shared storage



To set the SCSI identifier value

- 1 Determine the SCSI adapters on each system:

```
north # lsdev -C -c adapter | grep scsi
scsi0   Available 11-08   Wide/Ultra-2 SCSI I/O Controller
scsi1   Available 11-09   Wide/Ultra-2 SCSI I/O Controller
south # lsdev -C -c adapter | grep scsi
scsi0   Available 11-08   Wide/Ultra-2 SCSI I/O Controller
scsi1   Available 11-09   Wide/Ultra-2 SCSI I/O Controller
```

- 2 Verify the SCSI ID of each adapter:

```
north # lsattr -E -l scsi0 -a id
id 7 Adapter card SCSI ID True
north # lsattr -E -l scsi1 -a id
id 7 Adapter card SCSI ID True
south # lsattr -E -l scsi0 -a id
id 7 Adapter card SCSI ID True
south # lsattr -E -l scsi1 -a id
id 7 Adapter card SCSI ID True
```

- 3 If necessary, change the SCSI identifier on each system so that it is unique:

```
south # chdev -P -l scsi0 -a id=5
scsi0 changed
south # chdev -P -l scsi1 -a id=5
scsi1 changed
```

- 4 Shut down all systems in the cluster.

- 5 Cable the shared storage as illustrated in [Figure 4-3](#).
- 6 Restart each system. After all systems have booted, use the `lspv` command to verify that each system can see all shared devices needed by the application.

Setting up Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up Fibre Channel

- 1 Connect the Fibre Channel adapters and the shared storage devices to the same hub or switch.

All systems must see all the shared devices that are required to run the critical application. If you want to implement zoning for a fibre switch, make sure that no zoning prevents all systems from seeing all these shared devices.
- 2 Reboot each system:


```
# shutdown -Fr
```
- 3 After all systems have booted, use the `lspv` command to verify that each system can see all shared devices needed by the application.

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.

Configuring LLT interconnects to use Jumbo Frames

You can configure LLT interconnects to enable Jumbo Frames by increasing the maximum transmission unit (MTU) for physical systems and logical domains.

For physical systems enable Jumbo Frames at interface and LLT-level.

For logical domains enable jumbo Frames for LLT inside the logical domain. You need to ensure that Jumbo Frames are enabled for the virtual network (`vnet`), virtual switch (`vsw`) and the backend physical interface. If a physical switch is used between any cluster nodes to connect the interconnect, ensure that MTU value of switch is also set to a value that matches with other network components.

Perform these steps on all nodes of the cluster

- 1 Enable Jumbo frames at interface level.
- 2 If its a physical machine, run these steps for all the interfaces to be used by LLT

```
# chdev -Pl ifc-name -a jumbo_frames=yes
```

where, *ifc-name* is the interface name.

- 3 Run the command on physical as well as on LPARs.

```
# chdev -Pl ifc-name -a mtu=9000
```

where, *ifc-name* is the interface name.

- 4 Reboot the system

```
# shutdown -Fr
```

- 5 Modify llttab for 9000 MTU.

```
llttab:
```

```
set-node <hostname>
set-cluster <clus-id>
```

```
link ent1 /dev/dlpi/en:1 - ether - 9000
link ent2 /dev/dlpi/en:2 - ether 9000
```

Planning the installation setup for SF Oracle RAC 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 39.
- Planning the storage
See [“Planning the storage”](#) on page 42.
- Planning volume layout
See [“Planning volume layout”](#) on page 48.

- Planning file system design
See [“Planning file system design”](#) on page 49.

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 32 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.

For a11g and later versions, you must use UDP IPC for the database cache fusion traffic.

Veritas recommends using multiple private interconnects for load balancing the 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-miscount 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 41.

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.
For Oracle RAC 11.2.0.1 versions, the PrivNIC and MultiPrivNIC agents provide a reliable alternative when operating system limitations prevent you from using NIC bonding to provide high availability and increased bandwidth using multiple network interfaces. In the event of a NIC failure or link failure, the agent fails over the private IP address from the failed link to the connected or available LLT link. To use multiple links for database cache fusion for increased bandwidth, configure the `cluster_interconnects` initialization parameter with multiple IP addresses for each database instance and configure these IP addresses under MultiPrivNIC for high availability.
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	<p>Use a native NIC bonding solution to provide redundancy, in case of NIC failures.</p> <p>Make sure that a link configured under a aggregated link or NIC bond is not configured as a separate LLT link.</p> <p>When LLT is configured over a bonded interface, do one of the following steps to prevent GAB from reporting jeopardy membership:</p> <ul style="list-style-type: none">■ Configure an additional NIC under LLT in addition to the bonded NIC.■ Add the following line in the <code>/etc/llttab</code> file: <pre>set-dbg-minlinks 2</pre>

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

Options	Description
Using PrivNIC/MultiPrivNIC agents	<p>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 HAIP.</p> <p>Use the PrivNIC agent when operating system limitations prevent you from using NIC bonding to provide high availability using multiple network interfaces.</p> <p>Use the MultiPrivNIC agent when operating system limitations prevent you from using NIC bonding to provide high availability and increased bandwidth using multiple network interfaces.</p> <p>For more deployment scenarios that illustrate the use of PrivNIC/MultiPrivNIC deployments, see the appendix "SF Oracle RAC deployment scenarios" in this document.</p>

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.

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

Files	Type of storage
SF Oracle RAC binaries	Local
SF Oracle RAC 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 44.

- 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 46.

Planning the storage for OCR and voting disk

Depending on the Oracle RAC version and the type of redundancy you want for the OCR and voting disks, use one of the following storage options:

External redundancy	<p>Oracle RAC 11g Release 2 and later versions:</p> <ul style="list-style-type: none"> ■ Clustered File System ■ ASM disk groups created using CVM raw volumes <p>See “OCR and voting disk storage configuration for external redundancy” on page 44.</p>
Normal redundancy	<p>Clustered File System</p> <p>See “OCR and voting disk storage configuration for normal redundancy” on page 45.</p> <p>Note: It is recommended that you configure atleast resource dependency for high availability of the OCR and voting disk resources.</p>

Review the following notes before you proceed:

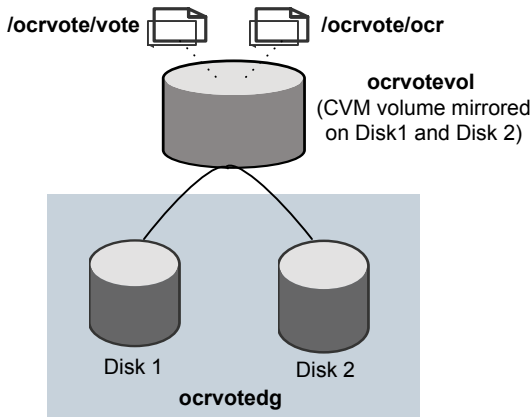
- Set the disk detach policy setting to (`local`) with `ioship` off for OCR and voting disk.
- Configure OCR and voting disk on non-replicated shared storage when you configure global clusters.
- If you plan to use FSS, configure OCR and voting disk on SAN storage.

OCR and voting disk storage configuration for external redundancy

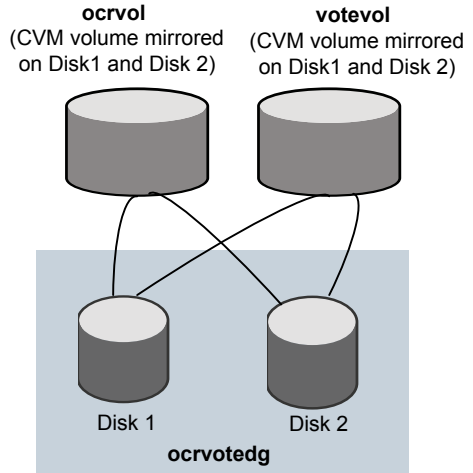
[Figure 4-4](#) illustrates the OCR and voting disk storage options for external redundancy.

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

Option 1: OCR and voting disk on CFS with two-way mirroring



Option 2: OCR and voting disk on CVM raw volume with two-way mirroring



- 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 CVM raw volumes or 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.

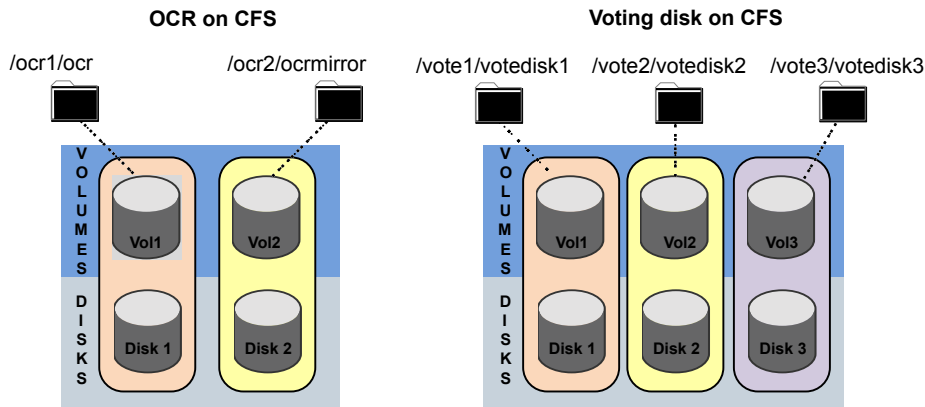
For both option 1 and option 2:

- 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 4-5 illustrates the OCR and voting disk storage options for normal redundancy.

Figure 4-5 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, online redo logs and archive log files, and backup files. Starting with Oracle RAC 11g Release 2, ASM also supports storage for OCR and voting disk.
Not supported by ASM	Oracle RAC 11g Release 2 and later versions: ASM does not support Oracle binaries, trace files, alert logs, export files, tar files, core files, and application binaries on ASM.

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 than 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
```

Updating the SCSI reserve ODM attribute settings for VIOS

This step applies to DMP.

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

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

```
# lsattr -El hdisk557 | grep res
reserve_policy single_path
```

```
Reserve Policy True
```

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

```
hdisk557 changed
```

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

```
# lsattr -El hdisk558 | grep reserve_lock
```

```
reserve_lock yes
```

```
Reserve Device on open True
```

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

```
hdisk558 changed
```

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](#)

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
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 product license keys”](#) on page 14.

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

Installing 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 file 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 filesets. Configuration can be performed at a later time using the <code>-configure</code> 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}{license}	CFG{keys}{vxkeyless} gives the list of keyless keys to be registered on the system. CFG{keys}{license} gives the list of user defined keys to be registered on the system List of Scalar: List Optional or required: Required.
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required

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 filesets that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
CFG{opt}{rsh}	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. List or scalar: scalar Optional or required: optional
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional

Sample response file for Veritas InfoScale installation

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

```
our %CFG;  
  
$CFG{accepteula}=1;  
$CFG{keys}{keyless}=[ qw(ENTERPRISE) ];  
$CFG{opt}{gco}=1;  
$CFG{opt}{install}=1;  
$CFG{prod}="ENTERPRISE731";  
$CFG{systems}=[ qw(system1 system2) ];  
  
1;
```

Installing Veritas Infoscale using operating system-specific methods

This chapter includes the following topics:

- [About installing Veritas InfoScale using operating system-specific methods](#)
- [Installing Veritas InfoScale using NIM and the installer](#)

About installing Veritas InfoScale using operating system-specific methods

On AIX, you can install Veritas InfoScale using the following methods:

- You can use the product installer along with Network Installation Manager (NIM) to install the Veritas InfoScale product, or to install the operating system with the Veritas InfoScale product.

See [“Installing Veritas InfoScale using NIM and the installer”](#) on page 59.

Installing Veritas InfoScale using NIM and the installer

You can use the installer together with the NIM to install the Veritas InfoScale product, or to install the operating system and the Veritas InfoScale product.

The instructions in this section assume a working knowledge of the Network Installation Management process. See the operating system documentation for detailed information on Network Installation Management.

In the following samples, the LPP resource uses LPP-73100-up2date and its relevant SPOT resource is spot-73100-up2date.

Preparing the installation bundle on the NIM server

You need to prepare the installation bundle on the NIM server before you use NIM to install Veritas InfoScale filesets. The following actions are executed on the NIM server.

Note: Make sure that the appropriate NIM LPP_SOURCE and SPOT resources are present on the NIM server.

To prepare the installation bundle

1 Insert and mount the installation media.

2 Choose an LPP source:

```
# lsnim |grep -i lpp_source  
LPP-73100-up2date resources lpp_source
```

3 Navigate to the product directory on the installation media and run the following command to prepare the bundle resource:

```
# ./installer -nim LPP-73100-up2date
```

The installation program copies the necessary filesets and patches to the LPP resource directory.

4 Enter a name for the bundle.

5 Run the `lsnim -l` command to check that the `installp_bundle` resource and `install_scripts` resource are both created successfully.

Installing Veritas InfoScale on the NIM client using SMIT on the NIM server

You can install Veritas InfoScale on the NIM client using the SMIT tool on the NIM server.

Perform these steps on each node to have Veritas InfoScale installed in a cluster.

To install Veritas InfoScale

- 1 On the NIM server, start SMIT.


```
# smitty nim
```
- 2 In the menu, select **Perform NIM Software Installation and Maintenance Tasks**.
- 3 In the menu, select **Install and Update Software**.
- 4 In the menu, select **Install Software Bundle**.
- 5 Select the systems from the list on which to install the software bundle.
- 6 In the menu, select the LPP_SOURCE. In this example, specify **LPP-73100-up2date**.
- 7 In the menu, select the bundle.
- 8 For the `installp` flags, specify that the ACCEPT new license agreements flag has a **yes** value.
- 9 If you want to use NIM to upgrade to another Veritas InfoScale product, make sure that the **Customization SCRIPT to run after installation** flag has the value `[install_scripts]`.
- 10 Press the Enter key to start the installation. Note that it may take some time to finish.
- 11 After the installation completes, configure Veritas InfoScale.

Installing Veritas InfoScale and the operating system on the NIM client using SMIT

You can install Veritas InfoScale and the operating system on the NIM client using the SMIT tool.

Perform these steps on each node to have Veritas InfoScale and AIX installed in a cluster.

To install Veritas InfoScale and the operating system

- 1 On the NIM server, start smitty for a NIM and operating system installation.


```
# smitty nim_bosinst
```
- 2 In the menu, select the standalone target.
- 3 In the menu, select **spot - Install a copy of a SPOT resource**.
- 4 In the menu, select the spot resource **spot-73100-up2date**.

- 5 In the menu, select the LPP_SOURCE. In this example, select **LPP-73100-up2date**.
- 6 In the menu, select the following options:
 - For the ACCEPT new license agreements option, specify **yes**.
- 7 For the `installp` flags, specify that the ACCEPT new license agreements flag has a **yes** value.
- 8 After the installation completes, configure Veritas InfoScale.

Completing the post installation tasks

This chapter includes the following topics:

- [Verifying product installation](#)
- [Setting environment variables](#)
- [Next steps after installation](#)

Verifying product installation

After the Veritas InfoScale products are installed, the packages should be in the COMMITTED state, as indicated by a C in the output:

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

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

To find out about the installed filesets 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 packages, and the status (success or failure) of each package, 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`.

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

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

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

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

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

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

The `nroff` versions of the online manual pages are not readable using the `man` command if the `bos.txt.tfs` fileset is not installed. However, the `VRTSvxvm` and `VRTSvxfs` filesets install ASCII versions in the `/opt/VRTS/man/cat*` and `/opt/VRTS/man/man*` directories that are readable without the `bos.txt.tfs` fileset.

Next steps after installation

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

[Table 8-1](#) lists the components and the respective Configuration and Upgrade guides that are available.

Table 8-1 Guides available for configuration

Component	Document name
Storage Foundation	See <i>Storage Foundation Configuration and Upgrade Guide</i> See <i>Storage Foundation Administrator's Guide</i>
Storage Foundation and High Availability	See <i>Storage Foundation and High Availability Configuration and Upgrade Guide</i>
Storage Foundation Cluster File System HA	See <i>Storage Foundation Cluster File System High Availability Configuration and Upgrade Guide</i> See <i>Storage Foundation Cluster File System High Availability Administrator's Guide</i>
Cluster Server	See <i>Cluster Server Configuration and Upgrade Guide</i> See <i>Cluster Server Administrator's Guide</i>
Storage Foundation for Oracle RAC	See <i>Storage Foundation for Oracle RAC Configuration and Upgrade Guide</i> See <i>Storage Foundation for Oracle RAC Administrator's Guide</i>

Uninstallation of Veritas InfoScale

- [Chapter 9. Uninstalling Veritas InfoScale using the installer](#)
- [Chapter 10. Uninstalling Veritas InfoScale using response files](#)

Uninstalling Veritas InfoScale using the installer

This chapter includes the following topics:

- [Preparing to uninstall a Veritas InfoScale product](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling Veritas InfoScale filesets using the installer](#)
- [Removing Storage Foundation products using SMIT](#)
- [Removing the Storage Foundation for Databases \(SFDB\) repository](#)

Preparing to uninstall a Veritas InfoScale product

Complete the following preparations to uninstall a Veritas InfoScale product.

Warning: Failure to follow the preparations that are outlined in this chapter can result in loss of data.

To remove Veritas InfoScale, complete the following preparations before the uninstallation:

- Back up all VxFS file systems in full and move the files in all VxFS file systems to native file systems backed with LVM logical volumes. Raw application data stored in VxVM logical volumes must be moved to LVM logical volumes.
- Remove all but one copy of file systems and databases.

- Remove all but one plex from volumes that contain multiple plexes (mirrors). To display a list of all volumes, use the command:

```
# vxprint -Ath
```

To remove a plex, use the command:

```
# vxplex -g diskgroup -o rm dis plex
```

- If a remaining plex contains multiple subdisks, consolidate the subdisks into a single subdisk using the commands:

```
# vxassist -g diskgroup mirror volume layout=contig  
# vxplex -g diskgroup -o rm dis plex
```

Sufficient space on another disk is required for this operation to complete.

- Modify `/etc/filesystems` to remove or change entries for VxFS file systems that were moved to native file systems.
- Move all data from volumes created from multiple regions of storage, including striped or spanned volumes, onto a single disk or appropriate LVM logical volume. This can be done using one of the following three methods:
 - Back up the system to tape or other media and recover the system from this.
 - Move volumes incrementally (evacuate) onto logical volumes. Evacuation moves subdisks from the source disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to LVM volumes. See [“Moving volumes to physical disks”](#) on page 68.

Moving volumes to physical disks

You can use the following steps to move data off of VxVM volumes.

To move data off of VxVM volumes

- 1 Evacuate as many disks as possible by using one of the following methods:
 - the "Remove a disk" option in `vxdiskadm`
 - the Veritas Enterprise Administrator

- the `vxevac` script from the command line.

- 2 Remove the evacuated disks from Veritas Volume Manager control using the following commands:

```
# vxdg -g diskgroup rmdisk disk_media_name
# /usr/lib/vxvm/bin/vxdiskunsetup -C disk_access_name
# vxdisk rm disk_access_name
```

For example:

```
# vxdg -g mydg rmdisk mydg01
# /usr/lib/vxvm/bin/vxdiskunsetup -C hdisk1
# vxdisk rm hdisk1
```

- 3 Decide which volume to move first. If the volume to be moved is mounted, unmount it. 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 has been synchronized.
- 4 On the free disk space, create an LVM logical volume that is the same size as the VxVM volume. If there is not enough free space for the logical volume, add a new disk to the system for the first volume to be removed. For subsequent volumes, you can use the free space generated by the removal of the first volume.
- 5 Copy the data on the volume onto the newly created LVM logical volume using the following command:

```
# dd if=/dev/vx/dsk/diskgroup/volume of=/dev/vgvol
```

where *diskgroup* is the name of a VxVM disk group, *volume* is the old volume in that disk group, and *vgvol* is a newly created LVM volume.

If the volume contains a VxFS file system, the user data managed by VxFS in the volume must be backed up or copied to a native AIX file system in an LVM logical volume.

- 6 The entries in `/etc/filesystems` for volumes holding VxFS file systems, that were copied to native file systems in step 5, must be modified according to the change in step 5.
- 7 Mount the disk if the corresponding volume was previously mounted.
- 8 Remove the volume from VxVM using the following command:

```
# vxedit -g diskgroup -rf rm volume
```

- 9 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 -g diskgroup -F "%sdnum" disk_media_name
```

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

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

- 11 Copy the data in the next volume to be removed to the newly created free space.
- 12 Reboot the system after all volumes have been converted successfully. Verify that no open volumes remain after the system reboot using the following command:

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

- 13 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:

```
# vxlink -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 filesets using the installer

Use the following procedure to remove Veritas InfoScale products.

Not all filesets 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 7.3.1 with a previous version of Veritas InfoScale.

To shut down and remove the installed Veritas InfoScale filesets

- 1 Disable DMP native support, if it is enabled. Run the following command to disable DMP native support


```
# vxddmpadm settune dmp_native_support=off  
# shutdown -Fr
```
- 2 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/filesystems`. Failing to remove these entries could result in system boot problems later.
- 3 Unmount all mount points for VxFS file systems.


```
# umount /mount_point
```
- 4 If the VxVM package (`VRTSvxxvm`) is installed, read and follow the uninstallation procedures for VxVM.

See [“Preparing to uninstall a Veritas InfoScale product”](#) on page 67.
- 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 filesets are uninstalled.

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

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

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

```
# installp -u VRTSvxvm
```

Removing Storage Foundation products using SMIT

Use the following procedure to remove Storage Foundation products using SMIT.

To remove the packages using SMIT

- 1 Stop the following SFCFSA modules: VCS, VxFEN, ODM, GAB, and LLT.

Run the following commands to stop the SFCFSA modules:

```
# hastop -all

# /etc/methods/glmkextadm unload

# /etc/rc.d/rc2.d/s99odm stop

# /etc/methods/gmskextadm unload

# /etc/init.d/vxfen.rc stop

# /etc/init.d/gab.rc stop

# /etc/init.d/llt.rc stop
```

Run the following commands to check if all the modules have been stopped:

```
# gabconfig -a

# lltconfig
```

- 2 Disable DMP native support, if it is enabled. Run the following command to disable DMP native support

```
# vxddmpadm settune dmp_native_support=off

# shutdown -Fr
```

- 3 Enter this command to invoke SMIT:

```
# smit
```

- 4 In SMIT, select **Software Installation and Maintenance > Software Maintenance and Utilities > Remove Installed Software**.
- 5 Under the **SOFTWARE name** menu, press F4 or Esc-4 to list all the software that is installed on the system.
- 6 Enter "/" for Find, type "VRTS" to find all filesets, and select the filesets that you want to remove.

- 7 Restart the system after removing all Storage Foundation filesets.
- 8 Depending on the choices that were made when Storage Foundation was originally installed, you may find that not all of the listed Storage Foundation filesets are installed on the system. You may also choose to remove the `VRTSvlic` licensing package unless some other Veritas InfoScale software requires it.

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.

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 filesets that are needed during the install. The default location is <code>/var/tmp</code> . 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 <code>/opt/VRTS/install/logs</code> . 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 filesets. 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}="ENTERPRISE731";
$CFG{systems}=[ qw("system1", "system2") ];

1;
```

Installation reference

- [Appendix A. Installation scripts](#)
- [Appendix B. Tunable files for installation](#)
- [Appendix C. Troubleshooting installation issues](#)

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 filesets required for the specified product. The filesets are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-comcleanup	The <code>-comcleanup</code> 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 <code>-comsetup</code> 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 package upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.
-hostfile <i>full_path_to_file</i>	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
<code>-keyfile <i>ssh_key_file</i></code>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-I <i>ssh_key_file</i></code> to every SSH invocation.
<code>-license</code>	Registers or updates product licenses on the specified systems.
<code>-logpath <i>log_path</i></code>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.
<code>-noipc</code>	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.
<code>-nolic</code>	Allows installation of product filesets without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
<code>-pkgtable</code>	Displays product's packages in correct installation order by group.
<code>-postcheck</code>	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
<code>-precheck</code>	Performs a preinstallation check to determine if systems meet all installation requirements. Veritas recommends doing a precheck before installing a product.
<code>-prod</code>	Specifies the product for operations.
<code>-component</code>	Specifies the component for operations.
<code>-redirect</code>	Displays progress details without showing the progress bar.
<code>-require</code>	Specifies an installer patch file.

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

Command Line Option	Function
<code>-requirements</code>	The <code>-requirements</code> option displays required OS version, required filesets and patches, file system space, and other system requirements in order to install the product.
<code>-responsefile <i>response_file</i></code>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> 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.
<code>-rolling_upgrade</code>	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.
<code>-rollingupgrade_phase1</code>	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel filesets get upgraded to the latest version.
<code>-rollingupgrade_phase2</code>	The <code>-rollingupgrade_phase2</code> option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent filesets upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.
<code>-rsh</code>	Specify this option when you want to use rsh and RCP for communication between systems instead of the default ssh and SCP.
<code>-serial</code>	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 <code>-tunablesfile</code> option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-timeout	The <code>-timeout</code> option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the <code>-timeout</code> option overrides the default value of 1200 seconds. Setting the <code>-timeout</code> option to 0 prevents the script from timing out. The <code>-timeout</code> option does not work with the <code>-serial</code> option.
-tmppath <i>tmp_path</i>	Specifies a directory other than <code>/var/tmp</code> as the working directory for the installation scripts. This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
-tunables	Lists all supported tunables and create a tunables file template.
-tunablesfile <i>tunables_file</i>	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 packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing filesets and patches where applicable. Lists the installed patches, patches, and available updates for the installed product if an Internet connection is available.

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 88.

- 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 89.

- 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 90.

See [“About response files”](#) on page 55.

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

See [“Tunables value parameter definitions”](#) on page 92.

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 92.

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 91.
- 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 92.

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 91.
- 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 92.

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 91.
- 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 92.

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 (<i>vxesd</i>) 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 tuneftab(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 tuneftab(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 tuneftab(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 tuneftab(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_alllds	(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_checkpoint_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.
vx_bc_bufhwm	(Veritas File System) VxFS metadata buffer cache high water mark. 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 tuneftab(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 tuneftab(4) manual page for setting this tunable for a specified block device.

Troubleshooting installation issues

This appendix includes the following topics:

- [Restarting the installer after a failed network connection](#)
- [Troubleshooting an installation on AIX](#)
- [Incorrect permissions for root on remote system](#)
- [Resource temporarily unavailable](#)
- [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.

Troubleshooting an installation on AIX

Save a copy of `/var/adm/ras/errtmpl` and `/etc/trcfmt` files before you install the product. If the filesets fail to install due to the template file is corrupted error message, replace `/var/adm/ras/errtmpl` file and `/etc/trcfmt` file with the ones that you had saved, uninstall all the filesets installed.

See [“Preparing to uninstall a Veritas InfoScale product”](#) on page 67.

Then reinstall.

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.

Resource temporarily unavailable

If the installation fails with the following error message on the console:

```
fork() failed: Resource temporarily unavailable
```

The value of maximum number of processes allowed per user may not be large enough. This kernel attribute is a tunable and can be changed on any node of the cluster.

To determine the current value of "Maximum number of PROCESSES allowed per user", enter:

```
# lsattr -H -E -l sys0 -a maxuproc
```

To see the default value of this tunable and its valid range of values, enter:

```
# odmget -q "attribute=maxuproc" PdAt
```

If necessary, you can change the value of the tunable using the smitty interface:

```
# smitty chgsys
```

You can also directly change the CuAt class using the following command:

```
# chdev -l sys0 -a maxuproc=600
```

Increasing the value of the parameter takes effect immediately; otherwise the change takes effect after a reboot.

See the `smitty` and `chdev` manual pages.

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 AIX system names separated by spaces: q,? (host1)
```

Suggested solution: Verify that you entered the system name correctly; use the `ping(1M)` command to verify the accessibility of the host.

Index

A

- about
 - installation using operating system-specific methods 59
 - response files 55
 - Veritas InfoScale 9
 - Veritas InfoScale Availability 11
 - Veritas InfoScale Enterprise 11
 - Veritas InfoScale Foundation 10
 - Veritas InfoScale product licensing 13
 - Veritas InfoScale Storage 11
 - VRTSvlic package 19
 - vxlicinstupgrade utility 17

C

- cabling shared devices 37
- components
 - Veritas InfoScale 11
- configuring
 - hardware 25
 - private network 31
 - rsh 29
 - ssh 29
 - switches 31
- controllers
 - private Ethernet 31

D

- deinstalling the Volume Manager 67
- disabling
 - external network connection attempts 30
- disk space
 - directories 25
 - required 25

E

- eprom
 - parameters 31
- Ethernet controllers 31

F

- fibre channel 25

H

- hardware
 - configuring network and storage 25
- Hardware requirements
 - Veritas InfoScale 22
- hubs 31

I

- installation
 - next steps 64
 - response file variables 57
 - sample response file 58
 - Veritas InfoScale 52
- installation script options 81
- installer patches
 - obtaining either manually or automatically 29
- installing
 - on NIM client using SMIT on NIM server 60
 - operating system on the NIM client using SMIT 61
 - required disk space 25
 - using NIM 59
 - using response files 56
 - Veritas InfoScale using operating system-specific methods 59
- ISO image
 - mounting 28

J

- Jumbo Frames
 - LLT 37

K

- keyless licensing
 - Veritas InfoScale 15

L

licensing
 registering Veritas InfoScale product license
 keys 14

LLT

interconnects 34–35
 Jumbo Frames 37

M

MAC addresses 31
 media speed 34
 optimizing 34
 MTU 35

N

network switches 31
 NIM
 installing 59
 preparing the installation bundle 60

O

obtaining
 installer patches either automatically or
 manually 29
 optimizing
 media speed 34
 Oracle RAC
 pre-installation
 setting up storage 43

P

parameters
 eeprom 31
 persistent reservations
 SCSI-3 35
 preparing to uninstall
 Veritas InfoScale 67
 private network
 configuring 31

R

RAM
 installation requirement 25
 release information 21
 removing
 the Replicated Data Set 70
 removing Storage Foundation products using SMIT 73

Replicated Data Set
 removing the 70
 requirements
 Ethernet controllers 25
 fibre channel 25
 hardware 25
 RAM Ethernet controllers 25
 SCSI host bus adapter 25
 response file variables
 installation 57
 uninstall 78
 response files
 about 55
 installation 56
 syntax 56
 uninstalling 77
 rsh
 configuration 29

S

sample response file
 installation 58
 uninstall 79
 SCSI
 changing initiator IDs 35
 SCSI host bus adapter 25
 SCSI ID
 changing 36
 verifying 36
 SCSI-3
 persistent reservations 35
 setting
 environment variables 64
 setting umask, before installing 49
 setup
 cabling shared devices 37
 SCSI Initiator ID 35
 Shared storage
 Fibre Channel 35
 shared storage
 setting SCSI initiator ID 35
 ssh
 configuration 29
 supported operating systems 27
 switches 31
 synchronizing time settings, before installing 37

T

- tunables file
 - about setting parameters 87
 - parameter definitions 92
 - preparing 91
 - setting for configuration 88
 - setting for installation 88
 - setting for upgrade 88
 - setting parameters 91
 - setting with no other operations 89
 - setting with un-integrated response file 90

U

- uninstall
 - moving volumes to physical disks 68
 - response file variables 78
 - sample response file 79
 - using the installer 72
- uninstalling
 - Storage Foundation products using SMIT 73
 - using response files 77
- updating licenses
 - Veritas InfoScale 17

V

- verifying
 - product installation 63
- Veritas InfoScale
 - about 9
 - components 11
 - Hardware requirements 22
 - keyless licensing 15
 - mounting ISO image 28
 - preparing to uninstall 67
 - product installer 52
 - registering Veritas InfoScale product license
 - keys 14
 - updating licenses 17
- Veritas InfoScale Availability
 - about 11
- Veritas InfoScale Enterprise
 - about 11
- Veritas InfoScale Foundation
 - about 10
- Veritas InfoScale installation
 - pre-installation tasks
 - setting umask 49
 - synchronizing time settings 37

Veritas InfoScale installation *(continued)*

- pre-installation tasks *(continued)*
 - verifying systems 31
- requirements
 - hardware 24

Veritas InfoScale Storage

- about 11

VIOS requirements 26

vradmin

- delpri 71
- stoprep 71