

Veritas InfoScale™ 7.4 Installation Guide - Solaris

Last updated: 2019-02-26

Legal Notice

Copyright © 2019 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:

infoscaledocs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Section 1	Planning and preparation	8
Chapter 1	Introducing Veritas InfoScale	9
	About the Veritas InfoScale product suite	9
	Components of the Veritas InfoScale product suite	9
	About the co-existence of Veritas InfoScale products	11
Chapter 2	Licensing Veritas InfoScale	12
	About Veritas InfoScale product licensing	12
	Licensing notes	13
	Registering Veritas InfoScale using permanent license key file	15
	Registering Veritas InfoScale using keyless license	16
	Managing InfoScale licenses	18
	About the <code>vxlicinstupgrade</code> utility	20
	Generating license report with <code>vxlicrep</code> command	21
Chapter 3	System requirements	22
	Important release information	22
	Disk space requirements	23
	Hardware requirements	23
	SF and SFHA hardware requirements	24
	SFCFS and SFCFSHA hardware requirements	24
	SF Oracle RAC and SF Sybase CE hardware requirements	25
	VCS hardware requirements	26
	Supported operating systems and database versions	27
	Veritas File System requirements	27
	Number of nodes supported	28
Chapter 4	Preparing to install	29
	Mounting the ISO image	29
	Setting up ssh or rsh for inter-system communications	30
	Obtaining installer patches	30
	Disabling external network connection attempts	31

Verifying the systems before installation	32
Setting up the private network	32
Optimizing LLT media speed settings on private NICs	35
Guidelines for setting the media speed for LLT interconnects	36
Guidelines for setting the maximum transmission unit (MTU) for LLT interconnects in Flexible Storage Sharing (FSS) environments	36
Setting up shared storage	37
Setting up shared storage: SCSI disks	37
Setting up shared storage: Fibre Channel	39
Synchronizing time settings on cluster nodes	40
Creating a root user	40
Configuring LLT interconnects to use Jumbo Frames	41
Creating the /opt directory	43
Planning the installation setup for SF Oracle RAC and SF Sybase CE systems	43
Planning your network configuration	43
Planning the storage	47
Planning volume layout	53
Planning file system design	54
Setting the umask before installation	54
Making the IPS publisher accessible	54
Preparing zone environments	56
Section 2	Installation of Veritas InfoScale
	57
Chapter 5	Installing Veritas InfoScale using the installer
	58
	Installing Veritas InfoScale using the installer
	58
Chapter 6	Installing Veritas InfoScale using response files
	61
	About response files
	61
	Syntax in the response file
	62
	Installing Veritas InfoScale using response files
	62
	Response file variables to install Veritas InfoScale
	63
	Sample response files for Veritas InfoScale installation
	64

Chapter 7	Installing Veritas Infoscale using operating system-specific methods	66
	About installing Veritas InfoScale using operating system-specific methods	66
	Installing Veritas InfoScale on Solaris 11 using Automated Installer	67
	About Automated Installation	67
	Using Automated Installer	67
	Using AI to install the Solaris 11 operating system and Veritas InfoScale products	68
	Manually installing packages on Solaris 11 systems	72
	Manually installing packages on Oracle Solaris 11 systems	72
	Manually installing packages on Solaris non-global zones	73
Chapter 8	Completing the post installation tasks	75
	Verifying product installation	75
	Setting environment variables	76
	Disabling the abort sequence on SPARC systems	76
	Next steps after installation	77
Section 3	Uninstallation of Veritas InfoScale	79
Chapter 9	Uninstalling Veritas InfoScale using the installer	80
	About removing Veritas InfoScale	80
	Preparing to uninstall	81
	Removing the Replicated Data Set	90
	Uninstalling Veritas InfoScale packages using the installer	91
	Uninstalling Veritas InfoScale using the <code>pkgrm</code> or <code>pkg uninstall</code> command	93
	Manually uninstalling Veritas InfoScale packages on non-global zones on Solaris 11	94
	Removing the Storage Foundation for Databases (SFDB) repository	94
Chapter 10	Uninstalling Veritas InfoScale using response files	96
	Uninstalling Veritas InfoScale using response files	96
	Response file variables to uninstall Veritas InfoScale	97
	Sample response file for Veritas InfoScale uninstallation	98

Section 4	Installation reference	99
Appendix A	Installation scripts	100
	Installation script options	100
Appendix B	Tunable files for installation	106
	About setting tunable parameters using the installer or a response file	106
	Setting tunables for an installation, configuration, or upgrade	107
	Setting tunables with no other installer-related operations	108
	Setting tunables with an un-integrated response file	109
	Preparing the tunables file	110
	Setting parameters for the tunables file	110
	Tunables value parameter definitions	111
Appendix C	Troubleshooting installation issues	119
	Restarting the installer after a failed network connection	119
	About the VRTSspt package troubleshooting tools	119
	Incorrect permissions for root on remote system	120
	Inaccessible system	121

Planning and preparation

- [Chapter 1. Introducing Veritas InfoScale](#)
- [Chapter 2. Licensing Veritas InfoScale](#)
- [Chapter 3. System requirements](#)
- [Chapter 4. Preparing to install](#)

Introducing Veritas InfoScale

This chapter includes the following topics:

- [About the Veritas InfoScale product suite](#)
- [Components of the Veritas InfoScale product suite](#)
- [About the co-existence of Veritas InfoScale products](#)

About the Veritas InfoScale product suite

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

The Veritas InfoScale product suite offers the following products:

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

Components of the Veritas InfoScale product suite

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

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

Table 1-1 Veritas InfoScale product suite

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

About the co-existence of Veritas InfoScale products

You can install an InfoScale product on a system where another InfoScale product is already installed.

The following table provides the supported co-existence scenarios.

Table 1-2 InfoScale products co-existence

Product installed	Supported co-existence			
	InfoScale Foundation	InfoScale Availability	InfoScale Storage	InfoScale Enterprise
InfoScale Foundation	Not applicable	Supported	Not supported	Not supported
InfoScale Availability	Supported	Not applicable	Supported	Not supported
InfoScale Storage	Not supported	Supported	Not applicable	Not supported
InfoScale Enterprise	Not supported	Not supported	Not supported	Not applicable

Licensing Veritas InfoScale

This chapter includes the following topics:

- [About Veritas InfoScale product licensing](#)
- [Licensing notes](#)
- [Registering Veritas InfoScale using permanent license key file](#)
- [Registering Veritas InfoScale using keyless license](#)
- [Managing InfoScale licenses](#)
- [Generating license report with vxlicrep command](#)

About Veritas InfoScale product licensing

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

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

- Install product with a permanent license
When you purchase a Veritas InfoScale product, you receive a License Key certificate. The certificate specifies the products and the number of product licenses purchased.
See [“Registering Veritas InfoScale using permanent license key file”](#) on page 15.
- Install product without a permanent license key (keyless licensing)
Installation without a license does not eliminate the need to obtain a license. The administrator and company representatives must ensure that a server or cluster is entitled to the license level for the products installed. Veritas reserves the right to ensure entitlement and compliance through auditing.

See “[Registering Veritas InfoScale using keyless license](#)” on page 16.

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

www.veritas.com/licensing/process

Licensing notes

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

- If you use a keyless license option, you must configure Veritas InfoScale Operations Manager within two months of product installation and add the node as a managed host to the Veritas InfoScale Operations Manager Management Server. Failing this, a warning message for non-compliance is displayed periodically.
For more details, refer to Veritas InfoScale Operations Manager product documentation.
- Note the following limitation in case of InfoScale Availability and InfoScale Storage co-existence:
If Keyless licensing type is selected during the product installation, checks performed to monitor the number of days of product installation are based on the InfoScale Storage component. As a result, if you do not enter a valid license key file or do not add the host as a managed host within 60 days of InfoScale Storage installation, a non-compliance error is logged every 4 hrs in the Event Viewer.
- The text-based license keys that are used in previous product versions are not supported when upgrading to version 7.4. If your current product is installed using a permanent license key and you do not have a permanent license key file for the newer InfoScale version, you can temporarily upgrade using the keyless licensing. Then you must procure a permanent license key file from the Veritas license certificate and portal within 60 days, and upgrade using the permanent license key file to continue using the product.
- The license key file must be present on the same node where you are trying to install the product.

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

- You can manage the license keys using the `vxlicinstupgrade` utility.

See [“Managing InfoScale licenses”](#) on page 18.

- Before upgrading the product, review the licensing details and back up the older license key. If the upgrade fails for some reason, you can temporarily revert to the older product using the older license key to avoid any application downtime.
- You can use the license assigned for higher Stock Keeping Units (SKU) to install the lower SKUs.

For example, if you have procured a license that is assigned for InfoScale Enterprise, you can use the license for installing any of the following products:

- InfoScale Foundation
- InfoScale Storage
- InfoScale Availability

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

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

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

Registering Veritas InfoScale using permanent license key file

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

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

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

Using the
installer

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

- Run the following command:

```
./installer
```

- During the installation, the following interactive message appears:

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

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

- Enter **1** to register the license key.
- Then provide the absolute path of the .slf license key file saved on the current node.

Example:

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

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

- Run the following command:

```
./installer
```

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

To install InfoScale using the installer:

See [“Installing Veritas InfoScale using the installer”](#) on page 58.

Manual

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

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

or

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

then,

```
# vxdctl license init
```

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

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

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

For more information:

See [“Managing InfoScale licenses”](#) on page 18.

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

Registering Veritas InfoScale using keyless license

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

Using the `installer`

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

- Run the following command:

```
./installer
```

- During the installation, the following interactive message appears:

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

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

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

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

- Run the following command:

```
./installer
```

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

To install InfoScale using the installer:

See [“Installing Veritas InfoScale using the installer”](#) on page 58.

Manual

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

- 1 Change your current working directory:

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

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

```
# vxkeyless displayall
```

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

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

Example:

```
# vxkeyless set ENTERPRISE
```

For more information:

See “Managing InfoScale licenses” on page 18.

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

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

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

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

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

Managing InfoScale licenses

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

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

Using the
`vxlicinstupgrade`

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

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

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

Example:

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

For more information on `vxlicinstupgrade` utility:

See [“About the `vxlicinstupgrade` utility”](#) on page 20.

For more information on permanent licensing:

See [“Registering Veritas InfoScale using permanent license key file”](#) on page 15.

Using the `vxkeyless`

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

1 Change your current working directory:

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

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

```
# vxkeyless displayall
```

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

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

Example:

```
# vxkeyless set ENTERPRISE
```

For more information on keyless licensing:

See [“Registering Veritas InfoScale using keyless license”](#) on page 16.

About the `vxlicinstupgrade` utility

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

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

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

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

Example: If the 7.4 Foundation Keyless license key is already installed on a system and the user tries to install another 7.4 Foundation Keyless license key, then `vxlicinstupgrade` utility shows an error message:

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

- If the current key is keyless and the newly entered license key file is a permanent license of the same product

Example: If the 7.4 Foundation Keyless license key is already installed on a system and the user tries to install 7.4 Foundation permanent license key file, then the `vxlicinstupgrade` utility installs the new license at `/etc/vx/licenses/lic` and the 7.4 Foundation Keyless key is deleted.

- The `vxlicinstupgrade` utility in Veritas InfoScale does not support managing the text-based license keys used in versions before 7.4.
- If the current key is of a lower version and the user tries to install a higher version license key.

Example: If 7.0 Storage license key is already installed on a system and the user tries to install 7.4 Storage license key file, then the `vxlicinstupgrade` utility installs the new license at `/etc/vx/licenses/lic` and the 7.0 Storage key is deleted.

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

Generating license report with vxlicrep command

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

To display a license report:

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

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

System requirements

This chapter includes the following topics:

- [Important release information](#)
- [Disk space requirements](#)
- [Hardware requirements](#)
- [Supported operating systems and database versions](#)
- [Veritas File System requirements](#)
- [Number of nodes supported](#)

Important release information

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

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

- For important updates regarding this release, review the Late-Breaking News TechNote on the Veritas Technical Support website:
https://www.veritas.com/support/en_US/article.000126340
- For the latest patches available for this release, go to:
<https://sort.veritas.com>
- The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware, visit the following URL:
https://www.veritas.com/support/en_US/article.000126344
- The software compatibility list summarizes each Veritas InfoScale product stack and the product features, operating system versions, and third-party products

it supports. For the latest information on supported software, visit the following URL:

https://www.veritas.com/support/en_US/article.000126342

Disk space requirements

Table 3-1 lists the minimum disk space requirements for each product.

Table 3-1 Disk space requirements

Product name	Solaris 11 SPARC (MB)
Veritas InfoScale Foundation	574
Veritas InfoScale Availability	780
Veritas InfoScale Storage	1270
Veritas InfoScale Enterprise	1364

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 24.
Storage Foundation Cluster File System (SFCFS) and Storage Foundation Cluster File System for High Availability (SFCFSHA)	See “ SFCFS and SFCFSHA hardware requirements ” on page 24.

Table 3-2 Hardware requirements for components in Veritas InfoScale
(continued)

Component	Requirement
Storage Foundation for Oracle RAC (SF Oracle RAC)	See “ SF Oracle RAC and SF Sybase CE hardware requirements ” on page 25.
Storage Foundation for Sybase CE (SF Sybase CE)	
Cluster Server (VCS)	See “ VCS hardware requirements ” on page 26.

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.

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 Solaris 11 SPARC operating systems as long as all the nodes in the cluster have the same CPU architecture.

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

Requirement	Description
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	<p>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.</p>
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 7.4 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	<p>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.</p>

SF Oracle RAC and SF Sybase CE 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>

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

Item	Description
RAM	Each system requires at least 2 GB.
Swap space	For SF Oracle RAC: See the Oracle Metalink document: 169706.1
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>

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 network interfaces for private interconnects.</p> <p>You can also configure aggregated interfaces.</p> <p>Veritas recommends that you turn off the spanning tree algorithm on the switches used to connect private network interfaces..</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.

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*.

Veritas File System requirements

Veritas File System requires that the values of the Solaris variables `lwp_default_stksize` and `svc_default_stksize` are at least 0x8000 (for Solaris 11). When you install the Veritas File System package, `VRTSvxfs`, the `VRTSvxfs` packaging scripts check the values of these variables in the kernel. If the values are less than the required values, `VRTSvxfs` increases the values and modifies the `/etc/system` file with the required values. If the `VRTSvxfs` scripts increase the values, the installation proceeds as usual except that you must reboot and restart the installation program. A message displays if a reboot is required.

To avoid an unexpected need for a reboot, verify the values of the variables before installing Veritas File System. Use the following commands to check the values of the variables:

```
For Solaris 11: # echo "lwp_default_stksize/X" | mdb -k
lwp_default_stksize:
lwp_default_stksize:      8000

# echo "svc_default_stksize/X" | mdb -k
svc_default_stksize:
svc_default_stksize:      8000
```

If the values shown are less than 8000 (for Solaris 11), you can expect a reboot after installation.

Note: The default value of the `svc_default_stksize` variable is 0 (zero), which indicates that the value is set to the value of the `lwp_default_stksize` variable. In this case, no reboot is required, unless the value of the `lwp_default_stksize` variable is too small.

To avoid a reboot after installation, you can modify the `/etc/system` file with the appropriate values. Reboot the system prior to installing the packages. Add the following lines to the `/etc/system` file:

```
For Solaris 11: set lwp_default_stksize=0x8000
set rpcmod:svc_default_stksize=0x8000
```

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](#)
- [Creating a root user](#)
- [Configuring LLT interconnects to use Jumbo Frames](#)
- [Creating the /opt directory](#)
- [Planning the installation setup for SF Oracle RAC and SF Sybase CE systems](#)
- [Making the IPS publisher accessible](#)
- [Preparing zone environments](#)

Mounting the ISO image

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

To mount the ISO image

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

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

Where:

<ISO_image_path> is the complete path to the ISO image

<block_device> is the complete path to the block device

- 3 Mount the image:

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

Setting up ssh or rsh for inter-system communications

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

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

Obtaining installer patches

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

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

To download installer patches automatically

- ◆ If you are running 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 31.

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-7.4P2-patches.tar.gz
```

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

```
# tar -xvf cpi-7.4P2-patches.tar
patches/
patches/CPI74P2.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/CPI74P2.pl
```

Disabling external network connection attempts

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

To disable external network connection attempts

- ◆ Disable inter-process communication (IPC).

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

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

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

Verifying the systems before installation

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

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

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

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

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

where *sys1*, *sys2* are the names of the nodes in the cluster.

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

Setting up the private network

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

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

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

For example, consider the following case where:

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

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

The following products make extensive use of the private cluster interconnects for distributed locking:


- Storage Foundation Cluster File System (SFCFS)
- Storage Foundation for Oracle RAC (SF Oracle RAC)

Veritas recommends network switches for the SFCFS and the SF Oracle RAC clusters due to their performance characteristics.

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

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

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

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

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

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

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

You can also manually configure LLT.

Optimizing LLT media speed settings on private NICs

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

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

Guidelines for setting the media speed for LLT interconnects

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

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

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

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

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

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

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

- For virtual NICs, all the components—the virtual NIC, the corresponding physical NIC, and the virtual switch—must be set to 9000 MTU.
- If a higher MTU cannot be configured on the public link (because of restrictions on other components such as a public switch), do not configure the public link in LLT. LLT uses the lowest of the MTU that is configured among all high priority and low priority links.

Setting up shared storage

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

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

Setting up shared storage: SCSI disks

When SCSI devices are used for shared storage, the SCSI address or SCSI initiator ID of each node must be unique. Since each node typically has the default SCSI address of "7," the addresses of one or more nodes must be changed to avoid a conflict. In the following example, two nodes share SCSI devices. The SCSI address of one node is changed to "5" by using `nvedit` commands to edit the `nvrामrc` script.

If you have more than two systems that share the SCSI bus, do the following:

- Use the same procedure to set up shared storage.
- Make sure to meet the following requirements:
 - The storage devices have power before any of the systems
 - Only one node runs at one time until each node's address is set to a unique value

To set up shared storage

- 1 Install the required SCSI host adapters on each node that connects to the storage, and make cable connections to the storage.

Refer to the documentation that is shipped with the host adapters, the storage, and the systems.

- 2 With both nodes powered off, power on the storage devices.
- 3 Power on one system, but do not allow it to boot. If necessary, halt the system so that you can use the ok prompt.

Note that only one system must run at a time to avoid address conflicts.

- 4 Find the paths to the host adapters:

```
{0} ok show-disks
...b) /sbus@6,0/QLGC,isp@2,10000/sd
```

The example output shows the path to one host adapter. You must include the path information without the "/sd" directory, in the `nvrामrc` script. The path information varies from system to system.

- 5 Edit the `nvrामrc` script on to change the `scsi-initiator-id` to 5. (The *Solaris OpenBoot 3.x Command Reference Manual* contains a full list of `nvedit` commands and keystrokes.) For example:

```
{0} ok nvedit
```

As you edit the script, note the following points:

- Each line is numbered, 0:, 1:, 2:, and so on, as you enter the `nvedit` commands.
- On the line where the `scsi-initiator-id` is set, insert exactly one space after the first quotation mark and before `scsi-initiator-id`.

In this example, edit the `nvrामrc` script as follows:

```
0: probe-all
1: cd /sbus@6,0/QLGC,isp@2,10000
2: 5 " scsi-initiator-id" integer-property
3: device-end
4: install-console
5: banner
6: <CTRL-C>
```

- 6 Store the changes you make to the `nvrामrc` script. The changes you make are temporary until you store them.

```
{0} ok nvstore
```

If you are not sure of the changes you made, you can re-edit the script without risk before you store it. You can display the contents of the `nvrामrc` script by entering:

```
{0} ok printenv nvrामrc
```

You can re-edit the file to make corrections:

```
{0} ok nvedit
```

Or, discard the changes if necessary by entering:

```
{0} ok nvquit
```

- 7 Instruct the OpenBoot PROM Monitor to use the `nvrामrc` script on the node.

```
{0} ok setenv use-nvrामrc? true
```

- 8 Reboot the node. If necessary, halt the system so that you can use the ok prompt.
- 9 Verify that the scsi-initiator-id has changed. Go to the ok prompt. Use the output of the show-disks command to find the paths for the host adapters. Then, display the properties for the paths. For example:

```
{0} ok show-disks
...b) /sbus@6,0/QLGC,isp@2,10000/sd
{0} ok cd /sbus@6,0/QLGC,isp@2,10000
{0} ok .properties
scsi-initiator-id      00000005
```

Permit the system to continue booting.

- 10 Boot the second node. If necessary, halt the system to use the ok prompt. Verify that the scsi-initiator-id is 7. Use the output of the show-disks command to find the paths for the host adapters. Then, display the properties for that paths. For example:

```
{0} ok show-disks
...b) /sbus@6,0/QLGC,isp@2,10000/sd
{0} ok cd /sbus@6,0/QLGC,isp@2,10000
{0} ok .properties
scsi-initiator-id      00000007
```

Permit the system to continue booting.

Setting up shared storage: Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up shared storage

- 1 Install the required FC-AL controllers.
- 2 Connect the FC-AL controllers 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.

- 3 Boot each system with the reconfigure devices option:

```
ok boot -r
```

- 4 After all systems have booted, use the `format (1m)` command to verify that each system can see all shared devices.

If Volume Manager is used, the same number of external disk devices must appear, but device names (`c##d##s#`) may differ.

If Volume Manager is not used, then you must meet the following requirements:

- The same number of external disk devices must appear.
- The device names must be identical for all devices on all systems.

Synchronizing time settings on cluster nodes

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

For instructions, see the operating system documentation.

Creating a root user

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

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

To change root role into a user

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

```
% su - root
```

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

```
# roles admin
```

```
root
```

```
# usermod -R " " admin
```


3 Change the root role into a user.

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

4 Verify the change.

```
■ # getent user_attr root
```

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

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

```
■ # userattr type root
```

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

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

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

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.

Physical systems

- 1 Enable Jumbo frames at interface level.
- 2 Set MTU value to 9000 or more.

- 3** If IP is configured on a datalink, you will have to remove the IP before setting the MTU.

- 4** Solaris 11: # `dladm set-linkprop -p mtu=9000 <datalink>`

If LLT is configured over UDP, set MTU for the IP interface

```
# ipadm set-ifprop -p mtu=4000 <interface>
```

- 5** Verify that MTU is set to desired value for datalink.

```
# dladm show-link <datalink>
```

- 6** Verify that MTU is set to desired value for the IP interface.

```
# ipadm show-ifprop -p mtu <interface>
```

- 7** Enable Jumbo Frames at LLT-level.

Logical Domains

- 1** Set MTU at virtual switch level.

Set MTU to 9000 or more while creating a virtual switch.

```
# ldm add-vsw mtu=9000 <vswitch-name> <io-ldom>
```

If virtual switch is already created, set MTU to the desired value.

```
# ldm set-vsw mtu=9000 <vswitch-name>
```

If vnets are already bound to this switch, you will have to reboot the node.

- 2** Stop the logical domain.

- 3** In some cases, you may have to set the mtu at virtual network layer.

```
# ldm set-vnet mtu=9000 <VIRTUAL-NETWORK-INTERFACE> <ldom>
```

- 4** Verify that MTU is set to desired value for the datalink at control domain.

```
# dladm show-link <DATA LINK>
```

- 5** Verify that mtu is set to desired value for the virtual switch.

```
# ldm list-services <io-ldom>
```

- 6** Verify that MTU is set at desired value for the virtual network interface.

```
# ldm list-bindings <ldom>
```

- 7 Verify that MTU is set at datalink inside the logical domain.

```
# dladm show-link <datalink>
```

- 8 Enable Jumbo Frames at LLT level.

For configuring LLT link with MTU size = 9000, change the "link" line in llttab as follows:

```
Solaris 11: link net1 /dev/net/net1 - ether - 9000
```

Creating the /opt directory

This topic applies to DMP, SFHA, and SFCFSA.

The directory `/opt` must exist, be writable, and must not be a symbolic link.

If you want to upgrade, you cannot have a symbolic link from `/opt` to an unconverted volume. If you have a symbolic link to an unconverted volume, the symbolic link does not function during the upgrade and items in `/opt` are not installed.

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

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

Review the following planning guidelines before you install Veritas InfoScale:

- Planning your network configuration
See [“Planning your network configuration”](#) on page 43.
- Planning the storage
See [“Planning the storage”](#) on page 47.
- Planning volume layout
See [“Planning volume layout”](#) on page 53.
- Planning file system design
See [“Planning file system design”](#) on page 54.

Planning your network configuration

The following practices are recommended for a resilient network setup:

- Configure the private cluster interconnect over multiple dedicated gigabit Ethernet links. All single point of failures such as network interface cards (NIC), switches, and interconnects should be eliminated.
- The NICs used for the private cluster interconnect should have the same characteristics regarding speed, MTU, and full duplex on all nodes. Do not allow the NICs and switch ports to auto-negotiate speed.
- Configure non-routable IP addresses for private cluster interconnects.
- The default value for LLT peer inactivity timeout is 16 seconds.

For SF Oracle RAC: The value should be set based on service availability requirements and the propagation delay between the cluster nodes in case of campus cluster setup. The LLT peer inactivity timeout value indicates the interval after which Veritas InfoScale on one node declares the other node in the cluster dead, if there is no network communication (heartbeat) from that node.

The default value for the CSS miss-count in case of Veritas InfoScale is 600 seconds. The value of this parameter is much higher than the LLT peer inactivity timeout so that the two clusterwares, VCS and Oracle Clusterware, do not interfere with each other's decisions on which nodes should remain in the cluster in the event of network split-brain. Veritas I/O fencing is allowed to decide on the surviving nodes first, followed by Oracle Clusterware. The CSS miss-count value indicates the amount of time Oracle Clusterware waits before evicting another node from the cluster, when it fails to respond across the interconnect.

For more information, see the Oracle Metalink document: 782148.1

Planning the public network configuration for Oracle RAC

Identify separate public virtual IP addresses for each node in the cluster. Oracle RAC requires one public virtual IP address for the Oracle RAC listener process on each node. Public virtual IP addresses are used by client applications to connect to the Oracle RAC database and help mitigate TCP/IP timeout delays.

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

Planning the private network configuration for Oracle RAC

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

For Oracle 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 46.

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 IPMP for Oracle Clusterware	If Oracle Clusterware interconnects are configured over IPMP, all the NICs that are configured under LLT must be configured under the IPMP group. In such a configuration, it is recommended not to manage these links using the PrivNIC/MultiPrivNIC agents.
Using link aggregation/ NIC bonding for Oracle Clusterware	<div>Use a native NIC bonding solution to provide redundancy, in case of NIC failures.</div> <div>Make sure that a link configured under a aggregated link or NIC bond is not configured as a separate LLT link.</div> <div>When LLT is configured over a bonded interface, do one of the following steps to prevent GAB from reporting jeopardy membership:</div> <div><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:<div>set-dbg-minlinks 2</div></div>

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 and SF Sybase CE.

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

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

Planning the storage for Oracle RAC

Review the storage options and guidelines for Oracle RAC:

- Storage options for OCR and voting disk

See [“Planning the storage for OCR and voting disk”](#) on page 49.

- 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 51.

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 49.</p>
Normal redundancy	<p>Clustered File System</p> <p>See “OCR and voting disk storage configuration for normal redundancy” on page 50.</p> <p>Note: It is recommended that you configure atleast resource dependency for high availability of the OCR and voting disk resources.</p>

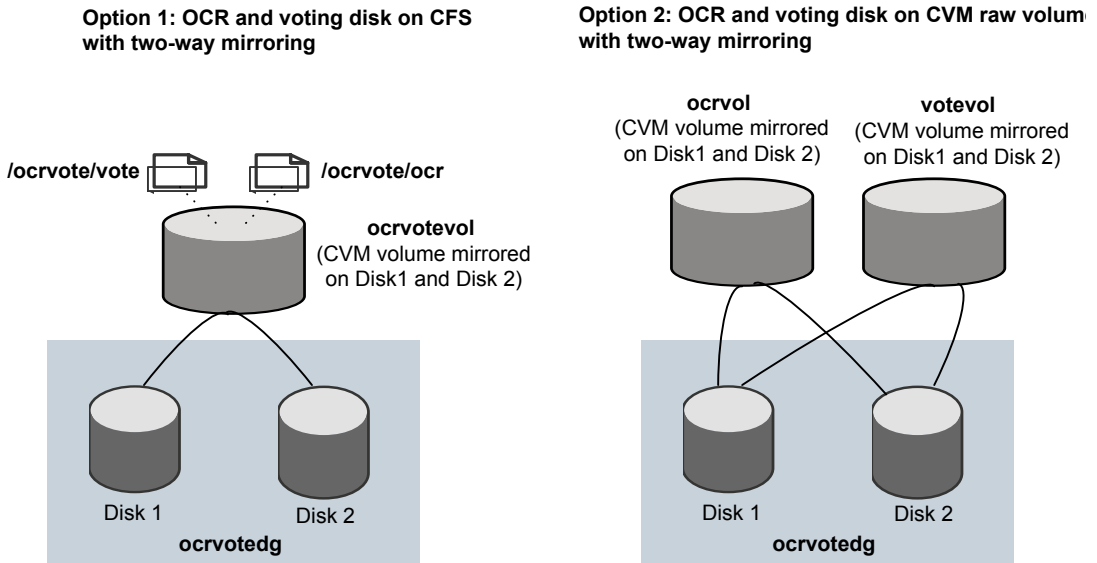
Review the following notes before you proceed:

- 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 4-3](#) illustrates the OCR and voting disk storage options for external redundancy.

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



- If you want to place OCR and voting disk on a clustered file system (option 1), you need to have two separate files for OCR and voting information respectively on CFS mounted on a CVM mirrored volume.
- If you want to place OCR and voting disk on 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 4-4 illustrates the OCR and voting disk storage options for normal redundancy.

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



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

Configure the storage as follows:

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

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

Planning the storage for Oracle RAC binaries and data files

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

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

Consider the following points while planning the installation:

- Local installations provide improved protection against a single point of failure and also allows for applying Oracle RAC patches in a rolling fashion.

- CFS installations provide a single Oracle installation to manage, regardless of the number of nodes. This scenario offers a reduction in storage requirements and easy addition of nodes.

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

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

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

Planning for Oracle RAC ASM over CVM

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

Supported by ASM	ASM provides storage for data files, control files, 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
```

Making the IPS publisher accessible

The topic applies to SFHA, SFCFSHA, SF Oracle RAC, and VCS.

The installation of Veritas InfoScale 7.4 fails on Solaris 11 if the Image Packaging System (IPS) publisher is inaccessible. The following error message is displayed:

CPI ERROR V-9-20-1273 Unable to contact configured publishers on <node_name>.

Solaris 11 introduces the new Image Packaging System (IPS) and sets a default publisher (solaris) during Solaris installation. When additional packages are being installed, the set publisher must be accessible for the installation to succeed. If the publisher is inaccessible, as in the case of a private network, then package

installation will fail. The following commands can be used to display the set publishers:

```
# pkg publisher
```

Example:

```
root@sol11-03:~# pkg publisher
PUBLISHER      TYPE      STATUS    URI
solaris        origin    online    http://pkg.oracle.com/solaris/release/
root@sol11-03:~# pkg publisher solaris
Publisher: solaris
      Alias:
      Origin URI: http://pkg.oracle.com/solaris/release/
      SSL Key: None
      SSL Cert: None
      Client UUID: 00000000-3f24-fe2e-0000-000068120608
      Catalog Updated: October 09:53:00 PM
      Enabled: Yes
      Signature Policy: verify
```

To make the IPS publisher accessible

- 1 Enter the following to disable the publisher (in this case, solaris):

```
# pkg set-publisher --disable solaris
```

- 2 Repeat the installation of Veritas InfoScale 7.4.

- 3 Re-enable the original publisher. If the publisher is still inaccessible (private network), then the `no-refresh` option can be used to re-enable it.

```
# pkg set-publisher --enable solaris
```

or

```
# pkg set-publisher --enable --no-refresh solaris
```

Note: Unsetting the publisher will have a similar effect, except that the publisher can only be re-set if it is accessible. See `pkg(1)` for further information on the `pkg` utility.

Preparing zone environments

The topic applies to SF, SFHA, SFCFSHA, SF Oracle RAC, and VCS.

You need to keep the following items in mind when you install or upgrade VCS using the installer in a zone environment on an Oracle Solaris 11 operating system.

- All the zones are upgraded (both global and non-global) unless they are detached and unmounted.
- After installing packages in the global zone, you have to install the required packages in the non-global zone for Oracle Solaris 11 if you are performing installation for the first time. On Oracle Solaris 11.1, if the non-global zone has an older version of Veritas InfoScale packages already installed then during the upgrade of the Veritas InfoScale packages in global zone, packages inside non-global zone are automatically upgraded provided the zone is running.
- Make sure that all non-global zones are booted and in the running state before you install or upgrade the VCS packages in the global zone. If the non-global zones are not running at the time of upgrade, you have to set the Veritas publisher and then attach the zone with the `-u` option to install or upgrade the VCS packages inside the non-global zone.

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

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

See [“Registering Veritas InfoScale using permanent license key file”](#) on page 15.

- 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 files for Veritas InfoScale installation](#)

About response files

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

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

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

Syntax in the response file

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

For example, in the case of a string value:

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

or, in the case of an integer value:

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

or, in the case of a list:

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

Installing Veritas InfoScale using response files

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

To install Veritas InfoScale using response files

- 1 Make sure the systems where you want to install Veritas InfoScale meet the installation requirements.
- 2 Make sure that the preinstallation tasks are completed.
- 3 Copy the response file to the system where you want to install Veritas InfoScale.
- 4 Edit the values of the response file variables as necessary.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installation from the system to which you copied the response file.
For example:

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

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

- 7 Complete the Veritas InfoScale post-installation tasks.

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

Response file variables to install Veritas InfoScale

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

Table 6-1 Response file variables for installing Veritas InfoScale

Variable	Description
CFG{opt}{install}	Installs Veritas InfoScale packages. 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}{licensefile}	CFG{keys}{vxkeyless} gives the keyless key to be registered on the system. CFG{keys}{licensefile} gives the absolute file path to the permanent license key to be registered on the system. List of Scalar: List Optional or required: Required.
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required

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

Variable	Description
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional
CFG{opt}{tmppath}	Defines the location where a working directory is created to store temporary files and the packages that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
CFG{opt}{rsh}	Defines that <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 files for Veritas InfoScale installation

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

```
our %CFG;

$CFG{accepteula}=1;
$CFG{keys}{keyless}=[ qw(ENTERPRISE) ];
$CFG{opt}{gco}=1;
$CFG{opt}{install}=1;
$CFG{prod}="ENTERPRISE74";
$CFG{systems}=[ qw(system1 system2) ];

1;
```


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

```
our %CFG;

$CFG{accepteula}=1;
$CFG{keys}{licensefile}=["<path_to_license_key_file>"];
$CFG{opt}{gco}=1;
$CFG{opt}{install}=1;
$CFG{prod}="ENTERPRISE74";
$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 on Solaris 11 using Automated Installer](#)
- [Manually installing packages on Solaris 11 systems](#)

About installing Veritas InfoScale using operating system-specific methods

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

- You can use the Oracle Solaris Automated Installer (AI) to install the Solaris 11 operating system and Storage Foundation product on multiple client systems in a network.
See [“Installing Veritas InfoScale on Solaris 11 using Automated Installer”](#) on page 67.
- You can also manually install Veritas InfoScale.
See [“Manually installing packages on Solaris 11 systems”](#) on page 72.

Installing Veritas InfoScale on Solaris 11 using Automated Installer

You can use the Oracle Solaris Automated Installer (AI) to install the Solaris 11 operating system and Veritas InfoScale product on multiple client systems in a network. AI performs a hands-free installation (automated installation without manual interactions) of SPARC systems. You can also use AI media to install the Oracle Solaris OS on a single SPARC platform. Oracle provides the AI bootable image and it can be downloaded from the Oracle website. All cases require access to a package repository on the network to complete the installation.

About Automated Installation

AI automates the installation of the Oracle Solaris 11 OS on one or more SPARC clients in a network. Automated Installation applies to Solaris 11 only. You can install the Oracle Solaris OS on many different types of clients. The clients can differ in:

- Architecture
- Memory characteristics
- MAC address
- IP address
- CPU

The installations can differ depending on specifications including network configuration and packages installed.

An automated installation of a client in a local network consists of the following high-level steps:

- 1 A client system boots and gets IP information from the DHCP server
- 2 Characteristics of the client determine which AI service and which installation instructions are used to install the client.
- 3 The installer uses the AI service instructions to pull the correct packages from the package repositories and install the Oracle Solaris OS on the client.

Using Automated Installer

To use Automated Installer to install systems over the network, set up DHCP and set up an AI service on an AI server. The DHCP server and AI server can be the same system or two different systems.

Make sure that the systems can access an Oracle Solaris Image Packaging System (IPS) package repository. The IPS package repository can reside on the AI server, on another server on the local network, or on the Internet.

An AI service is associated with a SPARC AI install image and one or more sets of installation instructions. The installation instructions specify one or more IPS package repositories from where the system retrieves the packages that are needed to complete the installation. The installation instructions also include the names of additional packages to install and information such as target device and partition information. You can also specify instructions for post-installation configuration of the system.

Consider the operating systems and packages you want to install on the systems. Depending on your configuration and needs, you may want to do one of the following:

- If two systems have different architectures or need to be installed with different versions of the Oracle Solaris OS, create two AI services. Then, associate each AI service with a different AI image
- If two systems need to be installed with the same version of the Oracle Solaris OS but need to be installed differently in other ways, create two sets of installation instructions for the AI service. The different installation instructions can specify different packages to install or a different slice as the install target.

The installation begins when you boot the system. DHCP directs the system to the AI install server, and the system accesses the install service and the installation instructions within that service.

For more information, see the *Oracle® Solaris 11 Express Automated Installer Guide*.

Using AI to install the Solaris 11 operating system and Veritas InfoScale products

Use the following procedure to install the Solaris 11 operating system and Veritas InfoScale products using AI.

To use AI to install the Solaris 11 operating system and Veritas InfoScale products

- 1** Follow the Oracle documentation to set up a Solaris AI server and DHCP server.

You can find the documentation at <http://docs.oracle.com>.

- 2** Set up the Veritas InfoScale package repository.

Run the following commands to startup necessary SMF services and create directories:

```
# svcadm enable svc:/network/dns/multicast:default
# mkdir /ai
# zfs create -o compression=on -o mountpoint=/ai rpool/ai
```

- 3 Run the following commands to set up IPS repository for Veritas SPARC packages:

```
# mkdir -p /ai/repo_vrts_sparc
# pkgrepo create /ai/repo_vrts_sparc
# pkgrepo add-publisher -s /ai/repo_vrts_sparc Veritas
# pkgrecv -s <media_sparc>/pkgs/VRTSpkgs.p5p -d
/ai/repo_vrts_sparc '*'
# svccfg -s pkg/server list
# svcs -a | grep pkg/server
# svccfg -s pkg/server add vrtssparc
# svccfg -s pkg/server:vrtssparc addpg pkg application
# svccfg -s pkg/server:vrtssparc setprop pkg/port=10003
# svccfg -s pkg/server:vrtssparc setprop pkg/inst_root=
/ai/repo_vrts_sparc
# svccfg -s pkg/server:vrtssparc addpg general framework
# svccfg -s pkg/server:vrtssparc addpropvalue general/complete
astring: vrtssparc
# svccfg -s pkg/server:vrtssparc addpropvalue general/enable
boolean: true
# svcs -a | grep pkg/server
# svcadm refresh application/pkg/server:vrtssparc
# svcadm enable application/pkg/server:vrtssparc
```

Or run the following commands to set up the private depot server for testing purposes:

```
# /usr/lib/pkg.depotd -d /ai/repo_vrts_sparc -p 10003 > /dev/null &
```

Check the following URL on IE or Firefox browser:

<http://<host>:10003>

4 Set up the install service on the AI server.

Run the following command:

```
# mkdir /ai/iso
```

Download the AI image from the Oracle website and place the `iso` in the `/ai/iso` directory.

Create an install service.

For example:

To set up the AI install service for SPARC platform::

```
# # installadm create-service -n sol11sparc -s \  
/ai/iso/sol-11-1111-ai-sparc.iso -d /ai/aiboot/
```

5 Run the installer to generate manifest XML files for all the Veritas InfoScale products that you plan to install.

```
# mkdir /ai/manifests  
# <media>/installer -ai /ai/manifests
```

6 For each system, generate the system configuration and include the host name, user accounts, and IP addresses. For example, enter one of the following:

```
# mkdir /ai/profiles  
# sysconfig create-profile -o /ai/profiles/profile_client.xml
```

or

```
# cp /ai/aiboot/auto-install/sc_profiles/sc_sample.xml  
/ai/profiles/profile_client.xml
```

7 Add a system and match it to the specified product manifest and system configuration.

Run the following command to add a SPARC system, for example:

```
# installadm create-client -e "<client_MAC>" -n sol11sparc  
# installadm add-manifest -n sol11sparc -f \  
/ai/manifests/tmp/aitmp/vrts_manifest_enterprise.xml  
# installadm create-profile -n sol11sparc -f \  
/ai/profiles/profile_client.xml -p profile_sc  
# installadm set-criteria -n sol11sparc -m \  
vrts_sfha -p profile_sc -c mac="<client_MAC>"  
# installadm list -m -c -p -n sol11sparc
```

- 8 Run the following command to restart the system and install the operating system and Storage Foundation products:

```
# boot net:dhcp - install
```

If the Solaris operating system version is 11.1 or later, DMP is enabled for the ZFS root device.

For more information about ZFS root support, see *Dynamic Multi-Pathing Administrator's Guide*.

- 9 When the system is up and running, run the installer command from the installation media to configure the Veritas InfoScale software.

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

Note: If you do not find the installer script, execute the `/opt/VRTSsfcp/bin/run-once` command.

Manually installing packages on Solaris 11 systems

The following sections describe how to install packages manually on Solaris 11 systems.

Manually installing packages on Oracle Solaris 11 systems

To install packages on Solaris 11 system

- 1 Copy the `VRTSpkgs.p5p` package from the `pkgs` directory from the installation media to the the system at `/tmp/install` directory..
- 2 Disable the publishers that are not reachable, as package install may fail if any of the already added repositories are unreachable.

```
# pkg set-publisher --disable <publisher name>
```

- 3 Add a file-based repository in the non-global zone.

```
# pkg set-publisher -p /tmp/install/VRTSpkgs.p5p Veritas
```

- 4 Install the required packages.

- 5 To configure an OracleVMServer logical domain for disaster recovery, install the following required packages inside the logical domain:

```
# pkg install --accept VRTSvcsnr
```

- 6 Remove the publisher on the non-global zone.

```
# pkg unset-publisher Veitas
```

- 7 Clear the state of the SMF service if non-global zones are present in the system. In presence of non-global zones, setting the file-based repository causes SMF service `svc:/application/pkg/system-repository:default` to go into maintenance state. .

```
# svcadm clear svc:/application/pkg/system-repository:default
```

- 8 Enable the publishers that were disabled earlier.

```
# pkg set-publisher --enable <publisher>
```

Manually installing packages on Solaris non-global zones

With Oracle Solaris 11, you must manually install Veritas InfoScale packages inside non-global zones.

To install packages manually on Solaris non-global zones

- 1 Ensure that the SMF service `svc:/application/pkg/system-repository:default` and `svc:/application/pkg/zones-proxyd:default` are online on the global zone.

```
global# svcs svc:/application/pkg/system-repository:default
```

```
global# svcs svc:/application/pkg/zones-proxyd:default
```

- 2 Log on to the non-global zone as a super user.
- 3 Ensure that the SMF service `svc:/application/pkg/zones-proxy-client:default` is online inside non-global zone

```
non-global# svcs svc:/application/pkg/zones-proxy-client:default
```

- 4 Copy the `VRTSpkgs.p5p` package from the `pkgs` directory from the installation media to the global zone (for example at `/tmp/install` directory).
- 5 Disable the publishers that are not reachable, as package install may fail if any of the already added repositories are unreachable.

```
global# pkg set-publisher --disable <publisher name>
```

- 6 Add a file-based repository in the global zone.

```
global# pkg set-publisher -g /tmp/install/VRTSpkgs.p5p Veritas
```

- 7 Log on to the non-global zone as a super user and install the required packages.

```
non-global# pkg install --accept <list of packages>
```

where *<list of packages>* is the names of the packages that are required to be installed.

- 8 Remove the publisher on the global zone.

```
global# pkg unset-publisher Veritas
```

- 9 Enable the publishers that were disabled earlier.

```
global# pkg set-publisher --enable <publisher>
```

Completing the post installation tasks

This chapter includes the following topics:

- [Verifying product installation](#)
- [Setting environment variables](#)
- [Disabling the abort sequence on SPARC systems](#)
- [Next steps after installation](#)

Verifying product installation

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

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

To find out about the installed packages 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
```

Disabling the abort sequence on SPARC systems

The topic applies to VCS.

Most UNIX operating systems provide a method to perform a "break" or "console abort." The inherent problem when you abort a hung system is that it ceases to heartbeat in the cluster. When other cluster members believe that the aborted node is a failed node, these cluster members may begin corrective action.

Keep the following points in mind:

- The only action that you must perform following a system abort is to reset the system to achieve the following:
 - Preserve data integrity
 - Prevent the cluster from taking additional corrective actions

- Do not resume the processor as cluster membership may have changed and failover actions may already be in progress.
- To remove this potential problem on SPARC systems, you should alias the `go` function in the OpenBoot eeprom to display a message.

To alias the `go` function to display a message

- 1 At the ok prompt, enter:

```
nvedit
```

- 2 Press Ctrl+L to display the current contents of the nvramrc buffer.
- 3 Press Ctrl+N until the editor displays the last line of the buffer.
- 4 Add the following lines exactly as shown. Press Enter after adding each line.

```
." Aliasing the OpenBoot 'go' command! "  
: go ." It is inadvisable to use the 'go' command in a clustered  
environment. " cr  
." Please use the 'power-off' or 'reset-all' commands instead. "  
cr  
." Thank you, from your friendly neighborhood sysadmin. " ;
```

- 5 Press Ctrl+C to exit the nvramrc editor.
- 6 To verify that no errors exist, type the `nvrn` command. You should see only the following text:

```
Aliasing the OpenBoot 'go' command!
```

- 7 Type the `nvstore` command to commit your changes to the non-volatile RAM (NVRAM) for use in subsequent reboots.
- 8 After you perform these commands, at reboot you see this output:

```
Aliasing the OpenBoot 'go' command! go isn't unique.
```

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	<p>See <i>Storage Foundation Configuration and Upgrade Guide</i></p> <p>See <i>Storage Foundation Administrator's Guide</i></p>
Storage Foundation and High Availability	See <i>Storage Foundation and High Availability Configuration and Upgrade Guide</i>
Storage Foundation Cluster File System HA	<p>See <i>Storage Foundation Cluster File System High Availability Configuration and Upgrade Guide</i></p> <p>See <i>Storage Foundation Cluster File System High Availability Administrator's Guide</i></p>
Cluster Server	<p>See <i>Cluster Server Configuration and Upgrade Guide</i></p> <p>See <i>Cluster Server Administrator's Guide</i></p>
Storage Foundation for Oracle RAC	<p>See <i>Storage Foundation for Oracle RAC Configuration and Upgrade Guide</i></p> <p>See <i>Storage Foundation for Oracle RAC Administrator's Guide</i></p>
Storage Foundation for Sybase SE	<p>See <i>Storage Foundation for Sybase ASE CE Configuration and Upgrade Guide</i></p> <p>See <i>Storage Foundation for Sybase ASE CE Administrator's Guide</i></p>

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:

- [About removing Veritas InfoScale](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling Veritas InfoScale packages using the installer](#)
- [Uninstalling Veritas InfoScale using the pkgm or pkg uninstall command](#)
- [Manually uninstalling Veritas InfoScale packages on non-global zones on Solaris 11](#)
- [Removing the Storage Foundation for Databases \(SFDB\) repository](#)

About removing Veritas InfoScale

This section covers uninstallation requirements and steps to uninstall the Veritas software.

Only users with superuser privileges can uninstall Veritas InfoScale.

Warning: Failure to follow the instructions in the following sections may result in unexpected behavior.

Preparing to uninstall

Review the following before removing the Veritas software.

Remote uninstallation

You must configure remote communication to uninstall Veritas InfoScale on remote systems. In a High Availability environment, you must meet the prerequisites to uninstall on all nodes in the cluster at one time.

The following prerequisites are required for remote uninstallation:

- Communication protocols must exist between systems. By default, the uninstall scripts use ssh.
- You must be able to execute ssh or rsh commands as superuser on all systems.
- The ssh or rsh must be configured to operate without requests for passwords or pass phrases

Preparing to remove Veritas Volume Manager

This section describes the steps you need to take before you remove Veritas Volume Manager (VxVM) to preserve the contents of the volumes.

Warning: Failure to follow the preparations in this section might result in unexpected behavior.

On Solaris 11, the SMF service `vxvm-configure` must be online in order to uninstall VRTSvxvm successfully.

To verify that the vxvm-configure service is online

- 1 Check the state of the `vxvm-configure` service:

```
# svcs -a | grep vxvm-configure
```

- 2 If the service is in disabled or maintenance state, use the following command to display information including the service log location:

```
# svcs -xv vxvm-configure
```

- 3 If there are no issues, use the following command to bring the vxvm-configure service online:

```
# svcadm enable vxvm-configure
```

Moving volumes from an encapsulated root disk

Use the following procedure to move volumes from an encapsulated root disk.

To uninstall VxVM if `root`, `swap`, `usr`, or `var` is a volume under Volume Manager control

- 1 Ensure that the `rootvol`, `swapvol`, `usr`, and `var` volumes have only one associated plex each.

The plex must be contiguous, non-striped, non-spanned, and non-sparse. To obtain this information, enter the following:

```
# vxprint -ht rootvol swapvol usr var
```

If any of these volumes have more than one associated plex, remove the unnecessary plexes using the following command:

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

- 2 Run the `vxunroot` command:

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

The `vxunroot` command changes the volume entries in `/etc/vfstab` to the underlying disk partitions for `rootvol`, `swapvol`, `usr`, and `var`. It also modifies `/etc/system` and prompts for a restart so that disk partitions are mounted instead of volumes for `root`, `swap`, `usr`, and `var`.

- 3 Once you have changed the `root`, `swap`, `usr`, and `var` volumes, move all remaining volumes to disk partitions.

You can do this using one of the following procedures:

- Back up the entire system to tape and then recover from tape.
- Back up each file system individually and then recover them all after you create new file systems on disk partitions.
- Move volumes incrementally to disk partitions.
See [“Moving volumes to disk partitions”](#) on page 82.
Otherwise, shut down VxVM.

Moving volumes to disk partitions

Use the following procedure to move volumes incrementally to disk partitions.

To move volumes incrementally to disk partitions

- 1 Evacuate disks using the `vxdiskadm` command, the VOM GUI, or the `vxevac` utility.

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

- 2 Remove the evacuated disks from VxVM control by entering:

```
# vxdg rmdisk diskname
# vxdisk rm devname
```

- 3 Decide which volume to move first, and if the volume is mounted, unmount it.
- 4 If the volume is used as a raw partition for database applications, make sure that the application does not update the volume. Also make sure that you have applied the `sync` command to the data on the volume.
- 5 Create a partition on free disk space of the same size as the volume using the `format` command.

If there is not enough free space for the partition, add a new disk to the system for the first volume removed. Subsequent volumes can use the free space that the removal of this first volume generates.

- 6 Copy the data on the volume onto the newly created disk partition using a command such as `dd`.

```
# dd if=/dev/vx/dsk/diskgroup/lhome of=/dev/dsk/c2t2d2s7
```

where `c2t2d2` is the disk outside of Volume Manager and `s7` is the newly created partition.

- 7 Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly created partition.
- 8 Mount the disk partition if the corresponding volume was previously mounted.
- 9 Stop and remove the volume from VxVM using the commands.

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

- 10** Remove any free disks (those disks that have no subdisks defined on them) by removing the volumes from VxVM control.

To check if there are still some subdisks remaining on a particular disk, use the `vxprint` command.

```
# vxprint -g diskgroup -F '%sdnum' diskname
```

If the output is not 0, there are still some subdisks on this disk that you need to remove. If the output is 0, remove the disk from VxVM control.

```
# vxdg rmdisk diskname
# vxdisk rm devname
```

Use the free space that is created for adding the data from the next volume you want to remove.

- 11** After you successfully convert all volumes into disk partitions, restart the system.
- 12** After the restart, make sure that none of the volumes are open by using the `vxprint` command.

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

- 13** If any volumes remain open, repeat the steps.

Example of moving volumes to disk partitions on Solaris

This example shows how to move the data on a volume to a disk partition. In the example, there are three disks: `disk1` and `disk2` are subdisks on volume `vol01` and `disk3` is a free disk. The data on `vol01` is copied to `disk3` using the `vxevac` command.

These are the contents of the disk group `voldg` before the data on `vol01` is copied to `disk3`.

```
# vxprint -g voldg -ht
DG NAME  NCONFIG  NLOG     MINORS   GROUP-ID
DM NAME  DEVICE   TYPE     PRIVLEN  PUBLLEN  STATE
RV NAME  RLINK_CNT KSTATE   STATE    PRIMARY  DATAVOL  SRL
RL NAME  RVG      KSTATE   STATE    REM_HOST REM_DG    REM_RLNK
V  NAME  RVG      KSTATE   STATE    LENGTH   READPOL   PREFPLEX  UTYPE
PL NAME  VOLUME   KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID  MODE
SD NAME  PLEX     DISK     DISKOFFS LENGTH    [COL/]OFF DEVICE    MODE
SV NAME  PLEX     VOLNAME  NVOLLAYR LENGTH    [COL/]OFF AM/NM     MODE
DC NAME  PARENTVOL LOGVOL
SP NAME  SNAPVOL  DCO
```

```
dg voldg default      default 115000
1017856044.1141.hostname.veritas.com
```

```
dm disk1 c1t12d0s2 sliced 2591      17900352 -
dm disk2 c1t14d0s2 sliced 2591      17899056 -
dm disk3 c1t3d0s2  sliced 2591      17899056 -
```

```
v  vol1 -            ENABLED ACTIVE  4196448 ROUND      -          fsgen
pl pl1  vol1         ENABLED ACTIVE  4196448 CONCAT    -          RW
sd sd1   pl1         disk1    0       2098224  0          c1t12d0  ENA
sd sd2   pl1         disk2    0       2098224  2098224   c1t14d0  ENA
```

Evacuate disk1 to disk3.

```
# /etc/vx/bin/vxevac -g voldg disk1 disk3
# vxprint -g voldg -ht
```

DG NAME	NCONFIG	NLOG	MINORS	GROUP-ID				
DM NAME	DEVICE	TYPE	PRIVLEN	PUBLEN	STATE			
RV NAME	RLINK_CNT	KSTATE	STATE	PRIMARY	DATAVOLS	SRL		
RL NAME	RVG	KSTATE	STATE	REM_HOST	REM_DG	REM_RLNK		
V NAME	RVG	KSTATE	STATE	LENGTH	READPOL	PREFPLEX	UTYPE	
PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE	
SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE	
SV NAME	PLEX	VOLNAME	NVOLLAYR	LENGTH	[COL/]OFF	AM/NM	MODE	
DC NAME	PARENTVOL	LOGVOL						
SP NAME	SNAPVOL	DCO						

```
dg voldg default      default 115000
1017856044.1141.hostname.veritas.com
```

```
dm disk1 c1t12d0s2 sliced 2591      17900352 -
dm disk2 c1t14d0s2 sliced 2591      17899056 -
dm disk3 c1t3d0s2  sliced 2591      17899056 -
```

```
v  vol1 -            ENABLED ACTIVE  4196448 ROUND      -          fsgen
pl pl1  vol1         ENABLED ACTIVE  4196448 CONCAT    -          RW
sd disk3-01l1      disk3    0       2098224  0          c1t3d0  ENA
sd sd2   pl1         disk2    0       2098224  2098224   c1t14d0  ENA
```

Evacuate disk2 to disk3.

```
# /etc/vx/bin/vxevac -g voldg disk2 disk3
# vxprint -g voldg -ht
```

DG NAME	NCONFIG	NLOG	MINORS	GROUP-ID			
DM NAME	DEVICE	TYPE	PRIVLEN	PUBLEN	STATE		
RV NAME	RLINK_CNT	KSTATE	STATE	PRIMARY	DATAVOLS	SRL	
RL NAME	RVG	KSTATE	STATE	REM_HOST	REM_DG	REM_RLNK	
V NAME	RVG	KSTATE	STATE	LENGTH	READPOL	PREFPLEX	UTYPE
PL NAME	VOLUME	KSTATE	STATE	LENGTH	LAYOUT	NCOL/WID	MODE
SD NAME	PLEX	DISK	DISKOFFS	LENGTH	[COL/]OFF	DEVICE	MODE
SV NAME	PLEX	VOLNAME	NVOLLAYR	LENGTH	[COL/]OFF	AM/NM	MODE
DC NAME	PARENTVOL	LOGVOL					
SP NAME	SNAPVOL	DCO					

```
dg voldg      default      default 115000
1017856044.1141.hostname.veritas.com
```

dm disk1	c1t12d0s2	sliced	2591	17900352	-
dm disk2	c1t14d0s2	sliced	2591	17899056	-
dm disk3	c1t3d0s2	sliced	2591	17899056	-

v vol1	-	ENABLED	ACTIVE	4196448	ROUND	-	fsgen
pl pl1	vol1	ENABLED	ACTIVE	4196448	CONCAT	-	RW
sd disk3-01	pl1	disk3	0	2098224	0	c1t3d0	ENA
sd disk3-02	pl1	disk3	2098224	2098224	2098224	c1t3d0	ENA

Remove the evacuated disks from VxVM control.

```
# vxdisk -g voldg list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c1t3d0s2	sliced	disk3	voldg	online
c1t12d0s2	sliced	disk1	voldg	online
c1t14d0s2	sliced	disk2	voldg	online

```
# vxdg rmdisk disk1
```

```
# vxdg rmdisk disk2
```

```
# vxdisk rm c1t12d0
```

```
# vxdisk rm c1t14d0
```

Verify that the evacuated disks have been removed from VxVM control.

```
# vxdisk -g voldg list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c1t3d0s2	sliced	disk3	voldg	online

Check to see whether the volume you want to move first is mounted.

```
# mount | grep vol1
/vol1 on /dev/vx/dsk/voldg/vol1
read/write/setuid/log/nolargefiles/dev=12dc138 on Wed Apr
3 10:13:11 2002
```

Create a partition on free disk space of the same size as the volume. In this example, a 2G partition is created on disk1 (c1t12d0s1).

```
# format
Searching for disks...done
```

AVAILABLE DISK SELECTIONS:

0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
1. c1t3d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
/sbus@1f,0/SUNW,fas@2,8800000/sd@3,0
2. c1t9d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
/sbus@1f,0/SUNW,fas@2,8800000/sd@9,0
3. c1t10d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
/sbus@1f,0/SUNW,fas@2,8800000/sd@a,0
4. c1t11d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
/sbus@1f,0/SUNW,fas@2,8800000/sd@b,0
5. c1t12d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
/sbus@1f,0/SUNW,fas@2,8800000/sd@c,0
6. c1t14d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
/sbus@1f,0/SUNW,fas@2,8800000/sd@e,0
7. c1t15d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
/sbus@1f,0/SUNW,fas@2,8800000/sd@f,0

Specify disk (enter its number): 5

selecting c1t12d0

[disk formatted]

FORMAT MENU:

- | | |
|-----------|-------------------------------------|
| disk | - select a disk |
| type | - select (define) a disk type |
| partition | - select (define) a partition table |
| current | - describe the current disk |
| format | - format and analyze the disk |
| repair | - repair a defective sector |
| label | - write label to the disk |
| analyze | - surface analysis |
| defect | - defect list management |
| backup | - search for backup labels |
| verify | - read and display labels |

```

save          - save new disk/partition definitions
inquiry       - show vendor, product and revision
volname       - set 8-character volume name
!<cmd>        - execute <cmd>, then return
quit
format> p

PARTITION MENU:
0          - change '0' partition
1          - change '1' partition
2          - change '2' partition
3          - change '3' partition
4          - change '4' partition
5          - change '5' partition
6          - change '6' partition
7          - change '7' partition
select     - select a predefined table
modify     - modify a predefined partition table
name       - name the current table
print      - display the current table
label      - write partition map and label to the disk
!<cmd>     - execute <cmd>, then return
quit

partition> 1
Part      Tag      Flag      Cylinders      Size      Blocks
  1 unassigned    wm         0           0      (0/0/0)         0
Enter partition id tag[unassigned]:
Enter partition permission flags[wm]:
Enter new starting cyl[0]:
Enter partition size[0b, 0c, 0.00mb, 0.00gb]: 2.00gb
partition> 1
Ready to label disk, continue? y

partition> p
Current partition table (unnamed):
Total disk cylinders available: 13814 + 2 (reserved cylinders)
Part      Tag      Flag      Cylinders      Size      Blocks
  0 unassigned    wm         0           0      (0/0/0)         0
  1 unassigned    wm      0 - 3236    2.00GB  (3237/0/0)    4195152
partition> q

```

Copy the data on vol101 to the newly created disk partition.


```
# dd if=/dev/vx/dsk/voldg/vol01 of=/dev/dsk/clt12d0s1
```

In the `/etc/vfstab` file, remove the following entry.

```
/dev/vx/dsk/voldg/voll /dev/vx/rdisk/voldg/voll /voll vxfs 4 yes rw
```

Replace it with an entry for the newly created partition.

```
/dev/dsk/clt12d0s1 /dev/rdsk/clt12d0s1 /vol01 vxfs 4 yes rw
```

Mount the disk partition.

```
# mount -F vxfs /dev/dsk/clt12d0s1 /vol01
```

Remove `vol01` from VxVM.

```
# vxedit -rf -g voldg rm /dev/vx/dsk/voldg/vol01
```

To complete the procedure, follow the remaining steps.

Preparing to remove Veritas File System

The `VRTSvxfs` package cannot be removed if there are any mounted VxFS file systems or Storage Checkpoints. Unmount the VxFS file systems and Storage Checkpoints before uninstalling Veritas InfoScale. After you remove the `VRTSvxfs` package, VxFS file systems are not mountable or accessible until another `VRTSvxfs` package is installed.

To unmount a file system

- 1 Check if any VxFS file systems are mounted.

```
# cat /etc/mnttab | grep vxfs
```

- 2 Unmount any file systems.

```
# umount special | mount_point
```

Specify the file system to be unmounted as a *mount_point* or *special* (the device on which the file system resides). See the `umount_vxfs(1M)` manual page for more information about this command and its available options.

You can use the `-a` option to unmount all file systems except `/`, `/usr`, `/usr/kvm`, `/var`, `/proc`, `/dev/fd`, and `/tmp`.

To unmount a Storage Checkpoint

- 1 Check if any Storage Checkpoints are mounted.

```
# cat /etc/mnttab | grep vxfs
```

- 2 Unmount any Storage Checkpoints.

```
# umount /checkpoint_name
```

Removing the Replicated Data Set

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

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

To remove the Replicated Data Set

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The RDS is removed.

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

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

Uninstalling Veritas InfoScale packages using the installer

Use the following procedure to remove Veritas InfoScale products.

Not all packages may be installed on your system depending on the choices that you made when you installed the software.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout version in Veritas InfoScale 7.4 with a previous version of Veritas InfoScale.

To shut down and remove the installed Veritas InfoScale packages

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.

- 2 Unmount all mount points for VxFS file systems.

```
# umount /mount_point
```

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

See [“Preparing to remove Veritas Volume Manager”](#) on page 81.

- 4 Make sure you have performed all of the prerequisite steps.

- 5 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
```

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

```
# cd /opt/VRTS/install
```

```
# ./installer -uninstall
```

- 7 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
```

- 8 The uninstall script prompts you to stop the product processes. If you respond yes, the processes are stopped and the packages are uninstalled.

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

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

- 10** To verify the removal of the packages, use the following commands:

Solaris 11:

```
# pkg list VRTS\*
```

- 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:

```
# pkgrm VRTSvxvm
```

Uninstalling Veritas InfoScale using the `pkgrm` or `pkg uninstall` command

Use the following procedure to uninstall Veritas InfoScale using the `pkgrm` command.

If you want to uninstall Veritas InfoScale using the `pkgrm` command, the packages must be removed in a specific order, or else the uninstallation fails. Removing the packages out of order results in some errors, including possible core dumps, although the packages are still removed.

To uninstall Veritas InfoScale

- 1** Unmount all mount points for file systems and Storage Checkpoints.

```
# umount /mount_point
```

Note: Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries can result in system boot problems later.

- 2** Stop all applications from accessing VxVM volumes, and close all volumes.
- 3** For Solaris 11.1 or later, if DMP native support is enabled, DMP controls the ZFS root pool. Turn off native support before removing Veritas InfoScale.

```
# vxddmpadm settune dmp_native_support=off
```

Note: If you do not disable native support, the system cannot be restarted after you remove DMP.

- 4 Stop any Veritas daemons that are running.
- 5 Remove the packages in the following order:
For Veritas InfoScale (Solaris 11):

```
# pkg uninstall VRTSodm VRTSdbed VRTSvcswiz \
VRTSvbs VRTSvcsea VRTSvcscsag VRTScps VRTSvcscs VRTSamf VRTSvxfen \
VRTSgab VRTS1lt VRTSfssdk VRTSfsadv VRTSvxfs VRTSsfmh VRTSaslapm \
VRTSvxvm VRTSspt VRTSsfcpv VRTSvlic VRTSperl
```

Manually uninstalling Veritas InfoScale packages on non-global zones on Solaris 11

- 1 Log on to the non-global zone as a super user.
- 2 Uninstall Veritas InfoScale packages from Solaris zones.

```
# pkg uninstall VRTSperl VRTSvlic VRTSvcscs VRTSvcscsag VRTSvcsea
VRTSvxfs
```

Note: If you have Veritas InfoScale packages installed inside non-global zones, perform the steps mentioned above to uninstall them from non-global zone before attempting to uninstall the packages from global zone.

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
```

- 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 packages 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 packages. 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}="ENTERPRISE74";
$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
-ai	The <code>-ai</code> option is supported on Solaris 11 only, and is used to generate Automated Installation manifest. This can be used by Solaris Automated Installation Server to install the Veritas InfoScale product, along with the Solaris 11 operation system.
-allpkgs	Displays all packages required for the specified product. The packages are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-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.

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

Command Line Option	Function
-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.
-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.

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

Command Line Option	Function
-patch5_path	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-keyfile <i>ssh_key_file</i>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-I ssh_key_file</code> to every SSH invocation.
-license	Registers or updates product licenses on the specified systems.
-logpath <i>log_path</i>	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.
-noipc	Disables the installer from making outbound networking calls to Veritas Services and Operations Readiness Tool (SORT) in order to automatically obtain patch and release information updates.
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
-pkgtable	Displays product's packages in correct installation order by group.
-postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Veritas recommends doing a precheck before installing a product.
-prod	Specifies the product for operations.
-component	Specifies the component for operations.
-redirect	Displays progress details without showing the progress bar.

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

Command Line Option	Function
-require	Specifies an installer patch file.
-requirements	The <code>-requirements</code> option displays required OS version, required packages and patches, file system space, and other system requirements in order to install the product.
-responsefile <i>response_file</i>	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.
-rolling_upgrade	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
-rollingupgrade_phase1	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel packages get upgraded to the latest version.
-rollingupgrade_phase2	The <code>-rollingupgrade_phase2</code> option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent packages upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.
-rootpath <i>root_path</i>	Specifies an alternative root directory on which to install packages. On Solaris operating systems, <code>-rootpath</code> passes <code>-R path</code> to <code>pkgadd</code> command.
-rsh	Specify this option when you want to use rsh and RCP for communication between systems instead of the default ssh and SCP.

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

Command Line Option	Function
<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.
<code>-settunables</code>	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.
<code>-start</code>	Starts the daemons and processes for the specified product.
<code>-stop</code>	Stops the daemons and processes for the specified product.
<code>-timeout</code>	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.
<code>-tmppath <i>tmp_path</i></code>	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 packages are copied on remote systems before installation.
<code>-tunables</code>	Lists all supported tunables and create a tunables file template.
<code>-tunablesfile <i>tunables_file</i></code>	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
<code>-uninstall</code>	This option is used to uninstall the products from systems.

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

Command Line Option	Function
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.
-version	Checks and reports the installed products and their versions. Identifies the installed and missing packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable. Lists the installed patches, patches, and available updates for the installed product if an Internet connection is available.

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 107.

- 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 108.

- 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 109.

See [“About response files”](#) on page 61.

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

See [“Tunables value parameter definitions”](#) on page 111.

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 111.

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 110.
- 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 111.

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 110.
- 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 111.

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 110.
- 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 111.

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_multipathing	(Dynamic Multi-Pathing) Whether DMP will intercept the I/Os directly on the raw OS paths or not.
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.

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

Tunable	Description
dmp_probe_threshold	(Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon.
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_era_nthreads	(Veritas File System) Maximum number of threads VxFS will detect read_ahead patterns on. 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](#)
- [About the VRTSspt package troubleshooting tools](#)
- [Incorrect permissions for root on remote system](#)
- [Inaccessible system](#)

Restarting the installer after a failed network connection

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

About the VRTSspt package troubleshooting tools

The VRTSspt package provides a group of tools for troubleshooting a system and collecting information on its configuration. If you install and use the VRTSspt package, it will be easier for Veritas Support to diagnose any issues you may have.

The tools can gather Veritas File System and Veritas Volume Manager metadata information and establish various benchmarks to measure file system and volume manager performance. Although the tools are not required for the operation of any

Veritas InfoScale product, Veritas recommends installing them should a support case be needed to be opened with Veritas Support. Use caution when you use the VRTSpt package, and always use it in concert with Veritas Support.

Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Failed to setup rsh communication on 10.198.89.241:
'rsh 10.198.89.241 <command>' failed
Trying to setup ssh communication on 10.198.89.241.
Failed to setup ssh communication on 10.198.89.241:
Login denied
```

```
Failed to login to remote system(s) 10.198.89.241.
Please make sure the password(s) are correct and superuser(root)
can login to the remote system(s) with the password(s).
If you want to setup rsh on remote system(s), please make sure
rsh with command argument ('rsh <host> <command>') is not
denied by remote system(s).
```

```
Either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

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

```
System verification did not complete successfully
```

```
The following errors were discovered on the systems:
```

```
The ssh permission denied on 10.198.89.241
rsh exited 1 on 10.198.89.241
either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

Suggested solution: You need to set up the systems to allow remote access using **ssh** or **rsh**.

Note: Remove remote shell permissions after completing the Veritas InfoScale installation and configuration.

Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
Verifying systems: 12% .....
Estimated time remaining: 0:10 1 of 8
Checking system communication ..... Done
System verification did not complete successfully
The following errors were discovered on the systems:
cannot resolve hostname host1
Enter the 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.