

Veritas Access Installation Guide

Linux

7.3.1

Veritas Access Installation Guide

Last updated: 2018-07-27

Document version: 7.3.1 Rev 0

Legal Notice

Copyright © 2018 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Veritas InfoScale, and NetBackup 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/licensing/process>

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 document version appears on page 2 of each guide. The latest documentation is available on the Veritas website:

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

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing Veritas Access	8
	About Veritas Access	8
Chapter 2	Licensing in Veritas Access	13
	About Veritas Access product licensing	13
Chapter 3	System requirements	16
	Important release information	16
	System requirements	16
	Linux requirements	17
	Software requirements for installing Veritas Access in a VMware ESXi environment	39
	Hardware requirements for installing Veritas Access virtual machines	39
	Management Server Web browser support	40
	Supported NetBackup versions	40
	Supported OpenStack versions	41
	Supported Oracle versions and host operating systems	41
	Supported IP version 6 Internet standard protocol	41
	Network and firewall requirements	42
	NetBackup ports	44
	OpenDedup ports and disabling the iptable rules	45
	CIFS protocols and firewall ports	46
	Maximum configuration limits	47
Chapter 4	Preparing to install Veritas Access	49
	Overview of the installation process	49
	Hardware requirements for the nodes	51
	About using LLT over the RDMA network for Veritas Access	51
	RDMA over InfiniBand networks in the Veritas Access clustering environment	52
	How LLT supports RDMA for faster interconnections between applications	52
	Configuring LLT over RDMA for Veritas Access	53

	How the Veritas Access installer configures LLT over RDMA	54
	LLT over RDMA sample /etc/littab	54
	Connecting the network hardware	55
	About obtaining IP addresses	57
	About calculating IP address requirements	58
	Reducing the number of IP addresses required at installation time	61
	About checking the storage configuration	62
Chapter 5	Deploying virtual machines in VMware ESXi for Veritas Access installation	63
	Setting up networking in VMware ESXi	63
	Creating a datastore for the boot disk and LUNs	64
	Creating a virtual machine for Veritas Access installation	65
Chapter 6	Installing and configuring a cluster	69
	Installation overview	70
	Summary of the installation steps	70
	Before you install	71
	Installing the operating system on each node of the cluster	72
	About the driver node	72
	Installing the operating system on the target Veritas Access cluster	73
	Installing the Oracle Linux operating system on the target Veritas Access cluster	75
	Obtaining the Red Hat Enterprise Linux compatible kernels	76
	Installing Veritas Access on the target cluster nodes	78
	Installing and configuring the Veritas Access software on the cluster	79
	Veritas Access Graphical User Interface	84
	About NIC bonding and NIC exclusion	85
	Excluding a NIC	86
	Including a NIC	90
	Creating a new NIC bond	94
	Removing a NIC bond	99
	Removing a NIC from the bond list	102
	About VLAN Tagging	104
	Adding a VLAN device on a particular NIC	104
	Limitations of VLAN Tagging	105
	Replacing an Ethernet interface card	105
	Configuring I/O fencing	106
	About configuring Veritas NetBackup	107

	About enabling kdump during an Veritas Access configuration	107
	Reconfiguring the Veritas Access cluster name and network	108
	Configuring a KMS server on the Veritas Access cluster	110
Chapter 7	Automating Veritas Access installation and configuration using response files	111
	About response files	111
	Performing a silent Veritas Access installation	112
	Response file variables to install and configure Veritas Access	112
	Sample response file for Veritas Access installation and configuration	121
Chapter 8	Displaying and adding nodes to a cluster	124
	About the Veritas Access installation states and conditions	124
	Displaying the nodes in the cluster	125
	Before adding new nodes in the cluster	127
	Adding a node to the cluster	129
	Deleting a node from the cluster	132
	Shutting down the cluster nodes	134
Chapter 9	Upgrading Veritas Access and operating system	135
	Upgrading the operating system and Veritas Access	135
Chapter 10	Upgrading Veritas Access using a rolling upgrade	142
	About rolling upgrades	142
	Supported rolling upgrade paths for upgrades on RHEL	145
	Performing a rolling upgrade using the installer	145
Chapter 11	Uninstalling Veritas Access	149
	Before you uninstall Veritas Access	149
	Uninstalling Veritas Access using the installer	151
	Removing Veritas Access 7.3.1 RPMs	151
	Running uninstall from the Veritas Access 7.3.1 disc	152
Appendix A	Installation reference	153
	Installation script options	153

Appendix B	Troubleshooting the LTR upgrade	155
	Locating the log files for troubleshooting the LTR upgrade	155
	Troubleshooting pre-upgrade issues for LTR	155
	Troubleshooting post-upgrade issues for LTR	156
Appendix C	Configuring the secure shell for communications	157
	Manually configuring passwordless secure shell (ssh)	157
	Setting up ssh and rsh connections using the pwutil.pl utility	160
Index		165

Introducing Veritas Access

This chapter includes the following topics:

- [About Veritas Access](#)

About Veritas Access

Veritas Access is a software-defined scale-out network-attached storage (NAS) solution for unstructured data that works on commodity hardware. Veritas Access provides resiliency, multi-protocol access, and data movement to and from the public or private cloud based on policies.

You can use Veritas Access in any of the following ways.

Table 1-1 Interfaces for using Veritas Access

Interface	Description
GUI	Centralized dashboard with operations for managing your storage. See the GUI and the Online Help for more information.
RESTful APIs	Enables automation using scripts, which run storage administration commands against the Veritas Access cluster. See the <i>Veritas Access RESTful API Guide</i> for more information.
Command-line interface (CLI or CLISH)	Single point of administration for the entire cluster. See the manual pages for more information.

[Table 1-2](#) describes the features of Veritas Access.

Table 1-2 Veritas Access key features

Feature	Description
Multi-protocol access	Veritas Access includes support for the following protocols: <ul style="list-style-type: none">■ Amazon S3■ CIFS■ FTP■ iSCSI target■ NFS■ Oracle Direct NFS■ SMB 3■ NFS with S3
WORM storage for Enterprise Vault Archiving	Veritas Access can be configured as WORM primary storage for archival by Enterprise Vault. Veritas Access is certified as a CIFS primary WORM storage for Enterprise Vault 12.1. For more information, see the <i>Veritas Access Enterprise Vault Solutions Guide</i> .
WORM support over NFS	Veritas Access supports WORM over NFS.
Creation of Partition Secure Notification (PSN) file for Enterprise Vault Archiving	A Partition Secure Notification (PSN) file is created at a source partition after the successful backup of the partition at the remote site. For more information, see the <i>Veritas Access Enterprise Vault Solutions Guide</i> .
Managing application I/O workloads using maximum IOPS settings	The MAXIOPS limit determines the maximum number of I/Os processed per second collectively by the storage underlying the file system.
Flexible Storage Sharing (FSS)	Enables cluster-wide network sharing of local storage.

Table 1-2 Veritas Access key features (*continued*)

Feature	Description
Scale-out file system	<p>The following functionality is provided for a scale-out file system:</p> <ul style="list-style-type: none"> ■ File system that manages a single namespace spanning over both on-premises storage as well as cloud storage, which provides better fault tolerance for large data sets. ■ Highly available NFS and S3 shares. You use scale-out file systems if you want to store a large capacity of data in a single namespace (3 PB is the maximum file system size). ■ Creation of CIFS shares. ■ File sharing for a scale-out file system using FTP.
Cloud as a tier for a scale-out file system	<p>Veritas Access supports adding a cloud service as a storage tier for a scale-out file system. You can move data between the tiers based on file name patterns and when the files were last accessed or modified. Use scheduled policies to move data between the tiers on a regular basis.</p> <p>Veritas Access moves the data from the on-premises tier to Amazon S3, Amazon Glacier, Amazon Web Services (AWS), GovCloud (US), Azure, Google cloud, Alibaba, Veritas Access S3, IBM Cloud Object Storage, and any S3-compatible storage provider based on automated policies. You can also retrieve data archived in Amazon Glacier.</p>
SmartIO	<p>Veritas Access supports both read and writeback caching on solid state drives (SSDs) for applications running on Veritas Access file systems.</p>
SmartTier	<p>Veritas Access's built-in SmartTier feature can reduce the cost of storage by moving data to lower-cost storage. Veritas Access storage tiering also facilitates the moving of data between different drive architectures and on-premises.</p>
Snapshot	<p>Veritas Access supports snapshots for recovering from data corruption. If files, or an entire file system, are deleted or become corrupted, you can replace them from the latest uncorrupted snapshot.</p>
Deduplication	<p>You can run post-process periodic deduplication in a file system, which eliminates duplicate data without any continuous cost.</p>

Table 1-2 Veritas Access key features (*continued*)

Feature	Description
Compression	You can compress files to reduce the space used, while retaining the accessibility of the files and having the compression be transparent to applications. Compressed files look and behave almost exactly like uncompressed files: the compressed files have the same name, and can be read and written as with uncompressed files.
Erasure-coding	Erasure-coding is configured with the <i>EC log</i> option for NFS use case.
IP load balancing	With IP load balancing, a single virtual IP is used to act as a load balancer IP which distributes the incoming requests to the different nodes in the Veritas Access cluster for the services that are run on an active-active cluster.
iSCSI target	Veritas Access as an iSCSI target can be configured to serve block storage. iSCSI target as service is hosted in an active-active mode in the Veritas Access cluster.
NetBackup integration	Built-in NetBackup client for backing up your file systems to a NetBackup master or media server. Once data is backed up, a storage administrator can delete unwanted data from Veritas Access to free up expensive primary storage for more data.
OpenStack plug-in	Integration with OpenStack: <ul style="list-style-type: none"> ■ OpenStack Cinder integration that allows OpenStack instances to use the storage hosted by Veritas Access. ■ OpenStack Manila integration that lets you share Veritas Access file systems with virtual machines on OpenStack Manila.
Quotas	Support for setting file system quotas, user quotas, and hard quotas.
Replication	Periodic replication of data over IP networks. See the <code>episodic(1)</code> man page for more information. Synchronous replication of data over IP networks See the <code>continuous(1)</code> man page for more information.

Table 1-2 Veritas Access key features (*continued*)

Feature	Description
Support for LDAP, NIS, and AD	Veritas Access uses the Lightweight Directory Access Protocol (LDAP) for user authentication.
Partition Directory	<p>With support for partitioned directories, directory entries are redistributed into various hash directories. These hash directories are not visible in the name-space view of the user or operating system. For every new create, delete, or lookup, this feature performs a lookup for the respective hashed directory and performs the operation in that directory. This leaves the parent directory inode and its other hash directories unobstructed for access, which vastly improves file system performance.</p> <p>By default this feature is not enabled. See the <code>storage_fs(1)</code> manual page to enable this feature.</p>
Isolated storage pools	Enables you to create an isolated storage pool with a self-contained configuration. An isolated storage pool protects the pool from losing the associated metadata even if all the configuration disks in the main storage pool fail.
Performance and tuning	<p>Workload-based tuning for the following workloads:</p> <ul style="list-style-type: none">■ Media server - Streaming media represents a new wave of rich Internet content. Recent advancements in video creation, compression, caching, streaming, and other content delivery technology have brought audio and video together to the Internet as rich media. You can use Veritas Access to store your rich media, videos, movies, audio, music, and photos.■ Virtual machine support■ Other workloads

Licensing in Veritas Access

This chapter includes the following topics:

- [About Veritas Access product licensing](#)

About Veritas Access product licensing

You have to obtain a license to install and use Veritas Access.

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

- Enter a valid perpetual license key file matching the functionality in use on the systems.
A perpetual license is like a permanent license for using Veritas Access.
- Enter a valid subscription license key file matching the functionality in use on the systems.
A subscription license is a license with validity of one year.
- Continue with evaluation mode, and complete system licensing later.
This license is a trialware which can be used for 60 days.
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.

To comply with the terms of the End User License Agreement, you have 60 days to either enter a valid subscription or perpetual license key or continue in evaluation mode.

You can invoke the license program using the `./installaccess -license` command.

During the installation, you get the following prompt:

- 1) Enter a valid perpetual or subscription license key file
- 2) Continue with evaluation mode and complete system licensing later

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

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

Enter the location of a valid ACCESS perpetual or subscription license key file: [b]

Note: Ensure that you place the license file in a folder on the system. You may get an invalid license key error if your file is placed in the “/” location.

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

www.veritas.com/licensing/process

The Veritas Access licensing has a few functional enforcements.

Table 2-1 Functional enforcements of Veritas Access licensing

Enforcement	Action
During Validity	None
During Grace period	Nagging message (in the GUI only)
Post Grace Period	<p>Before you restart the node, you can stop the NFS, CIFS, FTP, and S3 services, but you cannot start the services again (even if you have not restarted the node).</p> <p>After you restart the node, the NFS, CIFS, FTP, and S3 services do not come ONLINE on the restarted node.</p>

If you add the Veritas Access license using the GUI:

- When a node is restarted after the license has expired, the NFS, CIFS, FTP, and S3 services are stopped on that node. The status of the service appears ONLINE if the service is running anywhere in the cluster, even if it is OFFLINE on this node. Check the alerts on each node individually to see if the service is ONLINE or OFFLINE locally.

- An option to start, stop, and check the status of NFS, CIFS, and S3 services is available. You cannot start, stop, or check the status of the FTP service.
- You can only provide the license file from the local system, the `scp` path is not supported through the GUI.

If you add the Veritas Access license using the CLISH:

- When a node is restarted after the license has expired, the NFS, CIFS, FTP, and S3 services are stopped on that node. You can use the `support services show` command to display the node-wise status of the service.
- An option to start, stop, and check the status of NFS, CIFS, FTP, and S3 services is available.
- You can add the license using the `license add` command. The `license add` command provides support for `scp` path as well.
- The `license list` and `license list details` commands give details of the license installed on each node of the cluster.

System requirements

This chapter includes the following topics:

- [Important release information](#)
- [System requirements](#)
- [Network and firewall requirements](#)
- [Maximum configuration limits](#)

Important release information

Review the *Veritas Access Release Notes* for the latest information before you install the product.

For important updates regarding this release, review the Late-Breaking News tech note on the Veritas Technical Support website:

https://www.veritas.com/support/en_US/article.100040711

System requirements

[Table 3-1](#) lists the per-node system requirements for running the Veritas Access system software.

Table 3-1 System requirements for Veritas Access

Minimum	Recommended
Each Veritas Access node using a 64-bit Intel-based server architecture that is compatible with RHEL 6 Update 6, 7, and 8, RHEL 7 Update 3 and 4 or AMD64 and Intel EMT. Itanium is not supported.	Two nodes of dual or quad core processors at 2.0 GHz or later for optimal performance.
32 GB error-correcting code (ECC) random-access memory (RAM).	The recommended values depend on the expected workload.
One internal drive with size equal to size of RAM + 60GB.	Dual boot drives each of size RAM + 60 GB or more capacity. In an FSS-based environment, additional internal drives (SSD + HDD) are recommended.
Two gigabit Ethernet interfaces	Embedded Ethernet interfaces are recommended.
Two additional gigabit Ethernet interfaces (Embedded or peripheral component interconnect (PCI) based.)	N/A
One Fibre Channel Host Bus Adapters (HBA).	Two Fibre Channel Host Bus Adapters (HBAs) are recommended for high availability (HA) if you are using shared LUNs that need to be mapped over a Fibre Channel protocol. If the environment has only DAS or iSCSI disks, then the HBA requirement is optional.
Internal/external USB DVD-ROM DVD drive.	N/A
Redundant power supply.	Recommended, but not required.
SmartIO caching feature	A PCI-based SSD card is recommended if you want to use the SmartIO caching feature.
Minimum number of servers required is 1	N/A

Linux requirements

Veritas Access does not support the operating system on which Veritas Access runs. There are strict operating system versioning requirements for each release of Veritas Access.

The Veritas Access 7.3.1 release requires Red Hat Enterprise Linux (RHEL) or the Oracle Linux (OL).

The Veritas Access 7.3.1 release supports the following operating systems:

- Red Hat Enterprise Linux (RHEL)
 - RHEL 6 Update 6, 7, and 8
 - RHEL 7 Update 3 and 4
- Oracle Linux (only in RHEL compatible mode)
 - OL 6 Update 8
 - OL 7 Update 3 and 4

The minimum operating system requirements are enforced during the Veritas Access installation. A Kickstart file is also available on request for Veritas Access 7.3.1 to assist partners with the operating system installation requirements. Operating system patches, including security vulnerability patches, can be installed without requiring certification from Veritas. However, operating system Kernel RPMs should not be patched without specific approval from Veritas.

The certification of the Red Hat Enterprise Linux (RHEL) operating system updates can require a new minor version of Veritas Access. RHEL operating system updates cannot be installed without prior agreement with Veritas.

Veritas Access can be installed on computers running the following operating systems:

Requirement	Version	Version	Version	Version	Version
Red Hat Enterprise Linux version	RHEL 6 Update 6	RHEL 6 Update 7	RHEL 6 Update 8	RHEL 7 Update 3	RHEL 7 Update 4
Oracle Linux			OL 6 Update 8	OL 7 Update 3	OL 7 Update 4
Kernel version	2.6.32-504.el6	2.6.32-573.el6	2.6.32-642.el6	3.10.0-514.el7	3.10.0-693
	See “Kernel RPMs that are required to be installed with exact predefined RPM versions” on page 21.	See “Kernel RPMs that are required to be installed with exact predefined RPM versions” on page 21.	See “Kernel RPMs that are required to be installed with exact predefined RPM versions” on page 21.		

Requirement	Version	Version	Version	Version	Version
Required RPMs	See “Required operating system RPMs for RHEL 6.6” on page 29.	See “Required operating system RPMs for RHEL 6.7” on page 31.	See “Required operating system RPMs for RHEL 6.8” on page 32.	See “Required operating system RPMs for RHEL 7.3” on page 33.	See “Required operating system RPMs for RHEL 7.4” on page 36.
			See “Required operating system RPMs for OL 6.8” on page 23.	See “Required operating system RPMs for OL 7.3” on page 24.	See “Required operating system RPMs for OL 7.4” on page 27.

Operating system RPM installation requirements and operating system patching

Veritas has categorized the operating system RPMs that are required before you install Veritas Access into four groups:

Category 1

- This set of RPMs are kernel RPMs that are required to be installed with exact predefined RPM versions only.
- The required RPM versions are different for RHEL 6.6, RHEL 6.7, RHEL 6.8, RHEL 7.3, and RHEL 7.4.
- The required RPM versions are different for OL 6.8, OL 7.3, and OL 7.4.
- The RPMs in this category should not be patched without specific approval from Veritas.
- See [“Kernel RPMs that are required to be installed with exact predefined RPM versions”](#) on page 21.
- See [“OL kernel RPMs that are required to be installed with exact predefined RPM versions”](#) on page 22.

Category 2

- This set of RPMs include the OS libs and OS packages that must be installed with minimum predefined RPM versions.
- The required RPM versions are different for RHEL 6.6, RHEL 6.7, RHEL 6.8, RHEL 7.3, and RHEL 7.4.

- The required RPM versions are different for OL 6.8, OL 7.3, and OL 7.4.
- The RPMs in this category can be patched using official Red Hat patches.
- An approval or certification from Veritas is not required to patch these RPMs.
- See [“Required operating system RPMs for OL 6.8”](#) on page 23.
- See [“Required operating system RPMs for OL 7.3”](#) on page 24.
- See [“Required operating system RPMs for OL 7.4”](#) on page 27.
- See [“Required operating system RPMs for RHEL 6.6”](#) on page 29.
- See [“Required operating system RPMs for RHEL 6.7”](#) on page 31.
- See [“Required operating system RPMs for RHEL 6.8”](#) on page 32.
- See [“Required operating system RPMs for RHEL 7.3”](#) on page 33.
- See [“Required operating system RPMs for RHEL 7.4”](#) on page 36.

Category 3

- This set of RPMs are required by Category 2 RPMs as dependencies, their installation is enforced by Red Hat.
- Veritas Access does not require any specific versions of these RPMs to be installed.
- The versions of these RPMs are determined by Red Hat.
- The RPMs in this category can be patched using official Red Hat patches.
- An approval or certification from Veritas is not required to patch these RPMs.
- Veritas does not document these RPMs as required RPMs for Veritas Access.

Category 4

- This set of RPMs are third-party RPMs that are included in the Veritas Access ISO.
- These RPMs are not operating system RPMs. It includes Samba, Ganesha, and other third-party products.
- The RPMs in this category should not be patched without specific approval from Veritas.
- Veritas installs these RPMs as they are included in the Veritas Access ISO.

Kernel RPMs that are required to be installed with exact predefined RPM versions

After you install Red Hat Enterprise Linux operating system, install the following RPMs and then restart the system. Use the specified links to download the RPMs from the Red Hat site. The links require a Red Hat registration ID (user name and password).

The RHEL 6 Update 6 kernel packages:

- kernel-2.6.32-504.el6.x86_64.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1006010>
- kernel-debuginfo-2.6.32-504.el6.x86_64.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1005367>
- kernel-debuginfo-common-x86_64-2.6.32-504.el6.x86_64.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1006274>
- kernel-headers-2.6.32-504.el6.x86_64.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1006768>
- kernel-firmware-2.6.32-504.el6.noarch.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1005320>

The RHEL 6 Update 7 kernel packages:

- kernel-2.6.32-573.el6.x86_64.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1072387>
- kernel-debuginfo-2.6.32-573.18.1.el6.x86_64.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1072442>
- kernel-debuginfo-common-x86_64-2.6.32-573.18.1.el6.x86_64.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1072385>
- kernel-firmware-2.6.32-573.el6.noarch.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1012920>
- kernel-headers-2.6.32-573.el6.x86_64.rpm
<https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1072371>

The RHEL 6 Update 8 kernel packages:

- kernel-2.6.32-642.el6.x86_64.rpm
https://access.redhat.com/downloads/content/kernel/2.6.32-642.el6/x86_64/fd431d51/package
- kernel-debuginfo-2.6.32-642.el6.x86_64.rpm
https://access.redhat.com/downloads/content/kernel-debuginfo/2.6.32-642.el6/x86_64/fd431d51/package

- kernel-debuginfo-common-x86_64-2.6.32-642.el6.x86_64.rpm
https://access.redhat.com/downloads/content/kernel-debuginfo-common-x86_64/2.6.32-642.el6/x86_64/fd431d51/package
- kernel-headers-2.6.32-642.el6.x86_64.rpm
https://access.redhat.com/downloads/content/kernel-headers/2.6.32-642.el6/x86_64/fd431d51/package
- kernel-firmware-2.6.32-642.el6.noarch.rpm
<https://access.redhat.com/downloads/content/kernel-firmware/2.6.32-642.el6/noarch/fd431d51/package>
- <https://rhn.redhat.com/rhn/software/packages/details/Overview.do?pid=1072371>

The RHEL 7 Update 3 kernel packages:

The following RPMs are included in the DVD image under the **os_rpms** directory and are installed using the CPI installation.

- kernel-debuginfo-3.10.0-514.el7.x86_64.rpm
- kernel-headers-3.10.0-514.el7.x86_64.rpm
- kernel-debuginfo-common-x86_64-3.10.0-514.el7.x86_64.rpm

The RHEL 7 Update 4 kernel packages:

The following RPMs are included in the DVD image under the **os_rpms** directory and are installed using the CPI installation.

- kernel-headers-3.10.0-693.el7.x86_64.rpm

OL kernel RPMs that are required to be installed with exact predefined RPM versions

The OL environment should be one of the following with Red Hat compatible kernels only:

- OL 6.8
- OL 7.3
- OL 7.4

Note: For OL 6.x and 7.x, uek kernel is not supported.

Example:

```
[root@pilotvm13 ~]# cat /etc/redhat-release
Red Hat Enterprise Linux Server release 6.7 (Santiago)
[root@pilotvm13 ~]# cat /etc/oracle-release
```

```
Oracle Linux Server release 6.8
[root@pilotvm13 ~]# uname -r
2.6.32-573.el6.x86_64
```

Required operating system RPMs for OL 6.8

The RPM version numbers specified in these lists are the minimum required version numbers for these operating system RPMs.

Required OS lib rpms for OL6.8:

```
coreutils-8.4-43.0.1.el6.x86_64
ed-1.1-3.3.el6.x86_64
glibc-2.12-1.192.el6.x86_64
libstdc++-4.4.7-17.el6.i686
module-init-tools-3.9-25.0.1.el6.x86_64
parted-2.1-29.el6.x86_64
python-2.6.6-64.el6.x86_64
readline-6.0-4.el6.x86_64

findutils-4.4.2-9.el6.x86_64
libacl-2.2.49-6.el6.x86_64
libstdc++-4.4.7-17.el6.x86_64
ncurses-libs-5.7-4.20090207.el6.x86_64
perl-5.10.1-141.el6_7.1.x86_64
python-libs-2.6.6-64.el6.x86_64
zlib-1.2.3-29.el6.x86_64
```

```
glibc-2.12-1.192.el6.i686
libgcc-4.4.7-17.el6.x86_64
libudev-147-2.73.0.1.el6.i686
openssl-1.0.1e-48.el6_8.3.x86_64
policycoreutils-2.0.83-29.0.1.el6.x86_64
python-setuptools-0.6.10-3.el6
```

Required OS packages for OL6.8:

```
apr-devel 1.3.9-5.el6_2
bash 4.1.2-40.el6
glibc-devel.x86_64 2.12-1.192.el6
glibc-utils 2.12-1.192.el6
glibc.x86_64 2.12-1.192.el6
httpd 2.2.15-54.el6_8
httpd-manual 2.2.15-54.el6_8
infiniband-diags 1.6.4-1.el6
kernel-debuginfo-2.6.32-642.el6
kernel-firmware-2.6.32-642.el6
ksh 20120801-33.el6
libibverbs-devel 1.1.8-4.el6
libpcap 1.4.0
libyaml 0.1.3-4
memcached 1.4.4-3.el6
mod_wsgi 3.2-7.el6
net-snmp 5.5-57.el6

apr-util-devel 1.3.9-3.el6_0.1
glibc-common 2.12-1.192.el6
glibc-headers 2.12-1.192.el6
glibc.i686 2.12-1.192.el6
hal 0.5.14-14.el6
httpd-devel 2.2.15-54.el6_8
httpd-tools 2.2.15-54.el6_8
iscsi-initiator-utils 6.2.0.873-21.el6
kernel-debuginfo-common-x86_64-2.6.32-642.el6
kernel-headers-2.6.32-642.el6
libibumad 1.3.9-1.el6
libibverbs-utils 1.1.8-4.el6
libtirpc 0.2.1-11.el6
lsscsi 0.23
mod_ssl 2.2.15-54.el6_8
nc 1.84-24.el6
net-snmp-libs 5.5-57.el6
```

```
net-snmp-utils 5.5-57.e16
nscd 2.12-1.192.e16
openldap 2.4.40-12.e16
opensm 3.3.17-1.e16
openssl 1.0.1e-48.e16_8.1
perl 5.10.1
python-backports 1.0-5.e16
python-chardet 2.2.1-1.e16
python-paramiko 1.7.5-2.1.e16
python-six 1.9.0-2.e16
rdma 6.8_4.1-1.e16
scsi-target-utils 1.0.24-18.e16
sg3_utils-libs 1.28-12.e16
telnet 0.17-48.e16
wireshark 1.8.10-17.e16

nfs-utils 1.2.3-70.e16
nss-pam-ldapd 0.7.5-32.e16
openldap-clients 2.4.40-12.e16
opensm-libs 3.3.17-1.e16
openssl-devel 1.0.1e-48.e16_8.1
perl-Net-Telnet 3.03
python-backports-ssl_match_hostname 3.4.0.2-2.e16
python-memcached 1.43-6.e16
python-requests 2.6.0-3.e16
python-urllib3 1.10.2-1.e16
rrdtool 1.3.8-7.e16
sg3_utils 1.28-12.e16
sysstat 9.0.4-31.e16
vsftpd 2.2.2-21.e16
```

Required operating system RPMs for OL 7.3

The RPM version numbers specified in these lists are the minimum required version numbers for these operating system RPMs.

Required OS lib rpms for OL 7.3:

```
bc-1.06.95-13.e17.x86_64
ed-1.9-4.e17.x86_64
glibc-2.17-157.e17.x86_64
libgcc-4.8.5-11.e17.x86_64
openssl-libs-1.0.1e-60.e17.x86_64
perl-Socket-2.010-4.e17.x86_64
python-2.7.5-48.e17.x86_64
zlib-1.2.7-17.e17.x86_64

coreutils-8.22-18.e17.x86_64
findutils-4.5.11-5.e17.x86_64
libacl-2.2.51-12.e17.x86_64
libstdc++-4.8.5-11.e17.x86_64
perl-Exporter-5.68-3.e17.noarch
policycoreutils-2.5-8.e17.x86_64
python-libs-2.7.5-48.e17.x86_64
```

Required OS packages for OL 7.3:

```
apr-devel 1.4.8-3
apr-util-devel 1.5.2-6
arptables 0.0.4-8
at 3.1.13-22
autogen-libopts 5.18-5
avahi-libs 0.6.31-17
bash 4.2.46-20
binutils 2.25.1-22
cairo 1.14.2-1
coreutils 8.22-18
cups-libs 1.6.3-26
ethtool 4.5-3
```

fuse 2.9.2-7
fuse-devel 2.9.2-7
fuse-libs 2.9.2-7
glibc-common 2.17.157
glibc-devel.x86_64 2.17.157
glibc-headers 2.17.157
glibc-utils 2.17.157
glibc.i686 2.17.157
glibc.x86_64 2.17.157
httpd 2.4.6-45
httpd-devel 2.4.6-45
httpd-manual 2.4.6-45
httpd-tools 2.4.6-45
infiniband-diags 1.6.5-3
initscripts 9.49.37-1
iproute 3.10.0-74
ipvsadm 1.27-7
iscsi-initiator-utils 6.2.0.873-35
jansson 2.10-1
kernel-debuginfo =3.10.0-514.el7
kernel-debuginfo-common-x86_64 =3.10.0-514.el7
kernel-headers =3.10.0-514.el7
kmod 20-9
krb5-devel 1.14.1-26
krb5-libs 1.14.1-26
krb5-workstation 1.15.1-8
ksh 20120801-26.el7
libibumad 1.3.10.2-1
libibverbs-utils 1.2.1-1
libjpeg-turbo 1.2.90-5
libpcap 1.5.3-8
libtirpc 0.2.4-0.8
libyaml 0.1.4-11
lshw B.02.17-12
lsof 4.87-4
lsscsi 0.27-4
memcached 1.4.15-10
mlocate 0.26-6
mod_ssl 2.4.6-45
mod_wsgi 3.4-12
net-snmp 5.7.2-24
net-snmp-utils 5.7.2-24
net-tools 2.0-0.17

nfs-utils 1.3.0-0.33
nmap-ncat 6.40-7
nscd 2.17-157
nss-pam-ldapd 0.8.13-8
ntp 4.2.6p5-25
ntpdate 4.2.6p5-25
openldap 2.4.40-13
openldap-clients 2.4.40-13
opensm 3.3.19-1
opensm-libs 3.3.19-1
openssl 1.0.1e-60
openssl-devel 1.0.1e-60
pango 1.36.8-2
perl 5.16.3
perl-Convert-ASN1 0.26-4
perl-JSON 2.59-2
perl-LDAP 0.56-5
perl-Net-Telnet 3.03-19.e17
perl-XML-Parser 2.41-10
psmisc 22.20-11
python-backports 1.0-8
python-backports-ssl_match_hostname 3.4.0.2-4
python-chardet 2.2.1-1
python-memcached 1.48-4
python-paramiko 1.7.7.1-3
python-requests 2.6.0-1
python-setuptools 0.9.8-4
python-six 1.9.0-2
python-urllib3 1.10.2-2
PyYAML 3.10-11
rdma 7.3_4.7_rc2-5
rrdtool 1.4.8-9
rsh 0.17-76
sg3_utils 1.37-9
sg3_utils-libs 1.37-9
strace 4.8-11
sysstat 10.1.5-11
targetcli 2.1.fb41-3
telnet 0.17-60
traceroute 2.0.22-2
tzdata-java
unzip 6.0-16
vim-enhanced 7.4.160

```
vsftpd 3.0.2-21  
wireshark 1.10.14-10  
yp-tools 2.14-3  
ypbind 1.37.1-7  
zip 3.0-11
```

Required operating system RPMs for OL 7.4

The RPM version numbers specified in these lists are the minimum required version numbers for these operating system RPMs.

Required OS packages for OL 7.4:

```
PyYAML 3.10-11  
apr-devel 1.4.8-3  
apr-util-devel 1.5.2-6  
arptables 0.0.4-8  
at 3.1.13-22  
autogen-libopts 5.18-5  
avahi-libs 0.6.31-17  
bash 4.2.46-28  
binutils 2.25.1-31  
cairo 1.14.8-2  
coreutils 8.22-18  
cups-libs 1.6.3-29  
ethtool 4.8-1  
fuse 2.9.2-8  
fuse-devel 2.9.2-8  
fuse-libs 2.9.2-8  
glibc-common 2.17.196  
glibc-devel.x86_64 2.17.196  
glibc-headers 2.17.196  
glibc-utils 2.17.196  
glibc.i686 2.17.196  
glibc.x86_64 2.17.196  
httpd 2.4.6-67  
httpd-devel 2.4.6-67  
httpd-manual 2.4.6-67  
httpd-tools 2.4.6-67  
infiniband-diags 1.6.7-1  
initscripts 9.49.39-1  
iproute 3.10.0-87  
ipvsadm 1.27-7  
iscsi-initiator-utils 6.2.0.874-4
```

jansson 2.10-1
kmod 20-15
krb5-devel 1.15.1-8
krb5-libs 1.15.1-8
krb5-workstation 1.15.1-8
libibumad 13-7
libibverbs-utils 13-7
libjpeg-turbo 1.2.90-5
libpcap 1.5.3-9
libtirpc 0.2.4-0.10
libyaml 0.1.4-11
lshw B.02.18-7
lsof 4.87-4
lsscsi 0.27-6
memcached 1.4.15-10
mlocate 0.26-6
mod_ssl 2.4.6-67
mod_wsgi 3.4-12
net-snmp 5.7.2-28
net-snmp-utils 5.7.2-28
net-tools 2.0-0.22
nfs-utils 1.3.0-0.48
nmap-ncat 6.40-7
nscd 2.17-196
nss-pam-ldapd 0.8.13-8
ntp 4.2.6p5-25
ntpdate 4.2.6p5-25
openldap 2.4.44-5
openldap-clients 2.4.44-5
opensm 3.3.19-1
opensm-libs 3.3.19-1
openssl 1.0.2k-8
openssl-devel 1.0.2k-8
pango 1.40.4-1
perl 5.16.3
perl-Convert-ASN1 0.26-4
perl-JSON 2.59-2
perl-LDAP 0.56-5
perl-Net-Telnet 3.03-19.e17
perl-XML-Parser 2.41-10
psmisc 22.20-15
python-backports 1.0-8
python-backports-ssl_match_hostname 3.4.0.2-4

```
python-chardet 2.2.1-1
python-memcached 1.48-4
python-paramiko 1.7.7.1-3
python-requests 2.6.0-1
python-setuptools 0.9.8-7
python-six 1.9.0-2
python-urllib3 1.10.2-3
rrdtool 1.4.8-9
rsh 0.17-76
sg3_utils 1.37-12
sg3_utils-libs 1.37-12
strace 4.12-4
sysstat 10.1.5-12
targetcli 2.1.fb46-1
telnet 0.17-64
traceroute 2.0.22-2
tzdata-java
unzip 6.0-16
vim-enhanced 7.4.160
vsftpd 3.0.2-22
wireshark 1.10.14-14
yp-tools 2.14-5
ypbind 1.37.1-9
zip 3.0-11
```

Required operating system RPMs for RHEL 6.6

The RPM version numbers specified in these lists are the minimum required version numbers for these operating system RPMs.

Required OS lib rpms for RHEL6.6:

```
coreutils-8.4-37.el6.x86_64          ed-1.1-3.3.el6.x86_64
findutils-4.4.2-6.el6.x86_64        glibc-2.12-1.166.el6_7.7.i686
glibc-2.12-1.192.el6.i686           glibc-2.12-1.192.el6.x86_64
libgcc-4.4.7-11.el6.i686            libgcc-4.4.7-11.el6.x86_64
libstdc++-4.4.7-11.el6.x86_64       module-init-tools-3.9-24.el6_4.x86_64
openssl-1.0.1e-48.el6_8.1.x86_64    parted-2.1-25.el6.x86_64
policycoreutils-2.0.83-19.47.el6.x86_64 python-2.6.6-66.el6_8.x86_64
python-setuptools-0.6.10-3.el6

glibc-2.12-1.166.el6_7.7.x86_64
libacl-2.2.49-6.el6.x86_64
libstdc++-4.4.7-11.el6.i686
```

```
ncurses-libs-5.7-3.20090208.el6.x86_64
perl-5.10.1-136.el6.x86_64
python-libs-2.6.6-66.el6_8.x86_64
readline-6.0-4.el6.x86_64
```

Required OS packages for RHEL6.6:

```
apr-devel 1.3.9-5.el6_2
bash 4.1.2-29.el6
glibc-devel.x86_64 2.12-1.166.el6_7.7
glibc-utils 2.12-1.166.el6_7.7
glibc.x86_64 2.12-1.166.el6_7.7
httpd 2.2.15-54.el6_8
httpd-manual 2.2.15-54.el6_8
infiniband-diags 1.6.4-1.el6
kernel-debuginfo =2.6.32-504.el6
kernel-firmware =2.6.32-504.el6
ksh 20120801-21.el6
libibverbs-devel 1.1.8-3.el6
libpcap 1.4.0
libyaml 0.1.3-4
memcached 1.4.4-3.el6
mod_wsgi 3.2-7.el6
net-snmp 5.5-49.el6_5.3
net-snmp-utils 5.5-49.el6_5.3
nscd 2.12-1.166.el6_7.7
openldap 2.4.39-8.el6
opensm 3.3.17-1.el6
openssh 5.3p1-111.el6
openssh-clients 5.3p1-111.el6
openssl 1.0.1e-48.el6_8.1
perl 5.10.1
python-argparse 1.2.1-2.1.el6
python-backports-ssl_match_hostname 3.4.0.2-2.el6
python-memcached 1.43-6.el6
python-requests 2.6.0-3.el6
python-urllib3 1.10.2-1.el6
rrdtool 1.3.8-7.el6
sg3_utils 1.28-6.el6
sysstat 9.0.4-27.el6
vsftpd 2.2.2-12.el6_5.1
apr-util-devel 1.3.9-3.el6_0.1
glibc-common 2.12-1.166.el6_7.7
glibc-headers 2.12-1.166.el6_7.7
glibc.i686 2.12-1.166.el6_7.7
hal 0.5.14-14.el6
httpd-devel 2.2.15-54.el6_8
httpd-tools 2.2.15-54.el6_8
iscsi-initiator-utils 6.2.0.873-13.el6
kernel-debuginfo-common-x86_64 =2.6.32-504.el6
kernel-headers =2.6.32-504.el6
libibumad 1.3.9-1.el6
libibverbs-utils 1.1.8-3.el6
libtirpc 0.2.1-10.el6
lsscsi 0.23
mod_ssl 2.2.15-54.el6_8
nc 1.84-22.el6
net-snmp-libs 5.5-49.el6_5.3
nfs-utils 1.2.3-54.el6
nss-pam-ldapd 0.7.5-18.2.el6_4
openldap-clients 2.4.39-8.el6
opensm-libs 3.3.17-1.el6
openssh-askpass 5.3p1-111.el6
openssh-server 5.3p1-111.el6
openssl-devel 1.0.1e-48.el6_8.1
perl-Net-Telnet 3.03
python-backports 1.0-5.el6
python-chardet 2.2.1-1.el6
python-paramiko 1.7.5-2.1.el6
python-six 1.9.0-2.el6
rdma 6.6_3.15-1.el6
scsi-target-utils 1.0.24-16.el6
sg3_utils-libs 1.28-6.el6
telnet 0.17-48.el6
wireshark 1.8.10-7.el6_5
```

Required operating system RPMs for RHEL 6.7

The RPM version numbers specified in these lists are the minimum required version numbers for these operating system RPMs.

Required OS lib rpms for RHEL6.7:

```
coreutils-8.4-37.el6.x86_64          ed-1.1-3.3.el6.x86_64
findutils-4.4.2-6.el6.x86_64        glibc-2.12-1.166.el6_7.7.i686
glibc-2.12-1.192.el6.i686           glibc-2.12-1.192.el6.x86_64
libgcc-4.4.7-16.el6.i686            libgcc-4.4.7-16.el6.x86_64
libstdc++-4.4.7-16.el6.x86_64      module-init-tools-3.9-25.el6.x86_64
openssl-1.0.1e-48.el6_8.1.x86_64   parted-2.1-29.el6.x86_64
policycoreutils-2.0.83-24.el6.x86_64 python-2.6.6-64.el6.x86_64
python-setuptools-0.6.10-3.el6
```

```
glibc-2.12-1.166.el6_7.7.x86_64
libacl-2.2.49-6.el6.x86_64
libstdc++-4.4.7-16.el6.i686
ncurses-libs-5.7-4.20090207.el6.x86_64
perl-5.10.1-141.el6.x86_64
python-libs-2.6.6-64.el6.x86_64
readline-6.0-4.el6.x86_64
```

Required OS packages for RHEL6.7:

```
apr-devel 1.3.9-5.el6_2              apr-util-devel 1.3.9-3.el6_0.1
bash 4.1.2-33.el6                   glibc-common 2.12-1.166.el6_7.7
glibc-devel.x86_64 2.12-1.166.el6_7.7      glibc-headers 2.12-1.166.el6_7.7
glibc-utils 2.12-1.166.el6_7.7      glibc.i686 2.12-1.166.el6_7.7
glibc.x86_64 2.12-1.166.el6_7.7     hal 0.5.14-14.el6
httpd 2.2.15-54.el6_8              httpd-devel 2.2.15-54.el6_8
httpd-manual 2.2.15-54.el6_8        httpd-tools 2.2.15-54.el6_8
infiniband-diags 1.6.4-1.el6         iscsi-initiator-utils 6.2.0.873-14.el6
kernel-debuginfo-2.6.32-573.el6     kernel-debuginfo-common-x86_64-2.6.32-573.el6
kernel-firmware-2.6.32-573.el6     kernel-headers-2.6.32-573.el6
ksh-20120801-33.el6                libibumad 1.3.9-1.el6
libibverbs-devel 1.1.8-4.el6        libibverbs-utils 1.1.8-4.el6
libpcap 1.4.0                       libtirpc 0.2.1-10.el6
libyaml 0.1.3-4                     lsscsi 0.23
memcached 1.4.4-3.el6               mod_ssl 2.2.15-54.el6_8
mod_wsgi 3.2-7.el6                  nc 1.84-24.el6
net-snmp 5.5-54.el6                 net-snmp-libs 5.5-54.el6
net-snmp-utils 5.5-54.el6          nfs-utils 1.2.3-64.el6
nscd 2.12-1.166.el6_7.7            nss-pam-ldapd 0.7.5-20.el6_6.3
openldap 2.4.40-5.el6              openldap-clients 2.4.40-5.el6
```

```

opensm 3.3.17-1.el6
openssl 1.0.1e-48.el6_8.1
perl 5.10.1
python-backports 1.0-5.el6
python-chardet 2.2.1-1.el6
python-paramiko 1.7.5-2.1.el6
python-six 1.9.0-2.el6
rdma 6.7_3.15-5.el6
scsi-target-utils 1.0.24-16.el6
sg3_utils-libs 1.28-8.el6
telnet 0.17-48.el6
wireshark 1.8.10-17.el6

opensm-libs 3.3.17-1.el6
openssl-devel 1.0.1e-48.el6_8.1
perl-Net-Telnet 3.03
python-backports-ssl_match_hostname 3.4.0.2-2.el6
python-memcached 1.43-6.el6
python-requests 2.6.0-3.el6
python-urllib3 1.10.2-1.el6
rrdtool 1.3.8-7.el6
sg3_utils 1.28-8.el6
sysstat 9.0.4-27.el6
vsftpd 2.2.2-14.el6

```

Required operating system RPMs for RHEL 6.8

The RPM version numbers specified in these lists are the minimum required version numbers for these operating system RPMs.

Required OS lib rpms for RHEL6.8:

```

coreutils-8.4-43.el6.x86_64
findutils-4.4.2-9.el6.x86_64
libacl-2.2.49-6.el6.x86_64
libstdc++-4.4.7-17.el6.x86_64
openssl-1.0.1e-48.el6_8.3.x86_64
policycoreutils-2.0.83-29.el6.x86_64
python-setuptools-0.6.10-3.el6
ed-1.1-3.3.el6.x86_64
glibc-2.12-1.192.el6.i686
libgcc-4.4.7-17.el6.x86_64
module-init-tools-3.9-25.el6.x86_64
parted-2.1-29.el6.x86_64
python-2.6.6-64.el6.x86_64

```

```

glibc-2.12-1.192.el6.x86_64
libstdc++-4.4.7-17.el6.i686
ncurses-libs-5.7-4.20090207.el6.x86_64
perl-5.10.1-141.el6_7.1.x86_64
python-libs-2.6.6-64.el6.x86_64
readline-6.0-4.el6.x86_64
zlib-1.2.3-29.el6.x86_64

```

Required OS packages for RHEL6.8:

```

apr-devel 1.3.9-5.el6_2
bash 4.1.2-40.el6
glibc-devel.x86_64 2.12-1.192.el6
glibc-utils 2.12-1.192.el6
glibc.x86_64 2.12-1.192.el6
httpd 2.2.15-54.el6_8
httpd-manual 2.2.15-54.el6_8
infiniband-diags 1.6.4-1.el6
apr-util-devel 1.3.9-3.el6_0.1
glibc-common 2.12-1.192.el6
glibc-headers 2.12-1.192.el6
glibc.i686 2.12-1.192.el6
hal 0.5.14-14.el6
httpd-devel 2.2.15-54.el6_8
httpd-tools 2.2.15-54.el6_8
iscsi-initiator-utils 6.2.0.873-21.el6

```

kernel-debuginfo-2.6.32-642.el6	kernel-debuginfo-common-x86_64-2.6.32-642.el6
kernel-firmware-2.6.32-642.el6	kernel-headers-2.6.32-642.el6
ksh 20120801-33.el6	libibumad 1.3.9-1.el6
libibverbs-devel 1.1.8-4.el6	libibverbs-utils 1.1.8-4.el6
libpcap 1.4.0	libtirpc 0.2.1-11.el6
libyaml 0.1.3-4	lsscsi 0.23
memcached 1.4.4-3.el6	mod_ssl 2.2.15-54.el6_8
mod_wsgi 3.2-7.el6	nc 1.84-24.el6
net-snmp 5.5-57.el6	net-snmp-libs 5.5-57.el6
net-snmp-utils 5.5-57.el6	nfs-utils 1.2.3-70.el6
nscd 2.12-1.192.el6	nss-pam-ldapd 0.7.5-32.el6
openldap 2.4.40-12.el6	openldap-clients 2.4.40-12.el6
opensm 3.3.17-1.el6	opensm-libs 3.3.17-1.el6
openssl 1.0.1e-48.el6_8.1	openssl-devel 1.0.1e-48.el6_8.1
perl 5.10.1	perl-Net-Telnet 3.03
python-backports 1.0-5.el6	python-backports-ssl_match_hostname 3.4.0.2-2.el6
python-chardet 2.2.1-1.el6	python-memcached 1.43-6.el6
python-paramiko 1.7.5-2.1.el6	python-requests 2.6.0-3.el6
python-six 1.9.0-2.el6	python-urllib3 1.10.2-1.el6
rdma 6.8_4.1-1.el6	rrdtool 1.3.8-7.el6
scsi-target-utils 1.0.24-18.el6	sg3_utils 1.28-12.el6
sg3_utils-libs 1.28-12.el6	sysstat 9.0.4-31.el6
telnet 0.17-48.el6	vsftpd 2.2.2-21.el6
wireshark 1.8.10-17.el6	

Required operating system RPMs for RHEL 7.3

The RPM version numbers specified in this list are the minimum required version numbers for this operating system RPM.

Required OS lib rpms for RHEL 7.3:

bc-1.06.95-13.el7.x86_64	coreutils-8.22-18.el7.x86_64
ed-1.9-4.el7.x86_64	findutils-4.5.11-5.el7.x86_64
glibc-2.17-157.el7.x86_64	libacl-2.2.51-12.el7.x86_64
libgcc-4.8.5-11.el7.x86_64	libstdc++-4.8.5-11.el7.x86_64
openssl-libs-1.0.1e-60.el7.x86_64	perl-Exporter-5.68-3.el7.noarch
perl-Socket-2.010-4.el7.x86_64	policycoreutils-2.5-8.el7.x86_64
python-2.7.5-48.el7.x86_64	python-libs-2.7.5-48.el7.x86_64
zlib-1.2.7-17.el7.x86_64	

Required OS packages for RHEL 7.3:

apr-devel 1.4.8-3
 apr-util-devel 1.5.2-6
 arptables 0.0.4-8

at 3.1.13-22
autogen-libopts 5.18-5
avahi-libs 0.6.31-17
bash 4.2.46-20
binutils 2.25.1-22
cairo 1.14.2-1
coreutils 8.22-18
cups-libs 1.6.3-26
ethtool 4.5-3
fuse 2.9.2-7
fuse-devel 2.9.2-7
fuse-libs 2.9.2-7
glibc-common 2.17.157
glibc-devel.x86_64 2.17.157
glibc-headers 2.17.157
glibc-utils 2.17.157
glibc.i686 2.17.157
glibc.x86_64 2.17.157
httpd 2.4.6-45
httpd-devel 2.4.6-45
httpd-manual 2.4.6-45
httpd-tools 2.4.6-45
infiniband-diags 1.6.5-3
initscripts 9.49.37-1
iproute 3.10.0-74
ipvsadm 1.27-7
iscsi-initiator-utils 6.2.0.873-35
jansson 2.10-1
kernel-debuginfo =3.10.0-514.el7
kernel-debuginfo-common-x86_64 =3.10.0-514.el7
kernel-headers =3.10.0-514.el7
kmod 20-9
krb5-devel 1.14.1-26
krb5-libs 1.14.1-26
krb5-workstation 1.15.1-8
ksh 20120801-26.el7
libibumad 1.3.10.2-1
libibverbs-utils 1.2.1-1
libjpeg-turbo 1.2.90-5
libpcap 1.5.3-8
libtirpc 0.2.4-0.8
libyaml 0.1.4-11
lshw B.02.17-12

lsof 4.87-4
lsscsi 0.27-4
memcached 1.4.15-10
mlocate 0.26-6
mod_ssl 2.4.6-45
mod_wsgi 3.4-12
net-snmp 5.7.2-24
net-snmp-utils 5.7.2-24
net-tools 2.0-0.17
nfs-utils 1.3.0-0.33
nmap-ncat 6.40-7
nscd 2.17-157
nss-pam-ldapd 0.8.13-8
ntp 4.2.6p5-25
ntpdate 4.2.6p5-25
openldap 2.4.40-13
openldap-clients 2.4.40-13
opensm 3.3.19-1
opensm-libs 3.3.19-1
openssl 1.0.1e-60
openssl-devel 1.0.1e-60
pango 1.36.8-2
perl 5.16.3
perl-Convert-ASN1 0.26-4
perl-JSON 2.59-2
perl-LDAP 0.56-5
perl-Net-Telnet 3.03-19.e17
perl-XML-Parser 2.41-10
psmisc 22.20-11
python-backports 1.0-8
python-backports-ssl_match_hostname 3.4.0.2-4
python-chardet 2.2.1-1
python-memcached 1.48-4
python-paramiko 1.7.7.1-3
python-requests 2.6.0-1
python-setuptools 0.9.8-4
python-six 1.9.0-2
python-urllib3 1.10.2-2
PyYAML 3.10-11
rdma 7.3_4.7_rc2-5
rrdtool 1.4.8-9
rsh 0.17-76
sg3_utils 1.37-9

```
sg3_utils-libs 1.37-9
strace 4.8-11
sysstat 10.1.5-11
targetcli 2.1.fb41-3
telnet 0.17-60
traceroute 2.0.22-2
tzdata-java
unzip 6.0-16
vim-enhanced 7.4.160
vsftpd 3.0.2-21
wireshark 1.10.14-10
yp-tools 2.14-3
ypbind 1.37.1-7
zip 3.0-11
```

Required operating system RPMs for RHEL 7.4

The RPM version numbers specified in this list are the minimum required version numbers for this operating system RPM.

Required OS packages for RHEL 7.4:

```
PyYAML 3.10-11
apr-devel 1.4.8-3
apr-util-devel 1.5.2-6
arptables 0.0.4-8
at 3.1.13-22
autogen-libopts 5.18-5
avahi-libs 0.6.31-17
bash 4.2.46-28
binutils 2.25.1-31
cairo 1.14.8-2
coreutils 8.22-18
cups-libs 1.6.3-29
ethtool 4.8-1
fuse 2.9.2-8
fuse-devel 2.9.2-8
fuse-libs 2.9.2-8
glibc-common 2.17.196
glibc-devel.x86_64 2.17.196
glibc-headers 2.17.196
glibc-utils 2.17.196
glibc.i686 2.17.196
glibc.x86_64 2.17.196
```

httpd 2.4.6-67
httpd-devel 2.4.6-67
httpd-manual 2.4.6-67
httpd-tools 2.4.6-67
infiniband-diags 1.6.7-1
initscripts 9.49.39-1
iproute 3.10.0-87
ipvsadm 1.27-7
iscsi-initiator-utils 6.2.0.874-4
jansson 2.10-1
kmod 20-15
krb5-devel 1.15.1-8
krb5-libs 1.15.1-8
krb5-workstation 1.15.1-8
libibumad 13-7
libibverbs-utils 13-7
libjpeg-turbo 1.2.90-5
libpcap 1.5.3-9
libtirpc 0.2.4-0.10
libyaml 0.1.4-11
lshw B.02.18-7
lsodf 4.87-4
lsscsi 0.27-6
memcached 1.4.15-10
mlocate 0.26-6
mod_ssl 2.4.6-67
mod_wsgi 3.4-12
net-snmp 5.7.2-28
net-snmp-utils 5.7.2-28
net-tools 2.0-0.22
nfs-utils 1.3.0-0.48
nmap-ncat 6.40-7
nscd 2.17-196
nss-pam-ldapd 0.8.13-8
ntp 4.2.6p5-25
ntpdate 4.2.6p5-25
openldap 2.4.44-5
openldap-clients 2.4.44-5
opensm 3.3.19-1
opensm-libs 3.3.19-1
openssl 1.0.2k-8
openssl-devel 1.0.2k-8
pango 1.40.4-1

```
perl 5.16.3
perl-Convert-ASN1 0.26-4
perl-JSON 2.59-2
perl-LDAP 0.56-5
perl-Net-Telnet 3.03-19.e17
perl-XML-Parser 2.41-10
psmisc 22.20-15
python-backports 1.0-8
python-backports-ssl_match_hostname 3.4.0.2-4
python-chardet 2.2.1-1
python-memcached 1.48-4
python-paramiko 1.7.7.1-3
python-requests 2.6.0-1
python-setuptools 0.9.8-7
python-six 1.9.0-2
python-urllib3 1.10.2-3
rrdtool 1.4.8-9
rsh 0.17-76
sg3_utils 1.37-12
sg3_utils-libs 1.37-12
strace 4.12-4
sysstat 10.1.5-12
targetcli 2.1.fb46-1
telnet 0.17-64
traceroute 2.0.22-2
tzdata-java
unzip 6.0-16
vim-enhanced 7.4.160
vsftpd 3.0.2-22
wireshark 1.10.14-14
yp-tools 2.14-5
ypbind 1.37.1-9
zip 3.0-11
```

Software requirements for installing Veritas Access in a VMware ESXi environment

Table 3-2 Software requirements for installing Veritas Access in a VMware ESXi environment

Item	Description
Operating system (OS)	Red Hat Enterprise Linux (RHEL) 7.4
VMware environment	VMware ESXi 5.5, 6.0 (certified versions)
IP address	Nine IPs are required for a two-node cluster: <ul style="list-style-type: none"> ■ Four IP addresses are used to configure physical IPs. ■ Four IP addresses are used to configure virtual IPs. ■ One IP address is used for the management console.

Hardware requirements for installing Veritas Access virtual machines

Table 3-3 Hardware requirements for installing Veritas Access virtual machines

Item	Description
CPU	1 CPU – 64 bit, dual, or quad core, 2.0 GHz or later
RAM	<ul style="list-style-type: none"> ■ 32 GB of RAM for physical servers ■ 60 GB (or more) RAM size internally available storage capacity for boot disk
Network interface card (NIC)	Four NIC cards <ul style="list-style-type: none"> ■ Two NIC cards for public network (minimum) ■ Two NIC cards for private network
Fibre Channel HBA	Two-port Fibre Channel HBAs are required if you want to use shared LUNs. If the environment has only DAS disks, then the HBA requirement is optional.

Management Server Web browser support

The following are the supported Web browsers for Veritas Access:

Table 3-4

Browser	Version	Comments
Internet Explorer	■ IE 10	JavaScript: Enabled Cookies: Enabled
Firefox	Firefox 4.x and later	JavaScript: Enabled Cookies: Enabled
Google Chrome	Google Chrome 10 and later version	JavaScript: Enabled Cookies: Enabled

Additional considerations for supported Web browsers:

- Your browser must support JavaScript 1.2 or later.
- If you use pop-up blockers (including Yahoo Toolbar or Google Toolbar), either disable them or configure them to accept pop-ups from the Veritas Access node to which you connect.
- For Internet Explorer 8.0 on Windows Server 2003, download and install the hot fix from the following location:
<http://support.microsoft.com/kb/938397/en-gb>
- If you are unable to download the gendeploy script using Internet Explorer 9.0, visit the following location to resolve the issue:
<http://support.microsoft.com/kb/2549423>
- For Internet Explorer, enable the play animations in web pages option in the multimedia category of Advanced Internet options.
- For Internet Explorer, when popup-blocker is turned on, make sure that the filter Level is set to Medium or lower.
- For Internet Explorer, ensure that the site is included in the list of trusted sites.
- If you cannot add the site to the list of trusted sites, enable the Binary and script Behaviors option in security settings.
- You must install Adobe Flash plug-in version 10, or later.

Supported NetBackup versions

Veritas Access supports NetBackup version 7.7.3 or later.

Supported OpenStack versions

The OpenStack drivers, Cinder and Manila, are supported on the RHEL 7 OS and the OpenStack Kilo, Mitaka, Newton, or Ocata releases.

The Cinder and Manila drivers were tested with the following:

- OpenStack Kilo, Mitaka, Newton, or Ocata versions from the DevStack repository
- OpenStack RDO

Note: The Manila driver works only with kernel NFS. It does not work with NFS-Ganesha.

Supported Oracle versions and host operating systems

Veritas Access supports Oracle using Direct NFS. Veritas Access Direct NFS supports only NFS protocol version 3.

Veritas Access supports Oracle single instance only. OracleRAC is not supported.

The following are the supported Oracle versions for Veritas Access:

- Oracle version 11gR2 (11.2.0.4 or later)
- Oracle 12c (12.1.0.1)

The following are the supported Oracle host operating systems in the order of importance for Veritas Access:

- Linux
- AIX
- Solaris
- HP-UX
- Oracle Linux

Supported IP version 6 Internet standard protocol

[Table 3-5](#) describes the IP version 6 (IPv6) Internet standard protocol.

Table 3-5 IPv6 Internet standard protocol

Description	Example format
Preferred form	ABCD:EF01:2345:6789:ABCD:EF01:2345:6789
Compressed form	FF01::101

Table 3-5 IPv6 Internet standard protocol (*continued*)

Description	Example format
Mixed form	0:0:0:0:FFFF:129.144.52.38

Network and firewall requirements

Table 3-6 displays the default ports that Veritas Access uses to transfer information.

Table 3-6 Default Veritas Access ports

Port	Protocol or Service	Purpose	Impact if blocked
21	FTP	Port where the FTP server listens for connections. Note: Users can configure another port if desired.	FTP features are blocked.
22	SSH	Secure access to the Veritas Access server	Veritas Access is not accessible.
25	SMTP	Sending SMTP messages.	The SMTP messages that are sent from Veritas Access are blocked.
53	DNS queries	Communication with the DNS server	Domain name mapping fails.
111	rpcbind	RPC portmapper services	RPC services fail.
123	NTP	Communication with the NTP server	Server clocks are not synchronized across the cluster. NTP-reliant features (such as DAR) are not available.
139	CIFS	CIFS client to server communication	CIFS clients cannot access the Veritas Access cluster

Table 3-6 Default Veritas Access ports (*continued*)

Port	Protocol or Service	Purpose	Impact if blocked
161	SNMP	Sending SNMP alerts	SNMP alerts cannot be broadcast.
445	CIFS	CIFS client to server communication	CIFS clients cannot access the Veritas Access cluster.
514	syslog	Logging program messages	Syslog messages are not recorded.
756, 757, 755	statd	NFS statd port	NFS v3 protocol cannot function correctly.
2049	NFS	NFS client to server communication	NFS clients cannot access the Veritas Access cluster.
3172, 3173	ServerView	ServerView port	ServerView cannot work.
3260	iSCSI	iSCSI target and initiator communication	Initiator cannot communicate with the target.
4001	mountd	NFS mount protocol	NFS clients cannot mount file systems in the Veritas Access cluster.
4045	lockd	Processes the lock requests	File locking services are not available.
5634	HTTPS	Management Server connectivity	Web GUI may not be accessible.
56987	Replication	File synchronization, Veritas Access replication	Veritas Access replication daemon is blocked. Replication cannot work.
8088	REST server	REST client to server communication	REST client cannot access REST API of Veritas Access.

Table 3-6 Default Veritas Access ports (*continued*)

Port	Protocol or Service	Purpose	Impact if blocked
8143	S3	Data port for Veritas Access S3 server	User will not be able to use Veritas Access object server.
8144	ObjectAccess service	Administration port for Veritas Access S3 server.	User cannot create access or secret keys for using Objectaccess service.
11211	Memcached port	CLISH framework	CLISH cannot function correctly, and cluster configuration may get corrupted.
30000:40000	FTP	FTP passive port	FTP passive mode fails.
14161	HTTPS	Access Veritas Access GUI	User is unable to access Veritas Access GUI
51001	UDP	LLT over RDMA	LLT is not working.
51002	UDP	LLT over RDMA	LLT is not working.

NetBackup ports

NetBackup uses TCP/IP connections to communicate between one or more TCP/IP ports. Depending on the type of operation and configuration on the environment, different ports are required to enable the connections. NetBackup has different requirements for operations such as backup, restore, and administration.

[Table 3-7](#) shows some of the most-common TCP and UDP ports that Veritas Access NetBackup uses to transfer information. For more information, see the *Veritas NetBackup Security and Encryption Guide*.

Table 3-7 Default NetBackup TCP and UDP ports

Port Range	Protocol
1556	TCP, UDP
13701-13702, 13705-13706	TCP

Table 3-7 Default NetBackup TCP and UDP ports (*continued*)

Port Range	Protocol
13711, 13713, 13715-13717, 13719	TCP
13720-13722	TCP, UDP
13723	TCP
13724	TCP, UDP
13782-13783	TCP, UDP
13785	TCP

OpenDedup ports and disabling the iptable rules

This use case is specific to running OpenDedup on Veritas Access. Each time a SDFS volume is created and mounted on Veritas Access, it starts listening on a specific port. Initially, it starts with port 6442 and goes on incrementing +1 for further subsequent volumes.

Table 3-8 OpenDedup ports

Port Range	Protocol or Service	Purpose	Impact if Blocked
Starts from 6442 and increments +1 for subsequent volumes	TCP	Allows communication between Veritas Access and OpenDedup	Veritas Access cannot communicate with OpenDedup

To allow communication to the OpenDedup port running on Veritas Access, disable the iptable rules completely

- 1 Use the `df` command to show that the SDFS volume is mounted and on which port it is listening.

The SDFS volume is already mounted as part of the LTR script.

```
[root@ltrclust_02 ~]# df -h | tail -2
sdfs:/etc/sdfs/pool100-volume-cfg.xml:6442
11G    0    11G    0% /pool100
```

- 2 Use the `netstat` command to verify that the port is open.

```
[root@ltrclust_02 ~]# netstat -tulpn | grep 6442
tcp        0      0  :::6442    :::*       LISTEN
3761/jsvc.exec
```

- 3 Disable the `iptables` rules to allow communication to the OpenDedup port once the volume is mounted and to disallow traffic to this port once the volume is unmounted.

Use the following commands to disable the `iptables` rules:

```
[root@ltrclust_02 ~]# iptables -F

[root@ltrclust_02 ~]# /etc/init.d/iptables stop

[root@ltrclust_02 ~]# iptables -L
```

Use the `iptables -L` command to verify that all the `iptables` rules are disabled.

The `iptables` rules should be run on all the Veritas Access cluster nodes and on the NetBackup media server if OpenDedup is installed on it.

- 4 An alternative to disabling the `iptables` rules in Step 3 is to add an `iptables` rule to open the OpenDedup port, so that the existing `iptables` rules are also used.

Example:

```
[root@ltrclust_02 ~]# iptables -A INPUT -p tcp --dport 6442 -j ACCEPT
```

CIFS protocols and firewall ports

For the CIFS service to work properly in an Active Directory (AD) domain environment, the following protocols and firewall ports need be allowed or opened

to enable the CIFS server to communicate smoothly with Active Directory Domain Controllers and Windows/CIFS clients.

Internet Control Message Protocol (ICMP) protocol must be allowed through the firewall from the CIFS server to the domain controllers. Enable "Allow incoming echo request" is required for running the CIFS service.

[Table 3-9](#) lists additional CIFS ports and protocols.

Table 3-9 Additional CIFS ports and protocols

Port	Protocol	Purpose
53	TCP, UDP	DNS
88	TCP, UDP	Kerberos
139	TCP	DFSN, NetBIOS Session Service, NetLog
445	TCP, UDP	SMB, CIFS, SMB2, DFSN, LSARPC, NbtSS, NetLogonR, SamR, SrvSvc
464	TCP, UDP	Kerberos change or set a password
3268	TCP	LDAP GC
4379	TCP	CTDB in CIFS

[Table 3-10](#) lists the ports that are required for LDAP with SSL.

Table 3-10 LDAP with SSL ports

Port	Protocol	Purpose
636	TCP	LDAP SSL
3269	TCP	LDAP GC SSL

Maximum configuration limits

The maximum configuration limits for configuring the Veritas Access system software are as follows:

Table 3-11 Maximum configuration limits

Veritas Access system software	Configuration limit
File system size	512 TB for non-scale-out file system 3 PB for scale-out file system
Veritas Access nodes	20
Supported LUNs	The maximum number of disks is theoretically limited to the number that can be attached to the operating system. However, it has only been tested in the thousands.
Supported file systems	500
Tiers within a file system	2 (primary tier and secondary tier)

Preparing to install Veritas Access

This chapter includes the following topics:


- [Overview of the installation process](#)
- [Hardware requirements for the nodes](#)
- [About using LLT over the RDMA network for Veritas Access](#)
- [Connecting the network hardware](#)
- [About obtaining IP addresses](#)
- [About checking the storage configuration](#)

Overview of the installation process

The Veritas Access cluster is a set of connected servers called "nodes." Together these nodes form a unified entity called a cluster.

[Figure 4-1](#) shows an example of an Veritas Access cluster.

Figure 4-1 Sample of Veritas Access cluster overview



Note: The NIC names mentioned in [Figure 4-1](#) are only for examples. You need to determine the actual names of your NICs during the installation.

An overview of the Veritas Access software installation includes the following steps:

- Gather network information from your network administrator.
- Connect your network hardware.
- Install the operating system on each of the nodes.
- Install Veritas Access on the node. If the driver node is one of the nodes of the cluster, you must start the installer from the console of the node. If the driver node is not part of the cluster, the installer can be run from the driver node to install and configure the cluster over an ssh connection.

From the Veritas Access 7.2 release, the installer can be run from any node of the cluster.

See [“Installing and configuring the Veritas Access software on the cluster”](#) on page 79.

See [“About the driver node”](#) on page 72.

- Run the installation and configuration on the node to configure the entire cluster. Installation times vary depending on your configuration.

Hardware requirements for the nodes

The following table summarizes the hardware requirements for each node.

Table 4-1 Hardware requirements for the nodes

Item	Requirements
Network interface card (NIC)	<p>At least four NICs are required for each node.</p> <p>Two NICs connected to a private network.</p> <ul style="list-style-type: none"> ■ For a two-node cluster, either cross connect two private NICs on each node or use a switch. ■ If there are more than two nodes in the cluster, make sure that you have a dedicated switch (or a public or private switch with a dedicated VLAN) and that all the private NICs are connected to the switch. <p>Connect two public NICs from each node to the public network. The gateway must be reachable to each public NIC.</p>
IP address	<p>For a two-node cluster, make sure that you have nine IP addresses available.</p> <ul style="list-style-type: none"> ■ Four IP addresses are used to configure physical IPs. ■ Four IP addresses are used to configure virtual IPs. ■ One IP address is used to configure the Operations Manager console. ■ One IP address is used for replication, which is optional. <p>Make sure that these nine IP addresses are different from the IP addresses that are already assigned to the target cluster nodes to install Veritas Access over Secure Shell (ssh).</p>

About using LLT over the RDMA network for Veritas Access

Remote direct memory access (RDMA) allows server-to-server data movement directly between application memories with minimal CPU involvement. RDMA provides fast interconnection between user-space applications or file systems between nodes over InfiniBand networks with RDMA-enabled network cards and

switches. In a clustering environment, RDMA allows applications on separate nodes to transfer data at a faster rate with low latency and less CPU usage.

RDMA over InfiniBand networks in the Veritas Access clustering environment

Veritas Access uses Low Latency Transport (LLT) for data transfer between applications on nodes. LLT functions as a high-performance, low-latency replacement for the IP stack, and is used for all cluster communications. It distributes (load balances) internode communication across all available private network links. This distribution means that all cluster communications are evenly distributed across all private network links (maximum eight) for performance and fault resilience. If a link fails, traffic is redirected to the remaining links. LLT is also responsible for sending and receiving heartbeat traffic over network links. Using LLT data transfer over an RDMA network boosts performance of both file system data transfer and I/O transfer between nodes.

Network interface cards (NICs) and network switches that support RDMA are required to enable the faster application data transfer between nodes. You also need to configure the operating system and LLT for RDMA.

See [“Configuring LLT over RDMA for Veritas Access”](#) on page 53.

How LLT supports RDMA for faster interconnections between applications

Low Latency Transport (LLT) maintains two channels (RDMA and non-RDMA) for each of the configured RDMA links. Both RDMA and non-RDMA channels can transfer data between the nodes. LLT provides separate Application Program Interfaces (APIs) to the clients (such as CFS and CVM) to use these channels. The RDMA channel is mainly used for data transfer by the client; while the non-RDMA channel is created over the UDP layer, and LLT uses it mainly for sending and receiving heartbeats. Group Membership Services/Atomic Broadcast (GAB) decides cluster membership for the cluster according to the health of the non-RDMA channel. The connections of the RDMA and non-RDMA channels are under separate management, while the connect and disconnect operations for the RDMA channel are triggered based on the status of the non-RDMA channel.

If the non-RDMA channel is up while the RDMA channel is down, the data is transferred over the non-RDMA channel with lower performance until the RDMA channel is fixed. The system logs display a message when the RDMA channel is up or down.

LLT uses the Open Fabrics Enterprise Distribution (OFED) layer and the drivers on the operating system to communicate with the hardware. LLT over RDMA allows

applications running on one node to directly access the memory of an application running on another node over an RDMA-enabled network. While over a non-RDMA network, LLT clients have to create intermediate data copies to complete the read or write operation on the application. The RDMA network brings low latency, higher throughput, and minimized CPU host usage, and boosts application performance. LLT and GAB clients CFS and CVM can use LLT over RDMA.

Configuring LLT over RDMA for Veritas Access

During the Veritas Access installation, the installer automatically configures LLT over RDMA if there are InfiniBand NICs on the cluster nodes, unless the InfiniBand NICs are excluded.

This section describes the required hardware and configuration for LLT to support RDMA for Veritas Access. The high-level steps to configure LLT over RDMA are as follows:

1. Choose NICs, network switches, and cables that support RDMA.

Table 4-2 RDMA-enabled hardware

Hardware	Supported types	Reference
Network card	Mellanox-based Host Channel Adapters (HCAs) (VPI, ConnectX, ConnectX-2 and 3)	For detailed installation information, refer to the hardware vendor documentation.
Network switch	Mellanox, InfiniBand switches Ethernet switches must be Data Center Bridging (DCB) capable	For detailed installation information, refer to the hardware vendor documentation.
Cables	Copper and Optical Cables, InfiniBand cables	For detailed installation information, refer to the hardware vendor documentation.

2. Connect the first two non-excluded InfiniBand NICs as private NICs.

Note: Cross-links connection is not supported for private NICs in an RDMA environment.

3. Make sure that the required packages to enable RDMA, InfiniBand drivers, and utilities are installed with the base operating system. Or they can be installed from the yum repository.

Table 4-3 Drivers and utilities required for RDMA, InfiniBand, or an Ethernet network

Packages	Drivers and utilities
Device drivers for RDMA operations	<ul style="list-style-type: none"> ■ libmthca ■ libmlx4 ■ rdma ■ librdmacm-utils
OpenSM-related package	<ul style="list-style-type: none"> ■ opensm ■ opensm-libs ■ libibumad
InfiniBand troubleshooting and performance tests	<ul style="list-style-type: none"> ■ ibutils ■ infiniband-diags ■ perftest
libibverbs packages for InfiniBand operations	<ul style="list-style-type: none"> ■ libibverbs-devel ■ libibverbs-utils

How the Veritas Access installer configures LLT over RDMA

At a high level, the Veritas Access installer configures the InfiniBand NICs as LLT over RDMA for Veritas Access by the following steps:

- 1** After the InfiniBand NICs are detected, the installer installs the required operating system packages.
- 2** Choose InfiniBand NICs as private NICs, if the NIC is not excluded.
- 3** Assign static private IPs and configure LLT to use InfiniBand NICs.

LLT over RDMA sample /etc/llttab

The following is a sample of LLT over RDMA in the `etc/llttab` file.

```
rdma_01:~ # cat /etc/llttab
set-node rdma_01
set-cluster 54791
link priveth0 udp - rdma 51001 - 172.16.0.3 172.16.0.255
link priveth1 udp - rdma 51002 - 172.16.1.3 172.16.1.255
set-flow highwater:1000
set-flow lowwater:800
```

Connecting the network hardware

Before you install the Veritas Access software, you must assemble a cluster by configuring all the nodes with the required network hardware, and connecting the Ethernet interfaces to the private and the public networks.

To assemble the cluster, do the following:

- Determine a preferred location for the cluster.
- Make sure that each node has at least two redundant Ethernet interfaces (gigabit Ethernet) to connect to a private network for cluster internal control.
- Make sure that each node has at least two additional Ethernet interfaces (gigabit Ethernet) to connect to the public network. You can use the public Ethernet interfaces from the embedded interfaces on the motherboard or from the add-on (PCI) network adapter interfaces.
- To connect the public NICs, connect one end of the Ethernet cables to the Ethernet interfaces on the back of the nodes. Connect the other end of the Ethernet cables to your corporate network so that they can reach the gateway. At least two public interfaces are required for each node.
- To connect the private NICs, use the first two available NICs when sorted by NIC name. Available NICs are those not connected to the public network or excluded from the node.

For example, if your NICs are eth1, eth2, eth3, and eth4, and none of the NICs are connected to the public network or excluded, then use eth1 and eth2 as the private NICs.

Connect one end of the Ethernet cables to Ethernet interface 1 and 2 on the back of the nodes. For a 2-node cluster, connect the other end of the Ethernet cables to the corresponding Ethernet interfaces on the second node. For a cluster with more than 2 nodes, connect the other end of the Ethernet cables to a dedicated switch or VLAN.

Note: It is recommended to use InfiniBand NICs to configure LLT over RDMA for Veritas Access. Connect InfiniBand NICs as private or exclude the NICs when you install Veritas Access.

See [“About using LLT over the RDMA network for Veritas Access”](#) on page 51.

See [“Excluding a NIC”](#) on page 86.


- Ask your network administrator for a range of IP addresses to use in the Veritas Access installation. The number of IP addresses you need depends on the number of nodes and number of network interface cards in your cluster.

You need at least one IP address per node per public interface. For virtual IP addresses, you can configure the virtual IP addresses later in the CLISH if you input 0 for the number of virtual IP addresses per NIC during installation time. Veritas Access supports both Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6), but they cannot be mixed.

Physical IP address	An IP address that is associated with a specific Ethernet interface address and cannot automatically be failed over.
Virtual IP address (VIP)	An IP address whose association to a specific Ethernet interface (VIP) can be failed over to other interfaces on other nodes by the Veritas Access software.
Console IP address	A dedicated virtual IP address that is used to communicate with the Veritas Access cluster Management Console. This virtual IP address is assigned to the master node. If the master node fails, the Veritas Access software automatically selects a new master node from the cluster and fails the console IP address over to it.

Figure 4-2 shows a diagram of a four-node cluster.

Figure 4-2 Private network configurations: four-node cluster



Note: Two or more Veritas Access private networks cannot be configured on the same IPv4 network.

About obtaining IP addresses

The Veritas Access installation process lets you configure IP addresses for 1 to 20 nodes. The default is two nodes.

Note: You can configure either IPv4 addresses or IPv6 addresses (depending on what you use when installing Veritas Access), but not both. Do not use IP addresses starting with 172.16.X.X either as physical IP addresses or virtual IP addresses since this range of IP addresses are used for the private network.

You need to obtain a contiguous range of physical IP addresses, a contiguous range of virtual IP addresses, and a netmask for the chosen public network from the network administrator in charge of the facility where the cluster is located. All IP addresses (both physical and virtual) must be part of the same subnet and use the same netmask as the node's access IP.

By design, the installer does not support the use of the localhost (127.0.0.1) IP address during installation

Note: Netmask is used for IPv4 addresses. Prefix is used for IPv6 addresses. Accepted ranges for prefixes are 0-128 (integers) for IPv6 addresses.

The information you obtained from the network administrator is used to configure the following:

- Physical IP addresses
- Virtual IP addresses
- Console IP address
- Replication IP address (optional)
- IP address for the default gateway
- IP address for the Domain Name System (DNS) server
- DNS domain name
- IP address for the Network Time Protocol (NTP) server (optional)
- Virtual IP address for Veritas NetBackup (optional)

About calculating IP address requirements

This section provides an example of how to calculate IP addresses for a two-node cluster. In this example, all the nodes in the cluster have the same hardware configuration. Therefore, the number of network interface cards (NICs) is the same for all the nodes in the cluster.

- Two private NICs and two public NICs should be connected to respective networks.
- One public IP address should be assigned to one of the public interface for installation over ssh. None of the private interfaces should have the IP address in the same network segment.
- The public IP address must be made permanent by writing it to the network configuration file `/etc/sysconfig/network-scripts/ifcfg-ethX`.

Table 4-4 Example calculation of required IPs for a standard configuration

Number of IPs	Item
2	Number of nodes in the cluster
4	Number of interfaces on each node
2	Number of the private interfaces that are required for each node

After two private interfaces on each node are selected, all remaining interfaces act as public interfaces.

To calculate the number of public interfaces per node

- ◆ The total number of interfaces on the node, minus the number of private interfaces that are required on a node, is equal to the remaining number of public interfaces on the node.

```
Total number of interfaces (4)
- Number of private interfaces (2)
= Number of public interfaces
```

$$4 - 2 = 2$$

To calculate the physical and the virtual IP addresses for the cluster

- 1** The total number of physical IP addresses that are required for the cluster installation is equal to the number of nodes in the cluster multiplied by the number of public interfaces on each node:

```
Total number of nodes (2)
x Number of public interfaces per node (2)
= Total number of physical IP addresses

= 2 x 2 = 4
```

- 2** The number of nodes in the cluster multiplied by the number of public interfaces on each node is equal to the total number of virtual IP addresses that are required for the cluster installation:

```
Total number of nodes (2)
x Number of public interfaces per node (2)
= Total number of virtual IP addresses

= 2 x 2 = 4
```

- 3** The number of IP addresses required for the Veritas Access Operations Manager is equal to one (1).

To calculate the total number of public IP addresses for the cluster

- ◆ The number of physical IP addresses for the cluster, plus the number of virtual IP addresses for the cluster, plus the number of IP addresses for the Operations Manager is equal to the total number of public IP addresses that are required for the cluster.

```
Total number of physical IP addresses/cluster (4)
+ Total number of virtual IP addresses/cluster (4)
+ Number of IP addresses for the Management Console (1)
= Total number of public IP addresses required for the cluster

= 4 + 4 + 1 = 9
```

To request and specify IP addresses

- 1 Request the public IP addresses that you need from your Network Administrator.
- 2 For example, if the Network Administrator provides you with IP addresses 10.209.105.120 through 10.209.105.128, you can allocate the resources in the following manner:

```
Start of Physical IP address: 10.209.105.120
Start of Virtual IP address: 10.209.105.124
Management Console IP:"10.209.105.128"
```

This entry gives you four physical IP addresses (10.209.105.120 to 10.209.105.123), four virtual IP addresses (10.209.105.124 to 10.209.105.127), and one IP address for the Operations Manager (10.209.105.128).

10.209.105.120 and 10.209.105.121 are assigned to pubeth0 and pubeth1 as physical IP addresses on the first node.

10.209.105.122 and 10.209.105.123 are assigned to pubeth0 and pubeth1 as physical IP addresses on the second node.

10.209.105.124 to 10.209.105.127 are assigned to pubeth0 and pubeth1 as virtual IP addresses on the two nodes.

Reducing the number of IP addresses required at installation time

You can reduce the number of IP addresses required at installation time by not configuring any virtual IP addresses. During the Veritas Access installation, input 0 for the number of virtual IP addresses per NIC.

Virtual IP addresses are not required at installation time. You can configure the virtual IP addresses later using the `Network> ip addr add` command in the CLISH.

See the `network(1)` manual page for more information on adding NICs.

You need at least one IP address per node per public interface at installation time.

Table 4-5 Example configuration of required IP addresses at installation time for a two-node cluster with two public NICs per node

Number of IP addresses	Item
4	Number of physical IP addresses. The four IP addresses include the original physical IP addresses.
1	One IP address for the management console.

About checking the storage configuration

Warning: Do not connect the Fibre Channel HBAs until you finish installing the operating system. If the local disks are bad, connecting the Fibre Channel HBAs prevents the operating system from being installed on the local disks. Because the disk is scanned, it takes longer to install the software on a local disk.

Veritas Access supports Flexible Storage Sharing (FSS), which allows the users to configure and manage direct-attached storage on the Veritas Access appliance. After you install the operating system, check the storage configuration. If you don't want to use FSS, make sure that each node has the following:

- One or two Fibre Channel Host Bus Adapters (HBAs) for connection to the Storage Area Network (SAN) switch.
Two Fibre Channel HBAs are recommended, but only one is required. Having only one Fibre Channel HBA enables all the operations of the Fibre Channel (except high availability).
- An internal boot disk. Make sure that one is in place before you install the Veritas Access software.

If you want to use FSS, make sure that each node has attached at least two extra local data disks besides the internal boot disk.

Deploying virtual machines in VMware ESXi for Veritas Access installation

This chapter includes the following topics:

- [Setting up networking in VMware ESXi](#)
- [Creating a datastore for the boot disk and LUNs](#)
- [Creating a virtual machine for Veritas Access installation](#)

Setting up networking in VMware ESXi

Before you start, install the ESXi server. You can deploy the first virtual machine on your ESXi host by using the vSphere Client.

To set up a network in VMware ESXi

- 1 Start the vSphere Client and type the logon details for your host.
In the **IP address / Hostname** text box, enter the **ESXi server IP/hostname**.
In the **User name** text box, type **root**.
In the **Password** text box, type **my_esxi_password**.
- 2 Set up the networking requirements for Veritas Access.
- 3 To set up the public network virtual switch:
 - In the **Configuration** tab of the ESXi host, navigate to **Hardware > Networking**.
 - Click **Add Networking** on the top right corner.

- Select the connection type as **Virtual Machine** and click **Next**.
 - Select the NIC that is connected to the public network under the **Create a virtual switch** section.
 - Enter the appropriate network label for the public virtual switch.
 - Verify the summary and click **Finish** to create the public network virtual switch.
 - Repeat the steps for creating multiple public network switches.
- 4 To set up the private network virtual switch:
- In the **Configuration** tab of the ESXi host, navigate to **Hardware > Networking**.
 - Click **Add Networking** on the top right corner.
 - Select the connection type as **Virtual Machine** and click **Next**.
 - Deselect any NIC that is selected by default for creating the virtual switch.
 - Enter the appropriate network label for the private virtual switch.
 - Verify that the summary shows no-adapters under the physical adapters, and click **Finish** to create the first private network virtual switch.
 - Repeat the steps to create the second private network virtual switch.

Creating a datastore for the boot disk and LUNs

To create a datastore for the boot disk and LUNs

- 1 Create a datastore for vmrk files for virtual machines.
- 2
 - In the **Configuration** tab of the ESx host, navigate to **Hardware > Storage**.
 - Click **Add Storage** on the top right corner.
 - Select the storage type as **Disk/LUN** and click **Next**.
 - Select the disk that you want to use to create the virtual machine vmrk files.
 - Review the current disk layout and click **Next**.
 - Enter the datastore name of your choice and click **Next**.
 - Select the disk space that you want to dedicate for the datastore. The default option is to use the complete list.
 - Review the details and click **Finish**.

Creating a virtual machine for Veritas Access installation

To create a virtual machine for Veritas Access installation

- 1 After the networking configuration is complete and the datastore is defined, create the virtual machines.
 - Select the **ESXi host IP/hostname** in the top of the tree structure in the leftmost frame.
 - From the file menu, select **New Virtual Machine**, which opens a pop-up for creating the virtual machine.
 - Select the configuration as **Custom** and click **Next** to decide on the exact configuration of the virtual machine.
 - Enter the virtual machine name of your choice and click **Next**.
 - Select the datastore that stores the virtual machine vmdk file and click **Next**.
 - Select the virtual machine version that you want to use and click **Next**. Veritas recommends version 8.
 - Select the guest operating system as **Linux** and version as **Red Hat Enterprise Linux 6 or 7 (64-bit)** and click **Next**.

Select the number of CPUs. Veritas recommends eight cores that can be:

- Two virtual sockets and four cores per virtual socket.
 - One virtual socket and eight cores per virtual socket.
 - Any higher number of cores as per your workload.
- Select the memory configuration. Veritas recommends 32 GB.
- In the network configuration, select the number of NICs as four.
 - For NIC1, select the public network virtual switch and validate that the adapter is correct.
 - For NIC2, select the public network virtual switch and validate that the adapter is correct.
 - For NIC3, select the private network virtual switch 1 and validate that the adapter is correct.
 - For NIC4, select the private network virtual switch 2 and validate that the adapter is correct.
 - Select the SCSI controller as **VMware Paravirtual**.
 - In the disk configuration page, select **Create a new virtual disk** and click **Next**.

- Select the boot disk size. Veritas recommends 100 GB.
 - Select the disk provisioning type as **Thick Provision Eager Zeroed**.
 - Select the datastore as **Specify a data store or data store cluster** and click **Next**.
After selecting the datastore, click **Next**.
 - Select the **Virtual device node** as default (SCSI (0:0) for the boot disk) and click **Next**.
 - Review the virtual machine configuration and click **Finish** to create the virtual machine.
The virtual machine creation task is complete.
- 2** Select the virtual machine and click **Edit virtual machine settings** to validate the following:
- There should be four network adapters - two for the public network and two for the private network.
 - Verify that the memory and CPU configuration is correct.
- 3** Repeat Step 1 and Step 2 to create the second virtual machine, which is used to form the two-node Veritas Access cluster.
- 4** Add LUNs/DAS disks to the virtual machines.
- To add local DAS disks:
- Select the virtual machine and click **Edit virtual machine settings**.
 - Click the **Add** button.
 - Select the device type as **Hard Disk** and click **Next**.
 - Select **Create a new virtual disk** in the type of disk and click **Next**.
 - Select the DAS disk size. Veritas recommends 100 GB.
 - Select the disk provisioning type as **Thick Provision Eager Zeroed**.
 - Select the datastore as **Specify a data store or data store cluster** and click **Next**.
 - Select the **Virtual device node** as SCSI (1:0) for the first SAS disk and click **Next**.
- Once all the required DAS disk creation is complete, complete the following:
- Select the SCSI controller 1, which is used for DAS disks.
 - Set the SCSI Bus sharing mode as **Virtual**.
This mode is required so that DAS disks are claimed in VxVM enclosure-based naming (EBN) mode and host name is only prefixed

by VxVM when disks are in EBN mode, which distinguishes it from the shared LUNs present in the arrays.

- Click **OK** to create the DAS disk.
Repeat this step for creating the DAS disk for other Veritas Access nodes.

5 Map the shared disks to the LUNs. Mapping of LUNs from an array is only supported using Raw Device Mapping (RDM) mode.

Mapping shared LUNs to the first virtual machine:

- Select the first virtual machine and click **Edit virtual machine settings**.
- Click the **Add** button.
- Select the device type as **Hard Disk** and click **Next**.
- Select the LUN that you want to map and click **Next**.
- Select the datastore that stores the LUN mapping or select **Store with virtual machine**.
- Select the compatibility mode as **Physical** to access the array LUN hardware directly.
- Select the **Virtual device node** as SCSI (2:0) for the shared disk and click **Next**.
- Review the mapping of the disk and click **Finish** to map the array LUN disk to the virtual machine.
Repeat this Step for the number of LUNs that you want to map and update the **Virtual device node** to the next free SCSI controller port.

Once all the required LUNs are mapped, complete the following:

- Select the SCSI controller 2, which is used for shared LUNs.
- Set the SCSI Bus sharing mode as **Virtual**.
This mode is required so that the shared LUNs are claimed in VxVM enclosure-based naming (EBN) mode. This distinguishes it from the shared LUNs present in the arrays.
- Click **OK** to complete the mapping of LUNs in RDM mode.

Mapping shared LUNs to the second virtual machine:

- Select the first virtual machine and click **Edit virtual machine settings**.
- Click the **Add** button.
- Select the device type as **Hard Disk** and click **Next**.
- Select **Use an existing Virtual Disk** in the type of disk and click **Next**.

- Navigate to the corresponding disk path in the datastore where the shared disk was stored when they were mapped to the first virtual machine.
- Select the **Virtual device node** as SCSI (2:0) for the shared disk and click **Next**. Ensure that the sequence of disk mapping is the same as that of the first virtual machine and mapping has been done to the same SCSI controller to achieve a shared disk configuration.
- Review the mapping of the disk and click **Finish** to map the array LUN disk to the virtual machine.
Repeat this Step for the number of shared LUNs that you have mapped to other virtual machines and update the **Virtual device node** to the next free SCSI controller port.

Once all the required LUNs are mapped, complete the following:

- Select the SCSI controller 2, which is used for the shared LUNs.
 - Set the SCSI Bus sharing mode as **Virtual**.
This mode is required so that the shared LUNs are claimed in VxVM enclosure-based naming (EBN) mode. This distinguishes it from the shared LUNs present in the arrays.
 - Click **OK** to complete the mapping of LUNs in RDM mode.
The networking and storage configuration is complete for the virtual machines.
- 6** Install the RHEL 6 Update 6, 7, or 8 or RHEL 7 Update 3 or 4 (64-bit) operating system that is supported by the Veritas Access installer.

See [“Installing the operating system on the target Veritas Access cluster”](#) on page 73.

Installing and configuring a cluster

This chapter includes the following topics:

- [Installation overview](#)
- [Summary of the installation steps](#)
- [Before you install](#)
- [Installing the operating system on each node of the cluster](#)
- [Installing Veritas Access on the target cluster nodes](#)
- [About NIC bonding and NIC exclusion](#)
- [About VLAN Tagging](#)
- [Replacing an Ethernet interface card](#)
- [Configuring I/O fencing](#)
- [About configuring Veritas NetBackup](#)
- [About enabling kdump during an Veritas Access configuration](#)
- [Reconfiguring the Veritas Access cluster name and network](#)
- [Configuring a KMS server on the Veritas Access cluster](#)

Installation overview

Initially, you can install a two-node Veritas Access cluster. You can increase the cluster by adding nodes up to the maximum of 20 nodes. The recommended minimum value is two nodes. Adding nodes to the cluster does not disrupt service.

Summary of the installation steps

The Veritas Access software installation consists of two main pieces:

- Operating system installation.
Veritas Access requires Red Hat Enterprise Linux.
See See [“System requirements”](#) on page 16.
- Veritas Access software installation.

[Table 6-1](#) provides a brief summary of the installation steps. The summary includes cross references to where you can find more information about each task.

Table 6-1 Summary of installation steps

Task	Steps	For more information
Task 1: Install the operating system on each node of the cluster.	Steps include: <ul style="list-style-type: none">■ Automatic system discovery of USB devices, hard disk controllers, and so on.■ Select the installation device.■ Set the clock and the time zone.■ System preparation for automated installation.■ Automatic disk partitioning.■ Automatic package installation.■ Install the Red Hat Enterprise Linux kernel update.	See “Installing and configuring the Veritas Access software on the cluster” on page 79.

Table 6-1 Summary of installation steps (*continued*)

Task	Steps	For more information
Task 2: Install the Veritas Access software on the cluster.	<p>Steps include:</p> <ul style="list-style-type: none">■ Install the required Red Hat Enterprise Linux operating system RPMs. If yum is configured, then the installer helps to install the required RPMs during the precheck.■ Extract the Veritas Access tar file and run the installer.■ Enter network configuration information (cluster name, IP addresses, bond interface information, DNS information, and so on).■ Verify installation on the node.	See "Installing and configuring the Veritas Access software on the cluster" on page 79.

Before you install

Before you install the Veritas Access software:

- Make sure that no DHCP servers are running in the private network.
- Disable the USB Ethernet interface in BIOS for all nodes in the cluster.
- Make sure that there are at least two private and two or more public links between cluster nodes.
- Connect the Fibre Channel cable before installing Veritas Access on any node in the cluster.
- Prepare one public IP address for each cluster node. This IP address is used by the installer, and it will be reused as one of the public interface physical IP addresses after configuration.
- Configure the newly prepared IP on the public NIC and in the network config file, `/etc/sysconfig/network-scripts/ifcfg-XX` to make it as persistent. For example:

```
TYPE=Ethernet
HWADDR=00:50:56:3d:f1:3e
DEVICE=eth2
BOOTPROTO=none
IPADDR=10.200.56.214
NETMASK=255.255.252.0
NM_CONTROLLED=no
ONBOOT=yes
```

Installing the operating system on each node of the cluster

Before you install the Veritas Access software, you must install the Red Hat Enterprise Linux operating system and kernel version. The following procedure includes the instructions and download links.

To install the Red Hat Enterprise Linux operating system on each node of the cluster

- 1 Meet the requisite system requirements. Ensure that you have the correct version of the Red Hat Linux operating system and the kernel version.
- 2 Use one of the following information to install Red Hat Enterprise Linux operating system as required:
 - Refer to *Chapter 1. Obtaining Red Hat Enterprise Linux* in the *Red Hat Enterprise Linux 7 Install guide*:
<https://access.redhat.com/downloads/>
https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/pdf/installation_guide/Red_Hat_Enterprise_Linux-7-Installation_Guide-en-US.pdf

About the driver node

If you do not plan to install Veritas Access from the console of the nodes in the cluster (the local management console of your nodes), you need another server that is not a target node in the Veritas Access cluster to use in the Veritas Access installation. This server is called the driver node.

When you run the Veritas Access installation script, the Veritas Access installer helps set up the ssh connection between the driver node and the target Veritas Access cluster nodes.

Note: Installation may fail if it is triggered from one of the node in the cluster. Therefore, installation must be triggered from the driver node.

The driver node platform can be: RHEL 7, RHEL 6, SLES 11 SP2, or SLES 11 SP3.

Installing the operating system on the target Veritas Access cluster

This first task in the installation process is to install the Red Hat Enterprise Linux operating system on each node of the cluster.

To install the operating system

- 1 Insert the Red Hat Enterprise Linux operating system installation DVD, and boot the server from the DVD.

See [“Linux requirements”](#) on page 17.

You can also use an external USB DVD-ROM.

- 2 Disable the consistent network device naming feature when installing the Red Hat Enterprise operating system.

Note: By default, the consistent network device naming feature is enabled on Dell systems. To disable the feature, enter the following option on the boot command line: `linux biosdevname=0`

- 3 At the boot prompt, select the **Install and upgrade an existing system** option. Press **Enter**.
- 4 The installer asks you if you want to perform a media check or if you want to skip this option. Select **Skip** and continue with the installation.
- 5 The installation starts with the Red Hat Enterprise Linux banner, click **Next** to continue installation.
- 6 The installation displays a language selection screen. Veritas Access only supports English. Select English. English is used for the installation and as the system default. Click **Next** to continue.
- 7 Select the correct layout type for the keyboard you would prefer to use for the installation and as the system default. Once you have made your selection, click **Next** to continue.
- 8 You can install Red Hat Enterprise Linux on a large variety of storage devices. This screen lets you select either basic or specialized storage devices. Click **Next** to continue.

- 9 The installer automatically detects any existing installations of Red Hat Enterprise Linux. It asks you to choose whether you want to perform a **Fresh Installation** or **Upgrade an Existing Installation**.

If your system contains a Red Hat Enterprise Linux installation, a dialog appears asking whether you want to upgrade that installation. To perform an upgrade of an existing system, choose the appropriate installation from the drop-down list and select **Next**.

- 10 The installer prompts you to set the host name for the root user. It also asks if you want to configure the network. Set the IP address to a proper public NIC so that you can access the NIC through that IP after a restart. Click **Next** to proceed.
- 11 In the **Time Zone Configuration** screen, set your time zone by selecting the city closest to your computer's physical location. Click **Next** to proceed.
- 12 The installer prompts you to set a root password for your system. You cannot proceed to the next stage of the installation process without entering a root password. Enter the root password into the **Root Password** field. Red Hat Enterprise Linux displays the characters as asterisks for security. Type the same password into the **Confirm** field to ensure that it is set correctly. After you set the root password, select **Next** to proceed.
- 13 The installer asks you to choose the type of installation. Depending on your need, select the layout for installation.
- 14 In the **Disk Partitioning Setup** screen, you can choose to create the default partition layout in one of four different ways. Or you can choose to partition storage devices manually to create a custom layout.
- 15 If you selected the **Encrypt System** option, the installer prompts you for a pass phrase with which to encrypt the partitions on the system.
- 16 If you selected more than one storage device on the storage devices selection screen, the installer asks you to select which of these devices should be available for installation of the operating system, and which should only be attached to the file system for data storage. If you selected only one storage device, the installer does not present you with this screen. During installation, the devices that you identify as being for data storage only are mounted as part of the file system, but are not partitioned or formatted. When you have finished identifying devices to be used for installation, click **Next** to continue.
- 17 If no readable partition tables are found on existing hard disks, the installation program asks to initialize the hard disk. Click **Re-initialize drive**.

- 18 The installer prompts you to create a custom layout or modify the default layout. The installer also prompts you to confirm the partitioning options that you selected. Click **Write changes to disk** to allow the installer to partition your hard drive and install Red Hat Enterprise Linux.
- 19 The **Package Installation Defaults** screen appears and details the default packages for your Red Hat Enterprise Linux installation. This screen varies depending on the version of Red Hat Enterprise Linux you want to install.

Configure the boot loader (GRUB) and select the installation packages as per the requirements. The installer searches for the required packages from the installation media and installs those packages. Select the **Basic Server** for RHEL 6.6, RHEL 6.7, and RHEL 6.8 versions and select **Minimal Install** option for RHEL 7.3 and RHEL 7.4 versions for installation.
- 20 Once the package installation gets completed, you have to restart your system for post-installation tasks. Remove the install media and click on **Reboot** to continue.
- 21 Your Red Hat Enterprise Linux installation is now complete. You can follow the same steps that are shown in this section to install the operating system on other nodes of the cluster.

See the *Red Hat Enterprise Linux documentation* for the detailed procedure.
- 22 Disable SELinux on all nodes in the cluster.

Installing the Oracle Linux operating system on the target Veritas Access cluster

This first task in the installation process is to install the Oracle Linux (OL) operating system on each node of the cluster.

To install the OL operating system

- 1 Insert the OL operating system installation DVD, and start the server from the DVD, and press the **Enter** key.

See [“Linux requirements”](#) on page 17.

You can also use an external USB DVD-ROM.
- 2 Press the **Tab** key to move focus to the **Skip** key, then press the **Enter** key to continue.
- 3 On the **Welcome** screen, click the **Next** option.
- 4 Select **English** as the language, and then click the **Next** option.
- 5 Select **U.S. English** as the keyboard setting, and then click the **Next** option.
- 6 Select the **Basic Storage Devices** option, and then click the **Next** option.

- 7 Enter a fully qualified host name, then click the **Configure Network** option.
- 8 Highlight the relevant connection and click the **Edit** option.
- 9 Select the **Connect automatically** check box. If you are not using DHCP, click on the **IPv4 Settings** tab, set the method to **Manual**, click the **Add** option, and enter the appropriate network details. When you are satisfied with the details, click the **Apply** and **Close** options to return to the host name screen, and then click the **Forward** option.
- 10 Select the appropriate time zone by clicking on your nearest city on the map. Click the **Next** option.
- 11 Enter a root password for the server and click on **Next**.
- 12 Select the partitioning type you require. If you want to amend the default partitioning layout, select the **Review and modify partitioning layout** option. Click the **Next** option.
- 13 The OEL installer lists the default partitioning scheme for your size disk. Amend them as required and click the **Next** option. Click the **Format** and **Write changes to disk** options.
- 14 Accept the boot loader settings by clicking the **Next** option.
- 15 Select the **Basic Server** and the **Oracle Linux Server** options, and then click the **Next** option.
- 16 Wait for the installation to complete.
- 17 Click the **Reboot** option to complete the installation.
- 18 Veritas Access supports only Red Hat Enterprise Linux compatible kernels. Obtain the Red Hat Enterprise Linux compatible kernels directly after the OEL installation.

See [“Obtaining the Red Hat Enterprise Linux compatible kernels”](#) on page 76.

Obtaining the Red Hat Enterprise Linux compatible kernels

Veritas Access supports only Red Hat Enterprise Linux compatible kernels.

After installing the Oracle Linux (OL) operating system, obtain the Red Hat Enterprise Linux (RHEL) compatible kernels. To obtain the RHEL compatible kernels, check the `/etc/grub.conf` file and change the default from 0 to 1.

See [“Installing the Oracle Linux operating system on the target Veritas Access cluster”](#) on page 75.

To obtain the RHEL compatible kernels directly after the OL installation

- 1 Check the `/etc/grub.conf` file and change `default=0` to `default=1`.

```
default=1
timeout=5
splashimage=(hd0,0)/grub/splash.xpm.gz
hiddenmenu
title Oracle Linux Server Unbreakable Enterprise Kernel (3.8.13-68.3.4.el6uek.x86_64)
    root (hd0,0)
    kernel /vmlinuz-3.8.13-68.3.4.el6uek.x86_64 ro root=/dev/mapper/vg_pilotvm13-lv_root
rd_NO_LUKS rd_LVM_LV=vg_pilotvm13/lv_swap LANG=en_US.UTF-8 rd_LVM_LV=vg_pilotvm13/lv_root
rd_NO_MD SYSFONT=latarcyrheb-sun16 crashkernel=auto
KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet
    initrd /initramfs-3.8.13-68.3.4.el6uek.x86_64.img
title Oracle Linux Server Red Hat Compatible Kernel (2.6.32-573.el6.x86_64)
    root (hd0,0)
    kernel /vmlinuz-2.6.32-573.el6.x86_64 ro root=/dev/mapper/vg_pilotvm13-lv_root
rd_NO_LUKS rd_LVM_LV=vg_pilotvm13/lv_swap LANG=en_US.UTF-8 rd_LVM_LV=vg_pilotvm13/lv_root
rd_NO_MD SYSFONT=latarcyrheb-sun16 crashkernel=auto
KEYBOARDTYPE=pc KEYTABLE=us rd_NO_DM rhgb quiet
    initrd /initramfs-2.6.32-573.el6.x86_64.img
```

- 2 Before installing Veritas Access, check for the file

`/etc/yum.repos.d/public-yum-ol6.repo`.

Generally the latest version is enabled as:

```
[public_ol6_latest]
name=Oracle Linux $releasever Latest ($basearch)
baseurl=http://public-yum.oracle.com/repo/OracleLinux/OL6/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1
```

- 3 Because Veritas Access does not support RHEL 6.9 or OL 6.9, you need to make the following change to the file.

If

```
`uname -r` = 2.6.32-573.el6.x86_64
```

```
2.6.32-504.el6.x86_64 RHEL 6.6 kernel
```

```
2.6.32-573.el6.x86_64 RHEL 6.7 kernel
```

```
2.6.32-642.el6.x86_64 RHEL 6.8 kernel
```

then change

```
[public_ol6_latest]
name=Oracle Linux $releasever Latest ($basearch)
baseurl=http://public-yum.oracle.com/repo/OracleLinux/OL6/latest/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=0
```

and search for

```
public_ol6_u7_base
```

Mark the `enabled=0` to `enabled=1`.

```
[public_ol6_u7_base]
name=Oracle Linux $releasever Update 7 installation media copy ($basearch)
baseurl=http://public-yum.oracle.com/repo/OracleLinux/OL6/7/base/$basearch/
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle
gpgcheck=1
enabled=1
```

- 4 Run the following command to install Veritas Access:

```
# ./installaccess
```

Installing Veritas Access on the target cluster nodes

Installing the cluster is a one-time activity. You can install up to a 20-node cluster. Before you continue, be aware of the following parameters:

- If you do not allocate enough IP addresses for the cluster, the installation cannot proceed.

Note: You cannot mix IPv4 and IPv6 addresses; new IP addresses must be of the same version that you initially used when installing Veritas Access.

See [“About obtaining IP addresses”](#) on page 57.

It takes about 40 minutes to install a two-node cluster. Installation times may vary depending on your configuration and the number of nodes.

Installing and configuring the Veritas Access software on the cluster

To install and configure the cluster

Note: During the installation, the installer log is located at `/var/tmp`.

- 1 Enter one of the following commands to start the installation.

```
# ./installaccess node1_ip node2_ip
```

Where `node1_ip` and `node2_ip` are the public physical IP addresses that are already assigned to the target cluster nodes to install Veritas Access over ssh.

These are the current IPs assigned to the nodes for installation communication.

The example is used to install two nodes. To install another target node cluster, add `node3_ip` to the command line that is used in this step.

- 2 The installer checks for the operating system dependencies and automatically installs the required OS RPMs. In case the OS RPMs' dependencies are not sorted, then the Redhat subscription manager user id and password is required.
- 3 The installer installs the Veritas Access RPMs.
- 4 Choose the licensing method. Answer the licensing questions and follow the prompts.

```
1) Enter a valid perpetual or subscription license key file
2) Register with evaluation mode and complete system licensing later
```

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

- 5** The installer displays the firewall ports to be opened after the configuration, and asks if you want to open them:

Veritas Access needs to open the following ports:

```
111 Rpcbind (NFS)
11211 Memcached Port
123 NTP Service
139 CIFS Service
14161 GUI
161 SNMP Service
2049 NFS Service
21 FTP Port
22 SSH Service
25 SMTP Port
30000:40000 FTP Passive Port Range
3172,3173 Server View Ports
3260 iSCSI Port
4001 Mountd (NFS)
4045 NLM (NFS)
4379 CTDB Port
445 CIFS TCP Service
51001,51002 RDMA Service
514 Syslog Service
53 DNS Service
5634 VIOM
56987 Replication Service
756,757,755 Statd (NFS)
8088 REST Server
8143 Object Access Gateway
8144 Object Access Admin Gateway
Do you want to proceed? [y,n,q] (y)
```

- 6** The installer automatically configures the RDMA environment on the cluster nodes if there are InfiniBand NICs.

7 The installer asks the following information to configure the cluster:

```
Enter the cluster name: [q,?]
Enter the public IP starting address: [b,q,?]
Enter the netmask for the public IP address: [b,q,?] (255.255.240.0)
Enter the number of VIPs per interface: [b,q,?] (0)
Enter the virtual IP starting address: [b,q,?]
Enter the default gateway IP address: [b,q,?]
Enter the DNS IP address: [b,q,?] (10.182.128.34)
Enter the DNS domain name: [b,q,?] (engba.veritas.com)
Enter the console virtual IP address: [b,q,?]
Do you want to use the separate console port? [y,n,q,b,?] (n)
```

Note: Cluster names should be DNS-compatible. DNS-compliant bucket names should conform to the following naming conventions. Bucket names must be at least three and no more than 63 characters long. Allowed characters in a cluster name are 'a-z, 0-9, -' lowercase letters, numbers, and hyphens. Cluster name must start with a lowercase letter and must end with a lowercase letter or a number. Any other character is invalid. A bucket name should not be an IP address. A period ('.') is not recommended in a bucket name because of an SSL issue, and including a period in a bucket name is not supported. Also, if a separate console port is chosen, the first public NIC is chosen to work exclusively as a console port.

8 The installer asks if you want to configure the Network Time Protocol (NTP) server.

```
Do you want to configure the Network Time Protocol (NTP) server to
synchronize the system clocks? [y,n,q] y
Enter the Network Time Protocol server: [q,?]
```

If you enter **y**, you can type in your NTP server. If you enter **n**, the NTP server is not configured.

- 9** The installer detects the network devices. The installer asks if you want to configure NIC bonding or exclude NICs.

```
Do you want to configure NIC bonding or exclude NICs or configure  
VLAN tagging? [y,n,q] (n)
```

If you do not want to configure NIC bonding or exclude NICs, enter **n**. Go to step [To install and configure the cluster](#) .

If you do want to configure NIC bonding or exclude NICs, enter **y**.

See [“Excluding a NIC”](#) on page 86.

See [“Creating a new NIC bond”](#) on page 94.

See [“Adding a VLAN device on a particular NIC”](#) on page 104.

10 The installer prompts to verify the network configuration.

Verify that the configuration information such as the new IP addresses, host name, and other details are correct.

Configuration checklist:

System	Hostname	New Hostname
192.168.10.1	oldhostname01.example.com	newhostname01.example.com
192.168.10.2	oldhostname02.example.com	newhostname02.example.com

System	Gateway IP	DNS IP	Domain name
192.168.10.1	192.168.10.3	192.168.10.0	hostname.example.com
192.168.10.2	192.168.10.3	192.168.10.0	hostname.example.com

System	NIC name (previous name)	Physical IP
192.168.10.1	pubeth0 (eth10)	192.168.10.5
192.168.10.1	pubeth1 (eth11)	192.168.10.6
192.168.10.2	pubeth0 (eth10)	192.168.10.7
192.168.10.2	pubeth1 (eth11)	192.168.10.8

Virtual IP

192.168.10.10	192.168.10.11	192.168.10.12	192.168.10.13
---------------	---------------	---------------	---------------

System	NIC name (previous name)
192.168.10.1	priveth0 (eth8)
192.168.10.1	priveth1 (eth9)
192.168.10.2	priveth0 (eth8)
192.168.10.2	priveth1 (eth9)

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

- 11 The installer renames the NICs and host name, and assigns the IPs for the systems after the confirmation. The installer also checks the Low Latency Transport (LLT) link status and automatically selects them.

Note: The installer does not check the LLT link status if the InfiniBand NICs are chosen as private NICs. See [“About using LLT over the RDMA network for Veritas Access”](#) on page 51.

- 12 If you are installing the ISO image on any node of the cluster (not the driver node), your session will get disconnected due to the network configuration. Re-login using the IP address prompted by the installer to resume installation.

The system is performing network configuration tasks. The current session will be temporarily disconnected. Re-login with 10.209.193.248 to resume installation.

- 13 The installer prompts to ask if you want to configure I/O fencing during the installation.

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

If you do not want to configure I/O fencing, enter **n**. If you plan to use both shared disks and local disks in your cluster, we recommend that you do not configure I/O fencing with the installer. Go to step [To install and configure the cluster](#).

To configure I/O fencing, enter **y**.

See [“Configuring I/O fencing”](#) on page 106.

- 14 The installer automatically restarts the cluster nodes to enable the Kdump function for each node.
- 15 Check the log file to confirm the installation and configuration. Logs can be found in `/opt/VRTS/install/logs/`.

Note: After the installation, connect to the Veritas Access console using the console IP address you assigned earlier, then log on using the default user name `master` and the default password `master`.

Veritas Access Graphical User Interface

Veritas Access has a Graphical User Interface (GUI) that provides a dashboard for a specific Veritas Access cluster, as well as views for shares, storage infrastructure,

reports, and settings. The GUI lets the administrator perform tasks for the cluster and monitor the results. In this release, the GUI is part of Veritas Access.

After you complete I/O fencing configuration successfully, the link to the GUI appears on the screen.

Open the `https://<console IP>:14161` URL in your browser to start the Veritas Access GUI application.

About NIC bonding and NIC exclusion

When you install Veritas Access on a cluster, you can perform the following operations using the NICs:

- Exclude a NIC
See [“Excluding a NIC”](#) on page 86.
- Include a NIC
See [“Including a NIC”](#) on page 90.
- Create a new NIC bond and add a NIC to a bond
See [“Creating a new NIC bond”](#) on page 94.
- Remove a bond
See [“Removing a NIC bond”](#) on page 99.
- Remove a NIC from the bond list
See [“Removing a NIC from the bond list”](#) on page 102.
- Add a VLAN device on a particular NIC
See [“Adding a VLAN device on a particular NIC”](#) on page 104.

Note: The NIC bonding and NIC exclusion configuration options support both a single NIC or bond, and multiple NICs or bonds.

Note: When using the NIC exclusion feature, you can exclude any NIC on the first node. But if you want to exclude any NIC on the other nodes, you can choose to exclude NICs per node.

See [“Excluding a NIC”](#) on page 86.

Note: If you want to use the NIC bonding feature, make sure that the PCI IDs of the slave bond NICs of the first node is the same as the PCI IDs of the slave bond NICs of the other nodes.

Excluding a NIC

When you install Veritas Access on a cluster, you may want to use some of the NICs for other storage purposes. You can use the `Exclude a NIC` functionality to exclude some NICs that you do not want to use for Veritas Access.

Note: The NIC bonding/NIC exclusion configuration options support both a single NIC or bond, and multiple NICs or bonds.

To exclude a NIC

- 1 During Veritas Access installation, the installer asks if you want to configure NIC bonding or exclude NICs. Enter **y** if you want to exclude a NIC.

```
Do you want to configure NIC bonding or exclude NICs or configure VLAN
tagging? [y,n,q] (n)
```

- 2 The installer prompts you to enter your selection. Enter **1** to exclude a NIC.

```
Veritas Access 7.3.1 Configure Program
10.200.114.45 10.200.114.46
```

```
NIC bonding/NIC exclusion configuration
```

```
NIC bonding supports only public NICs. Make sure the NICs you choose
are connected to public network.
```

```
NIC  PCI ID          bond status    If excluded
=====
```

```
eth2 0000:02:03.0 (physical NIC) N
```

```
eth3 0000:02:04.0 (physical NIC) N
```

```
eth4 0000:02:05.0 (physical NIC) N
```

```
eth5 0000:02:06.0 (physical NIC) N
```

```
eth6 0000:02:07.0 (physical NIC) N
```

```
eth7 0000:02:08.0 (physical NIC) N
```

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Create VLAN device
- 8) Delete VLAN device
- 9) Save and Continue

```
Select the NIC option to be configured in this cluster: [1-9,q] 1
```

- 3** The installer prompts you to select the NIC that you want to exclude. Enter your choice.

Choose NICs for exclusion

- 1) eth2 0000:02:03.0 (physical NIC)
- 2) eth3 0000:02:04.0 (physical NIC)
- 3) eth4 0000:02:05.0 (physical NIC)
- 4) eth5 0000:02:06.0 (physical NIC)
- 5) eth6 0000:02:07.0 (physical NIC)
- 6) eth7 0000:02:08.0 (physical NIC)
- 7) Exclude NICs per node
- b) Back to previous menu

Choose NICs: [1-7,b,q] 1 2(1,2)

- 4 The installer goes back to the previous menu. You can choose another NIC for exclusion. Enter **1** to exclude another NIC. Or you can save your configurations and continue with the installation of Veritas Access.

If you want to save your configurations, enter **9** :

```
Veritas Access 7.3.1 Configure Program
10.200.114.45 10.200.114.46
```

```
NIC bonding/NIC exclusion configuration
```

NIC bonding supports only public NICs. Make sure the NICs you choose are connected to public network.

NIC	PCI ID	bond status	If excluded
eth2	0000:02:03.0	(physical NIC)	Y
eth3	0000:02:04.0	(physical NIC)	Y
eth4	0000:02:05.0	(physical NIC)	N
eth5	0000:02:06.0	(physical NIC)	N
eth6	0000:02:07.0	(physical NIC)	N
eth7	0000:02:08.0	(physical NIC)	N

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Create VLAN device
- 8) Delete VLAN device
- 9) Save and Continue

Select the NIC option to be configured in this cluster: [1-9,q] 9

- 5 If you want to exclude NICs per node, in Step [To exclude a NIC](#) enter 7. The NICs with inconsistent PCI IDs are listed:

```
Choose NICs for exclusion
```

- 1) eth2 0000:02:03.0 (physical NIC)
- 2) eth3 0000:02:04.0 (physical NIC)
- 3) eth4 0000:02:05.0 (physical NIC)
- 4) eth5 0000:02:06.0 (physical NIC)
- 5) eth6 0000:02:07.0 (physical NIC)
- 6) eth7 0000:02:08.0 (physical NIC)
- 7) Exclude NICs per node
- b) Back to previous menu

```
Choose NICs: [1-7,b,q] 7
```

```
Choose items: [1-1,b,q] 1
```

- 1 0000:02:00.0 (10.198.95.214)
- 2 0000:02:01.0 (10.198.95.214)
- 3 0000:02:06.0 (10.198.95.212)
- 4 0000:02:09.0 (10.198.95.214)
- 5 0000:02:14.0 (10.198.95.212)
- 6 0000:02:15.0 (10.198.95.212)
- b) Back to previous menu

```
Choose NICs: [1-6,b,q] 1 2 3 4 5 6
```

Note: NIC exclusion function is supported on InfiniBand NICs, but all the NICs with the same PCI ID are excluded during the exclusion operation. See [“About using LLT over the RDMA network for Veritas Access”](#) on page 51.

Including a NIC

When you install Veritas Access on a cluster, you may want to include one or more NICs that you had previously excluded. You can use the `Include a NIC` functionality to include NICs that you want to use for Veritas Access.

To include a NIC

- 1 If you have excluded some NICs and not saved your configuration, it is possible to include a NIC again. When the installer asks you to select the NIC option that you want to configure in the cluster, enter **2** if you want to include a NIC.

```
Veritas Access 7.3.1 Configure Program
10.200.114.45 10.200.114.46
```

```
NIC bonding/NIC exclusion configuration
```

NIC bonding supports only public NICs. Make sure the NICs you choose are connected to public network.

NIC	PCI ID	bond status	If excluded
eth2	0000:02:03.0	(physical NIC)	Y
eth3	0000:02:04.0	(physical NIC)	Y
eth4	0000:02:05.0	(physical NIC)	N
eth5	0000:02:06.0	(physical NIC)	N
eth6	0000:02:07.0	(physical NIC)	N
eth7	0000:02:08.0	(physical NIC)	N

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Create VLAN device
- 8) Delete VLAN device
- 9) Save and continue

```
Select the NIC option to be configured in this cluster: [1-9,q] 2
```

- 2** The installer prompts you to select the NIC that you want to include. Enter your choice.

Choose NICs for inclusion

- 1) eth2 0000:02:03.0 (excluded NIC)
- 2) eth3 0000:02:04.0 (excluded NIC)
- 3) Include NICs per node
- b) Back to previous menu

Choose NICs: [1-6,b,q] 1

- 3** The installer goes back to the previous menu. You can choose another NIC for inclusion. Enter **2** to include another NIC. Or you can save your configurations and continue with the installation of Veritas Access.

If you want to save your configurations, enter **9**.

```
Veritas Access 7.3.1 Configure Program
10.200.114.45 10.200.114.46
```

```
NIC bonding/NIC exclusion configuration
```

NIC bonding supports only public NICs. Make sure the NICs you choose are connected to public network.

```
NIC  PCI ID          bond status    If excluded
=====
eth2  0000:02:03.0 (physical NIC) N
eth3  0000:02:04.0 (physical NIC) Y
eth4  0000:02:05.0 (physical NIC) N
eth5  0000:02:06.0 (physical NIC) N
eth6  0000:02:07.0 (physical NIC) N
eth7  0000:02:08.0 (physical NIC) N
```

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Create VLAN device
- 8) Delete VLAN device
- 9) Save and continue

Select the NIC option to be configured in this cluster: [1-9,q]

- 4** If you want to include NICs per node, in Step 2 enter **3**.

Note: NIC inclusion function is supported on InfiniBand NICs, but all the NICs with same PCI ID are included during the include operation. See [“About using LLT over the RDMA network for Veritas Access”](#) on page 51.

Creating a new NIC bond

An administrator can create a bond NIC interface from a given list of public NIC interfaces during Veritas Access installation. This feature allows an administrator to save a number of physical IP addresses that are used for installation and post-installation bond creation.

- The bond interface feature is available for network interface card (NIC) bonding of public interfaces only. Bonding of private interfaces is not supported.
- You cannot bond InfiniBand NICs since the PCI IDs are identical. See [“About using LLT over the RDMA network for Veritas Access”](#) on page 51.
- You can create a bond interface for public NICs only if the PCI IDs of the public NICs are correspondingly same among the nodes.

If you do not want to create a bond interface, continue with the installation.

See [“About obtaining IP addresses”](#) on page 57.

See [“About calculating IP address requirements”](#) on page 58.

To create a bond

- 1 During the Veritas Access installation, the installer asks if you want to configure NIC bonding or exclude NICs. Enter **y** if you want to configure a NIC bond.

```
Do you want to configure NIC bonding or exclude NICs? [y,n,q] (n) y
```

- 2 The installer prompts you to enter your selection. Enter **3** to create a new bond.

```
Veritas Access 7.3.1 Configure Program  
10.200.114.45 10.200.114.46
```

```
NIC bonding/NIC exclusion configuration
```

```
NIC bonding supports only public NICs. Make sure the NICs you choose  
are connected to public network.
```

```
NIC  PCI ID          bond status   If excluded  
=====
```

eth2	0000:02:03.0	(physical NIC)	N
eth3	0000:02:04.0	(physical NIC)	N
eth4	0000:02:05.0	(physical NIC)	N
eth5	0000:02:06.0	(physical NIC)	N
eth6	0000:02:07.0	(physical NIC)	N
eth7	0000:02:08.0	(physical NIC)	N

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Save and Continue

```
Select the NIC option to be configured in this Cluster: [1-7,q] 3
```

- 3** The installer prompts you to select the bond mode of the new bond. Enter your choice.

Configure the mode for the NIC bonding:

- 1) balance-rr
- 2) active-backup
- 3) balance-xor
- 4) broadcast
- 5) 802.3ad
- b) Back to previous menu

Select the mode of bond: [1-5,b,q] 3

bond0 is created.

Press [Enter] to continue

- 4** If you choose **3** or **5**, the installer prompts you to choose the bond option for the bond mode:

- 1) layer2
- 2) layer3+4
- 3) default

Select the bonding option: [1-3,b,q] 1

5 The installer prompts you to select the NIC option that you want to configure in the cluster.

```
Veritas Access 7.3.1 Configure Program
10.200.114.45 10.200.114.46
```

```
NIC bonding/NIC exclusion configuration
```

```
NIC bonding supports only public NICs. Make sure the NICs you choose
are connected to public network.
```

```
NIC  PCI ID          bond status   If excluded
=====
eth2 0000:02:03.0 (physical NIC) N
eth3 0000:02:04.0 (physical NIC) N
eth4 0000:02:05.0 (physical NIC) N
eth5 0000:02:06.0 (physical NIC) N
eth6 0000:02:07.0 (physical NIC) N
eth7 0000:02:08.0 (physical NIC) N
```

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Save and Continue

```
Select the NIC option to be configured in this Cluster: [1-7,q] 4
```

6 The installer prompts you to select the NIC.

- 1) eth2 0000:02:03.0 (physical NIC)
- 2) eth3 0000:02:04.0 (physical NIC)
- 3) eth4 0000:02:05.0 (physical NIC)
- 4) eth5 0000:02:06.0 (physical NIC)
- 5) eth6 0000:02:07.0 (physical NIC)
- 6) eth7 0000:02:08.0 (physical NIC)
- b) Back to previous menu

```
Choose NICs: [1-6,b,q] 1
```

7 The installer prompts you to choose a bond name to which you want to add the NIC.

- 1) bond0
- b) Back to previous menu

Choose a bond: [1-1,b,q] 1
 Adding 0000:02:03.0 to bond0 was successful

Press [Enter] to continue:

8 The installer prompts you to select the NIC option that you want to configure in the cluster.

Enter **4** if you want to add another NIC to the bond . Or you can enter **7** to save your configurations and continue with the installation of Veritas Access.

Veritas Access 7.3.1 Configure Program
 10.200.114.45 10.200.114.46

NIC bonding/NIC exclusion configuration

NIC bonding supports only public NICs. Make sure the NICs you choose are connected to public network.

NIC	PCI ID	bond status	If excluded
eth2	0000:02:03.0	(Slave of bond0)	N
eth3	0000:02:04.0	(Slave of bond0)	N
eth4	0000:02:05.0	(physical NIC)	N
eth5	0000:02:06.0	(physical NIC)	N
eth6	0000:02:07.0	(physical NIC)	N
eth7	0000:02:08.0	(physical NIC)	N

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Save and Continue

Select the NIC option to be configured in this Cluster: [1-7,q] 7

Removing a NIC bond

An administrator can remove a bond that has been already created.

To remove a NIC bond

- 1 During the Veritas Access installation, the installer prompts you to enter your selection. Enter **5** to remove an existing bond.

```
Veritas Access 7.3.1 Install Program
10.200.114.45 10.200.114.46
```

```
NIC bonding/NIC exclusion configuration
```

```
NIC bonding supports only public NICs. Make sure the NICs you choose
are connected to public network.
```

```
NIC  PCI ID          BOND status      If excluded
=====
eth0  0000:02:01.0 (physical NIC)   N
eth1  0000:02:02.0 (physical NIC)   N
eth2  0000:02:03.0 (Slave of bond0) N
eth3  0000:02:04.0 (Slave of bond0) N
eth4  0000:02:05.0 (Slave of bond2) N
eth5  0000:02:06.0 (Slave of bond1) N
eth6  0000:02:07.0 (Slave of bond1) N
eth7  0000:02:08.0 (Slave of bond2) N
```

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Save and Continue

```
Select the NIC option to be configured in this Cluster: [1-7,q] 5
```

2 The installer prompts you to select the bond which you want to remove. Enter your choice.

- 1) bond0
- 2) bond1
- 3) bond2
- b) Back to previous menu

Choose bonds: [1-3,b,q] 3

Deleting NIC bonding bond2 succeeded

Press [Enter] to continue:

- 3** The installer prompts you to select the NIC option that you want to configure in the cluster. Enter **5** if you want to remove another bond . Or you can enter **7** to save your configurations and continue with the installation of Veritas Access.

```
Veritas Access 7.3.1 Install Program
10.200.114.45 10.200.114.46
```

```
NIC bonding/NIC exclusion configuration
```

```
NIC bonding supports only public NICs. Make sure the NICs you choose
are connected to public network.
```

```
NIC  PCI ID          BOND status      If excluded
=====
eth0 0000:02:01.0 (physical NIC)   N
eth1 0000:02:02.0 (physical NIC)   N
eth2 0000:02:03.0 (Slave of bond0) N
eth3 0000:02:04.0 (Slave of bond0) N
eth4 0000:02:05.0 (physical NIC)   N
eth5 0000:02:06.0 (Slave of bond1) N
eth6 0000:02:07.0 (Slave of bond1) N
eth7 0000:02:08.0 (physical NIC)   N
```

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Save and Continue

```
Select the NIC option to be configured in this Cluster: [1-7,q]
```

Removing a NIC from the bond list

During installation, an administrator can remove a NIC which is already a slave of a bond before the configuration is saved.

To remove a NIC from the bond list

- 1** During the Veritas Access installation, the installer prompts you to enter your selection. Enter **6** to remove a NIC from the bond list.

Note: The NIC bonding/NIC exclusion configuration options support both a single NIC or bond, and multiple NICs or bonds.

```
Veritas Access 7.3.1 Install Program
10.200.114.45 10.200.114.46
```

```
NIC bonding/NIC exclusion configuration
```

NIC bonding supports only public NICs. Make sure the NICs you choose are connected to public network.

NIC	PCI ID	BOND status	If excluded
eth0	0000:02:01.0	(physical NIC)	N
eth1	0000:02:02.0	(physical NIC)	N
eth2	0000:02:03.0	(Slave of bond0)	N
eth3	0000:02:04.0	(Slave of bond0)	N
eth4	0000:02:05.0	(Slave of bond0)	N
eth5	0000:02:06.0	(Slave of bond1)	N
eth6	0000:02:07.0	(Slave of bond1)	N
eth7	0000:02:08.0	(physical NIC)	N

- 1) Exclude NICs
- 2) Include NICs
- 3) Create a new bond
- 4) Add NICs to a bond
- 5) Remove bonds
- 6) Remove NICs from the bond list
- 7) Save and Continue

Select the NIC option to be configured in this Cluster: [1-7,q] 6

- 2 The installer prompts you to select the NIC that you want to delete from the NIC bonding. Enter your choice.

```
Choose a NIC to be deleted from the NIC bonding
```

- ```
1) eth2 0000:02:03.0 (Slave of bond0)
2) eth3 0000:02:04.0 (Slave of bond0)
3) eth4 0000:02:05.0 (Slave of bond0)
4) eth5 0000:02:06.0 (Slave of bond1)
5) eth6 0000:02:07.0 (Slave of bond1)
b) Back to previous menu
```

```
Choose a NIC: [1-8,b,q] 5
```

```
The NICs with the PCI id 0000:02:05.0 has been removed from /
the NIC bonding
```

```
Press [Enter] to continue:
```

- 3 The installer prompts you to select the NIC option that you want to configure in the cluster. Enter **6** if you want to remove another NIC from the bond list . Or you can enter **7** to save your configurations and continue with the installation of Veritas Access.

## About VLAN Tagging

When VLANs (Virtual Local Area Network) span multiple switches, VLAN Tagging is required. A VLAN is a way to create independent logical networks within a physical network. VLAN Tagging is the practice of inserting a VLAN ID into a packet header to identify which VLAN the packet belongs to.

VLAN Tagging feature includes the following:

- Ability to create a VLAN device during installation.
- Create a VLAN device on the specified bond interface.  
You need to create a bond interface first.  
See [“Adding a VLAN device on a particular NIC”](#) on page 104.

## Adding a VLAN device on a particular NIC

See [“About VLAN Tagging”](#) on page 104.

**To add a VLAN device for a particular NIC**

- 1 Start the Veritas Access installation.
- 2 Select the NICs for exclusion.
- 3 Create a new bond interface.
- 4 Configure the bonding mode.
- 5 Add the NICs to the bond interface.
- 6 Choose the NICs for bonding.
- 7 Choose a bond name to add the NICs.
- 8 Select the `Create VLAN device` option.
- 9 Choose the NICs to create the VLAN device on.
- 10 Enter the VLAN ID for the device.
- 11 Select the `Save and continue` option.

## Limitations of VLAN Tagging

Note the following limitations for using VLAN Tagging:

- Support only for a fresh installation. VLAN Tagging is not supported for reconfiguration with the `-updateparameter` option and add node configuration.
- Support only for creating a VLAN device on a bonded NIC.
- Support only for creating one VLAN device at installation time.

## Replacing an Ethernet interface card

In some cases, you may need to replace an Ethernet interface card on a node. This section describes the steps to replace the card.

---

**Note:** This procedure works for replacing an existing Ethernet interface card. It does not work for adding an Ethernet interface card to the cluster. If the Ethernet interface card you add needs a new device driver, install the new device driver before installing the Ethernet interface card on the node.

---

### To replace an Ethernet interface card

- 1 Use the `Cluster> shutdown` command to shut down the node.

For example:

```
Cluster> shutdown access_03
Stopping Cluster processes on access_03.....done
Sent shutdown command to access_03
```

- 2 Use the `Cluster> del` command to delete the node from the cluster.

For example:

```
Cluster> del access_03
```

- 3 Install the replacement Ethernet interface card on the node.
- 4 Turn on the node.
- 5 Make sure that the Ethernet interface card is active and online.
- 6 Use the `Cluster> add` command to add the node back into the cluster.

For example:

```
Cluster> add 172.16.113.118
```

For details on the `Cluster> add` and `Upgrade>` commands that are described in this section, see the relevant man pages.

## Configuring I/O fencing

Veritas Access supports two fencing modes: disk-based fencing for a cluster with shared disks, and majority-based fencing for a cluster with local DAS disks.

If you intend to use both shared disks (SAN) and local disks, majority-based fencing must be used. Veritas recommends that you do not configure I/O fencing through the installer.

- 1 During the Veritas Access configuration, after the product is started, the installer asks whether to configure fencing:

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

- 2 Enter `y` to configure fencing.
  - If the cluster does not include initialized shared disks, majority-based fencing mode is configured.

The I/O fencing would be configured in majority mode.

- If shared disks are connected and initialized, disk-based I/O fencing is configured. You are prompted to choose disks.

---

**Note:** You can choose three available VxVM disks or initialize three disks as VxVM disks to form the fencing disk group. You must choose exactly three disks.

---

- 3 The installer stops the product, and applies the fencing configuration before restart.

## About configuring Veritas NetBackup

If you use Veritas NetBackup, to comply with the NetBackup End-User License Agreement (EULA), you have to purchase and enter valid license keys on the external NetBackup master server before you configure NetBackup to work with Veritas Access. For more information on entering the NetBackup license keys on the NetBackup master server, see the *Veritas NetBackup Installation Guide*.

If you use NetBackup, configure the virtual IP address using the `Backup> virtual-ip` command so that it is different from all of the virtual IP addresses, including the console server IP address and the physical IP addresses that are used to install the Veritas Access software.

## About enabling kdump during an Veritas Access configuration

During the Veritas Access configuration, the Veritas Access installer tries to enable kdump on your cluster node. To meet the Veritas Access software requirements, the installer modifies the `/etc/kdump.conf` and `/boot/grub/grub.conf` files by using the following options:

- `/boot/grub/grub.conf`  
`crashkernel = 512M-2G:64M, 2G-:256M`
- `/etc/kdump.conf`  
`path /opt/VRTSsnas/core/kernel/  
core_collector makedumpfile -c --message-level 1 -d 31`

# Reconfiguring the Veritas Access cluster name and network

After you install and configure Veritas Access, you can reconfigure the cluster name and network, if required.

Before you reconfigure the cluster, you have to enable the *support* user for the nodes because the root user access authority is forbidden. The *support* user default password is *veritas*. You can change the password after you log on the first time.

## To reconfigure the Veritas Access cluster name and network

- 1 Log on to the host console using the *support* user name and password.
- 2 Ensure that all the service groups are offline. Enter the following command:

```
/opt/VRTS/install/installaccess -updateparameter
```

**3** Enter the private IPs of the systems.

```
172.16.0.3 172.16.0.4
```

---

**Note:** Only the private IPs of the systems must be entered. Public IPs should not be used here.

---

**4** Enter the cluster name and network information.

```
Enter the cluster name:
Enter the public IP starting address:
Enter the netmask for the public IP address:
Enter the number of VIPs per interface:
Enter the virtual IP starting address:
Enter the default gateway IP address:
Enter the DNS IP address:
Enter the DNS domain name:
Enter the console virtual IP address:
Do you want to use the separate console port? [y,n,q] (n):
Do you want to configure the Network Time Protocol (NTP) server to
synchronize the system clocks? [y,n,q] (n) y:
Enter the Network Time Protocol server:
```

The installer confirms that the information that you entered is correct. The configuration is completed and the new cluster and IPs are configured on the cluster.

The installer displays the location of the log and summary files. If required, view the files to confirm the configuration status.

---

**Note:** The cluster name can contain only alpha characters, numbers, or underscores. The cluster name must start with a letter of the alphabet and can have a length of maximum 15 characters. Also, if a separate console port is chosen, the first public NIC is chosen to work exclusively as a console port.

---

**Note:** If your cluster has DAS disks, limit the cluster name to 10 characters. After formatting the DAS disks, do not change the cluster name.

---

# Configuring a KMS server on the Veritas Access cluster

You can configure a KMS server on the Veritas Access cluster.

## To configure a KMS server on the Veritas Access cluster

- 1 Obtain the KMS server's SSL public key (in base64 format) and its port number. This key is used for communication between the Veritas Access cluster and the KMS server.
- 2 Generate a self-signed SSL key-pair on the Veritas Access cluster:

```
System> kms certificate generate
```

- 3 Import the KMS server's public key.

```
System> kms certificate import_server_cert
```

- 4 Configure the KMS server. Provide the SSL public key that was obtained in step 1 as input here.

```
System> kms config server <server_ip> <server_port>
```

Where *server\_ip* is the KMS server IP

*server\_port* is the KMS server port number.

- 5 KMS admin now sets up a trust certificate using its admin GUI to allow communication between the KMS server and Veritas Access cluster.

For more information, see the `system_kms` man page.

# Automating Veritas Access installation and configuration using response files

This chapter includes the following topics:

- [About response files](#)
- [Performing a silent Veritas Access installation](#)
- [Response file variables to install and configure Veritas Access](#)
- [Sample response file for Veritas Access installation and configuration](#)

## About response files

The installer 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 an installation or uninstallation.

See [“Installation script options”](#) on page 153.

## Performing a silent Veritas Access installation

A silent installation and configuration is based on a response file that you prepare so that the Veritas Access software can be installed without prompts. This feature is useful if you want to install the Veritas Access software on a large number of nodes.

Before performing a silent Veritas Access installation and configuration, you have to manually configure a secure shell (ssh) communication between the nodes.

See [“Manually configuring passwordless secure shell \(ssh\)”](#) on page 157.

You can get the Veritas Access example response file from the root directory of the ISO image.

### To use the Veritas Access silent installation feature

- ◆ Enter the following command:

```
./installaccess -responsefile access.responsefile
```

### To generate the access.response example file

- 1 Install and configure the Veritas Access software without any errors.
- 2 Get the `access.response` example file from the log directory.

### To use the access.response example file

- 1 Rename the Veritas Access example response file to `access.responsefile`.
- 2 Modify the file by changing the cluster name, IP address ranges, and other parameters, as necessary for your configuration.

Installation times may vary depending on your configuration.

See [“Installing and configuring the Veritas Access software on the cluster”](#) on page 79.

## Response file variables to install and configure Veritas Access

[Table 7-1](#) lists the response file variables that you can define to install and configure Veritas Access.

**Table 7-1** Response file variables for installing Veritas Access

| Variable                           | Description                                                                                                                                                                                                                                             |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{bondmode}{bond<n>}             | Defines the bond modes for BOND.<br>List or scalar: list<br>Optional or required: optional                                                                                                                                                              |
| CFG{bondname}                      | List of bond names for BOND.<br>List or scalar: list<br>Optional or required: optional                                                                                                                                                                  |
| CFG{bondpool}{bond<n>}             | List of the PCI IDs of the slave NICs.<br>List or scalar: list<br>Optional or required: optional                                                                                                                                                        |
| CFG{config_majority_based_fencing} | Enables majority fencing. The value is 1. It cannot be used with I/O fencing variables 'fencing_scsi3_disk_policy', 'fencing_newdg_disks', and 'fencing_dgname'.<br>List or scalar: scalar<br>Optional or required: required for majority-based fencing |
| CFG{exclusion}                     | List of PCI IDs of excluded NICs.<br>List or scalar: list<br>Optional or required: optional                                                                                                                                                             |
| CFG{fencing_dgname}                | Specifies the disk group for I/O fencing. The value is <code>sfscoorddg</code> .<br>List or scalar: scalar<br>Optional or required: required for I/O fencing                                                                                            |
| CFG{fencing_newdg_disks}           | Defines the fencing disks.<br>List or scalar: list<br>Optional or required: required for I/O fencing                                                                                                                                                    |

**Table 7-1** Response file variables for installing Veritas Access (*continued*)

| Variable                       | Description                                                                                                                                                                                |
|--------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{fencing_option}            | <p>Specifies the I/O fencing configuration mode. The value is 2 for disk-based I/O fencing.</p> <p>List or scalar: scalar</p> <p>Optional or required: required for I/O fencing</p>        |
| CFG{fencing_scsi3_disk_policy} | <p>Specifies the SCSI-3 disk policy to use I/O fencing. The value is <code>dmp</code>.</p> <p>List or scalar: scalar</p> <p>Optional or required: required for I/O fencing</p>             |
| CFG{fencingenabled}            | <p>Defines whether fencing is enabled. The value is 1 if enabled.</p> <p>List or scalar: scalar</p> <p>Optional or required: required for I/O fencing</p>                                  |
| CFG{opt}{licensefile}          | <p>Specifies the location of the Veritas perpetual or subscription license key file.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                               |
| CFG{keys}{"node_ip"}           | <p>Specifies the Veritas Access license for each node.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                                             |
| CFG{newnodes}                  | <p>Specifies the new access IP for the cluster nodes. The value should be the first public IP address for each node.</p> <p>List or scalar: list</p> <p>Optional or required: required</p> |
| CFG{opt}{comcleanup}           | <p>Cleans up the ssh connection that is added by the installer after the configuration. The value is 1.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>            |

**Table 7-1** Response file variables for installing Veritas Access (*continued*)

| Variable                 | Description                                                                                                                                                                                       |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{confignic}      | <p>Performs the NIC configuration with all the network variable values. The value is 1.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                   |
| CFG{opt}{configure}      | <p>Performs the configuration if the packages are already installed.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                                      |
| CFG{opt}{install}        | <p>Installs Veritas Access RPMs. Configuration can be performed at a later time using the <code>-configure</code> option.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p> |
| CFG{opt}{installallpkgs} | <p>Instructs the installer to install all the Veritas Access RPMs based on the variable that has the value set to 1.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>      |
| CFG{opt}{noipc}          | <p>Disables the connection to SORT for updates check. The value is 0.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                                     |
| CFG{opt}{ssh}            | <p>Determines whether to use ssh for communication between systems. The value is 1 if enabled.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                            |
| CFG{prod}                | <p>Defines the product to be installed or uninstalled.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                                                    |

**Table 7-1** Response file variables for installing Veritas Access (*continued*)

| Variable                          | Description                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{publicbond}                   | <p>List of PCI IDs of the bonded NICs.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>                                                                                                                                                                                                                       |
| CFG{publicnetmaskarr}             | <p>List of netmasks that are assigned to public NICs or bonds.</p> <p>List or scalar: list</p> <p>Optional or required: required</p>                                                                                                                                                                                               |
| CFG{publicparr}                   | <p>List of public IPs that are assigned to public NICs or bonds.</p> <p>List or scalar: list</p> <p>Optional or required: required</p>                                                                                                                                                                                             |
| CFG{redhat_subscription_username} | <p>Specifies the user name to register with Red Hat subscription management.</p> <p>List or scalar: scalar</p> <p>Optional or required: required if some required OS rpms are missing on the systems</p> <p>The user name should be enclosed in single quotes (for example : '1234@abc') if it contains any special character.</p> |
| CFG{redhat_subscription_password} | <p>Specifies the password to register with Red Hat subscription management.</p> <p>List or scalar: scalar</p> <p>Optional or required: required if some required OS rpms are missing on the systems</p> <p>The password should be enclosed in single quotes (for example : '1234@abc') if it contains any special character.</p>   |
| CFG{snas_clustername}             | <p>Defines the cluster name of the product.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                                                                                                                                                                                                |

**Table 7-1** Response file variables for installing Veritas Access (*continued*)

| <b>Variable</b>         | <b>Description</b>                                                                                                        |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------|
| CFG{snas_consoleip}     | Defines the console IP of the product.<br>List or scalar: scalar<br>Optional or required: required                        |
| CFG{snas_defgateway}    | Defines the gateway of the product.<br>List or scalar: scalar<br>Optional or required: required                           |
| CFG{snas_dnsdomainname} | Defines the DNS domain name of the product.<br>List or scalar: scalar<br>Optional or required: required                   |
| CFG{snas_dnsip}         | Defines the DNS IP of the product.<br>List or scalar: scalar<br>Optional or required: required                            |
| CFG{snas_ntpserver}     | Defines the NTP server name of the product.<br>List or scalar: scalar<br>Optional or required: required                   |
| CFG{snas_nvip}          | Defines the number of VIPs on each NIC.<br>List or scalar: scalar<br>Optional or required: required                       |
| CFG{snas_pipprefix}     | Defines the prefix of public IPs (only in IPV6 environments).<br>List or scalar: scalar<br>Optional or required: required |
| CFG{snas_pipstart}      | Defines the the initial IP of the public IPs.<br>List or scalar: scalar<br>Optional or required: required                 |

**Table 7-1** Response file variables for installing Veritas Access (*continued*)

| Variable                 | Description                                                                                                                                                           |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{snas_pnmaskstart}    | <p>Defines the netmask of public IPs (only in IPV4 environments).</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                             |
| CFG{snas_sepconsoleport} | <p>Defines if use of separate console port. 1 for yes, 0 for no.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                              |
| CFG{snas_vipprefix}      | <p>Defines the prefix of virtual IPs (only in IPV6 environments).</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                             |
| CFG{snas_vipstart}       | <p>Defines the the initial IP of the virtual IPs.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                             |
| CFG{snas_vnmaskstart}    | <p>Defines the netmask of virtual IPs (only in IPV4 environments).</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                            |
| CFG{systems}             | <p>List of systems on which the product is to be installed or uninstalled.</p> <p>List or scalar: list</p> <p>Optional or required: required</p>                      |
| CFG{vcs_allowcomms}      | <p>Indicates whether to start LLT or GAB when the user wants to set up a single node cluster.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p> |

**Table 7-1** Response file variables for installing Veritas Access (*continued*)

| Variable                                | Description                                                                                                                                                                                                                    |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_clusterid}                      | Defines the unique cluster ID with a string number.<br><br>List or scalar: scalar<br><br>Optional or required: required                                                                                                        |
| CFG{vcs_lltlink<n>}{new_node_ip}        | Defines the NIC name for the first heartbeat link.<br><br>List or scalar: scalar<br><br>Optional or required: required                                                                                                         |
| CFG{vcs_rdmalink1_address}{new_node_ip} | Specifies the RDMA private link1 IP address, the value follows that node_01 is 172.16.0.3, node_02 is 172.16.0.4, and so on.<br><br>List or scalar: scalar<br><br>Optional or required: required for RDMA NICs as private NICs |
| CFG{vcs_rdmalink1_netmask}{new_node_ip} | Specifies the RDMA private link1 IP netmask, the value is 255.255.255.0.<br><br>List or scalar: scalar<br><br>Optional or required: required for RDMA NICs as private NICs                                                     |
| CFG{vcs_rdmalink1_port}{new_node_ip}    | Specifies the port number for the RDMA private link1, the value is 51001.<br><br>List or scalar: scalar<br><br>Optional or required: required for RDMA NICs as private NICs                                                    |
| CFG{vcs_rdmalink2_address}{new_node_ip} | Specifies the RDMA private link2 IP address, the value follows that node_01 is 172.16.1.3, node_02 is 172.16.1.4, and so on.<br><br>List or scalar: scalar<br><br>Optional or required: required for RDMA NICs as private NICs |

**Table 7-1** Response file variables for installing Veritas Access (*continued*)

| Variable                                | Description                                                                                                                                                                        |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_rdmalink2_netmask}{new_node_ip} | <p>Specifies the RDMA private link2 IP netmask, the value is 255.255.255.0.</p> <p>List or scalar: scalar</p> <p>Optional or required: required for RDMA NICs as private NICs</p>  |
| CFG{vcs_rdmalink2_port}{new_node_ip}    | <p>Specifies the port number for the RDMA private link2, the value is 51002.</p> <p>List or scalar: scalar</p> <p>Optional or required: required for RDMA NICs as private NICs</p> |
| CFG{vcs_userenpw}                       | <p>Defines the encrypted user password.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                                                    |
| CFG{vcs_username}                       | <p>Defines the added username for VCS.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                                                     |
| CFG{vcs_userpriv}                       | <p>Defines the user privilege.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                                                             |
| CFG{virtualiparr}                       | <p>List of virtual IPs that will be assigned to public NICs or bonds.</p> <p>List or scalar: list</p> <p>Optional or required: required</p>                                        |
| CFG{virtualnetmaskarr}                  | <p>List of netmasks that will be assigned to public NICs or bonds.</p> <p>List or scalar: list</p> <p>Optional or required: required</p>                                           |

# Sample response file for Veritas Access installation and configuration

The following example shows a response file for installing and configuring Veritas Access.

```

our %CFG;
#Installs Product packages.
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{comsetup}=1;
$CFG{opt}{noipc}=1;
$CFG{opt}{ssh}=1;
$CFG{prod}="SNAS73";
$CFG{opt}{licensefile}="<absolute_path_of_licfile>";

#Performs the configuration if the packages are already installed
$CFG{opt}{configure}=1;

#the PCI IDs of slave NICs
$CFG{bondpool}{bond0}=[qw(0000:02:09.0 0000:02:07.0)];
$CFG{bondpool}{bond1}=[qw(0000:02:04.0 0000:02:08.0)];

#mode of each bond
$CFG{bondmode}{bond0}=5;
$CFG{bondmode}{bond1}=6;

#names of bond
$CFG{bondname}=[qw(bond0 bond1)];

#the PCI IDs of excluded NICs
$CFG{exclusion}=[qw(0000:02:03.0 0000:02:0a.0)];

#the PCI IDs of all the bonded NICs
$CFG{publicbond}=[qw(0000:02:03.0 0000:02:04.0 0000:02:07.0
0000:02:08.0)];

#public IPs
$CFG{publiciparr}=[qw(10.200.58.100 10.200.58.101 10.200.58.102
10.200.58.103 10.200.58.104 10.200.58.105 10.200.58.106 10.200.58.107)];

#netmask for public IPs
```

```
$CFG{publicnetmaskarr}=[qw(255.255.255.0 255.255.255.0 255.255.255.0
255.255.255.0 255.255.255.0 255.255.255.0 255.255.255.0 255.255.255.0)];

#the user name to register with Red Hat subscription management
$CFG{redhat_subscription_username}="rhel_user";

#the password to register with Red Hat subscription management
$CFG{redhat_subscription_password}="rhel_password";

#clustername of SNAS
$CFG{snas_clustername}="testsnas";

#console IP of SNAS
$CFG{snas_consoleip}="10.200.58.220";

#default gateway of SNAS
$CFG{snas_defgateway}="10.200.58.1";

#domain name of DNS
$CFG{snas_dnsdomainname}="cdc.veritas.com";

#IP of DNS
$CFG{snas_dnship}="10.200.58.3";

#NTP server name
$CFG{snas_ntpserver}="ntp.veritas.com";

#number of VIPs on each NIC
$CFG{snas_nvip}=1;

#netmask of public IPs(only ipv4 environment)
$CFG{snas_pnmaskstart}=255.255.255.0;

#the initial IP of public IPs
$CFG{snas_pipstart}="10.200.58.100";

#if use separate console port, 1 for yes, 0 for no
$CFG{snas_sepconsoleport}="0";

#netmask of virtual IPs(only ipv4 environment)
$CFG{snas_vnmaskstart}=255.255.255.0;

#the initial IP of virtual IPs
```

```

$CFG{snas_vipstart}="10.200.58.108";

#virtual IPs
$CFG{virtualiparr}=[qw(10.200.58.108 10.200.58.109
10.200.58.110 10.200.58.111 10.200.58.112
10.200.58.113 10.200.58.114 10.200.58.115)];

#netmask for virtual IPs
$CFG{virtualnetmaskarr}=[qw(255.255.255.0 255.255.255.0 255.255.255.0
255.255.255.0 255.255.255.0 255.255.255.0 255.255.255.0 255.255.255.0)];

#target systems
$CFG{systems}=[qw(10.200.58.66 10.200.58.82)];

#indicates whether to start llc/gab when user wants to setup a single
node cluster
$CFG{vcs_allowcomms}=1;

#define the unique cluster id with a string number
$CFG{vcs_clusterid}=325;

#define the cluster name with a string
$CFG{vcs_clustername}="testsnas";

#define the nic name for the first heartbeat link.
$CFG{vcs_lltlink1}{"10.200.58.100"}="priveth0";
$CFG{vcs_lltlink1}{"10.200.58.104"}="priveth0";
$CFG{vcs_lltlink2}{"10.200.58.100"}="priveth1";
$CFG{vcs_lltlink2}{"10.200.58.104"}="priveth1";

#define the encrypted user password
$CFG{vcs_userenpw}=[qw(GPQiPKpMQLQQoYQkPN)];

#define the added username for VCS
$CFG{vcs_username}=[qw(admin)];

#define the user privilege
$CFG{vcs_userpriv}=[qw(Administrators)];

1;
#####

```

# Displaying and adding nodes to a cluster

This chapter includes the following topics:

- [About the Veritas Access installation states and conditions](#)
- [Displaying the nodes in the cluster](#)
- [Before adding new nodes in the cluster](#)
- [Adding a node to the cluster](#)
- [Deleting a node from the cluster](#)
- [Shutting down the cluster nodes](#)

## About the Veritas Access installation states and conditions

[Table 8-1](#) describes the Veritas Access installation states.

**Table 8-1** Veritas Access installation states

| Installation state | Description                                                                     |
|--------------------|---------------------------------------------------------------------------------|
| RUNNING            | Node is part of the cluster and the Veritas Access processes are running on it. |
| FAULTED            | Node is down and/or the Veritas Access processes are not running on it.         |
| LEAVING            | Node is leaving the cluster gracefully                                          |

**Table 8-1** Veritas Access installation states (*continued*)

| Installation state | Description                                  |
|--------------------|----------------------------------------------|
| EXITED             | Node has exited the cluster gracefully       |
| UNKNOWN            | Exact state of the node cannot be determined |

Depending on the cluster condition as described in [Table 8-2](#), output for the `Cluster> show` command changes.

**Table 8-2** Cluster conditions and states

| Condition                                                                                                                        | Description                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| If the node is configured and part of the cluster, but the node is powered off.                                                  | State displays as FAULTED, and there is no installation state or network statistics.             |
| If the node is configured and part of the cluster, but the node is physically removed from the cluster.                          | State displays as FAULTED, and there is no installation state or network statistics.             |
| If the node is configured and part of the cluster, but the node is shutdown using the <code>Cluster&gt; shutdown</code> command. | State changes from LEAVING to EXITED.                                                            |
| If the node is configured and part of the cluster, and you use the <code>Cluster&gt; del</code> command.                         | Node is deleted from the cluster, and information about the deleted node is no longer available. |

## Displaying the nodes in the cluster

You can display all the nodes in the cluster, their states, CPU load, and network load during the past 15 minutes.

If you use the `Cluster> show currentload` option, you can display the CPU and network loads collected from now to the next five seconds.

**To display a list of nodes in the cluster**

- 1 To display a list of nodes that are part of a cluster, and the systems that are available to add to the cluster, enter the following:

```
Cluster> show
```

Command output includes the following information. See examples below.

**Node** Displays the node name if the node has already been added to the cluster. Displays the IP address of the node if it is still in the process of being added to the cluster.

Example:

```
node_01
```

or

```
10.200.58.202
```

**State** Displays the state of the node or the installation state of the system along with an IP address of the system if it is installed.

See [“About the Veritas Access installation states and conditions”](#) on page 124.

**CPU** Indicates the CPU load.

**pubethX** Indicates the network load for the Public Interface X.

**bondX** Indicates the network load for bond NIC X.

- 2 For nodes already in the cluster, the following is displayed:

| Node    | State   | CPU(15 min)<br>% | pubeth0(15 min)<br>rx(MB/s) tx(MB/s) | pubeth1(15 min)<br>rx(MB/s) tx(MB/s) |
|---------|---------|------------------|--------------------------------------|--------------------------------------|
| snas_01 | RUNNING | 1.35             | 0.00 0.00                            | 0.00 0.00                            |
| snas_02 | RUNNING | 1.96             | 0.00 0.00                            | 0.00 0.00                            |

- 3 For the nodes that are being added to the cluster, for the nodes that are being deleted from the cluster, and for the nodes that is getting upgraded, the following progress is displayed:

Nodes in Transition

| Node/IP         | Operation       | State   | Description             |
|-----------------|-----------------|---------|-------------------------|
| -----           | -----           | -----   | -----                   |
| 10.200.58.202   | Add node        | FAILED  | Installing packages     |
| snas_03         | Delete node     | ONGOING | Removing node           |
| snas_01,snas_02 | Rolling upgrade | ONGOING | Rolling upgrade phase 2 |

---

**Note:** The `add node` and `delete node` operations cannot be performed at the same time.

---

- 4 To display the CPU and network loads collected from now to the next five seconds, enter the following:

```
Cluster> show currentload
```

Example output:

| Node    | State   | CPU(5 sec)<br>% | pubeth0(5 sec) |          | pubeth1(5 sec) |          |
|---------|---------|-----------------|----------------|----------|----------------|----------|
| ----    | -----   | -----           | rx(MB/s)       | tx(MB/s) | rx(MB/s)       | tx(MB/s) |
| snas_01 | RUNNING | 0.26            | 0.01           | 0.00     | 0.01           | 0.00     |
| snas_02 | RUNNING | 0.87            | 0.01           | 0.00     | 0.01           | 0.00     |
| snas_03 | RUNNING | 10.78           | 27.83          | 12.54    | 0.01           | 0.00     |

Statistics for network interfaces are shown for each public interface available on the cluster nodes.

## Before adding new nodes in the cluster

After you have installed the operating system, you can install and configure a multiple node Veritas Access cluster at one time. If you want to add additional nodes to the cluster after that, you need to complete the following procedures:

- Install the appropriate operating system software on the additional nodes. See “[Installing the operating system on each node of the cluster](#)” on page 72.
- Disable SELinux on the new node.
- You do not need to install the Veritas Access software on the additional node before you add the node. The Veritas Access software is installed when you add the nodes. If the Veritas Access software is already installed, it is uninstalled and the product (same version as the cluster) is installed after that. The reason to uninstall and then install the product is to make sure that the new node is installed with exactly the same version, and patch level (if any) as the other cluster nodes. The packages are stored in the cluster nodes so the product image is not needed during the addition of the new node.
- Verify that the existing cluster has sufficient physical IP addresses for the new nodes. You can add additional IP addresses with the CLISH command: .

```
Network> ip addr add command
```

For example:

```
Network> ip addr add 10.200.58.107 255.255.252.0 physical
ACCESS ip addr SUCCESS V-288-1031 ip addr add successful.
```

```
Network> ip addr show
```

| IP            | Netmask/Prefix | Device  | Node       | Type     | Status          |
|---------------|----------------|---------|------------|----------|-----------------|
| 10.200.58.101 | 255.255.252.0  | pubeth0 | snas_01    | Physical |                 |
| 10.200.58.102 | 255.255.252.0  | pubeth1 | snas_01    | Physical |                 |
| 10.200.58.103 | 255.255.252.0  | pubeth0 | snas_02    | Physical |                 |
| 10.200.58.104 | 255.255.252.0  | pubeth1 | snas_02    | Physical |                 |
| 10.200.58.105 | 255.255.252.0  |         | ( unused ) | Physical |                 |
| 10.200.58.107 | 255.255.252.0  |         | ( unused ) | Physical |                 |
| 10.200.58.231 | 255.255.252.0  | pubeth0 | snas_01    | Virtual  | ONLINE (Con IP) |
| 10.200.58.62  | 255.255.252.0  | pubeth1 | snas_01    | Virtual  | ONLINE          |
| 10.200.58.63  | 255.255.252.0  | pubeth1 | snas_01    | Virtual  | ONLINE          |
| 10.200.58.64  | 255.255.252.0  | pubeth1 | snas_01    | Virtual  |                 |

In the example, the unused IP addresses 10.200.58.105 and 10.200.58.107 can be used by the new node as physical IP addresses.

- If you want to add nodes to a cluster that has RDMA-based LLT links, disable iptables on the cluster nodes using the `service iptables stop` command.

For example:

```
service iptables stop
iptables: Setting chains to policy ACCEPT: filter [OK]
```

```
iptables: Flushing firewall rules: [OK]
iptables: Unloading modules: [OK]
```

---

**Note:** Before proceeding, make sure that all of the nodes are physically connected to the private and public networks.

---

- Add the node to your existing cluster.  
See [“Adding a node to the cluster”](#) on page 129.

## Adding a node to the cluster

The operating system has to be installed on the nodes before you add nodes to a cluster.

If you use disk-based fencing, the coordinator disks must be visible on the newly added node as a prerequisite for I/O fencing to be configured successfully. Without the coordinator disks, I/O fencing will not load properly and the node will not be able to obtain cluster membership.

If you use majority-based fencing, the newly added node doesn't have to have shared disks.

If you want to add a new node and want to exclude some unique PCI IDs, add the unique PCI IDs to the `/opt/VRTSsnas/conf/net_exclusion_dev.conf` file on each cluster node manually. For example:

```
[root@bob_01 ~]# cat /opt/VRTSsnas/conf/net_exclusion_dev.conf
0000:42:00.0 0000:42:00.1
```

---

**Note:** Writeback cache is supported for two-node clusters only, so adding nodes to a two-node cluster changes the caching to read-only.

---



---

**Note:** Newly added nodes should have the same configuration of InfiniBand NICs. See [“About using LLT over the RDMA network for Veritas Access”](#) on page 51.

---

If your cluster has a configured the FSS pool, and the FSS pool's node group is missing a node, then the newly added node is added into the FSS node group, and the installer adds the new node's local data disks into the FSS pool.

### To add the new node to the cluster

- 1 Log in to Veritas Access using the `master` or the `system-admin` account.
- 2 In CLISH, enter the `cluster` command to enter the `cluster>` mode.

**3** To add the new nodes to the cluster, enter the following:

```
Cluster> add node1ip, node2ip....
```

where *node1ip, node2ip, ....* are the IP address list of the additional nodes for the ssh connection.

It is important to note that:

- The node IPs should not be the IPs which are allocated to the new nodes as physical IPs or virtual IPs.
- The physical IPs of new nodes are usable IPs found from the configured public IP starting addresses.
- The virtual IPs are re-balanced to the new node but additional virtual IPs are not assigned.

Go to step 7 to add new virtual IP addresses to the cluster after adding a node.

- The IPs that are accessible to the new nodes should be given.
- The accessible IPs of the new nodes should be in the public network, they should be able to ping the public network's gateway successfully.

For example:

```
Cluster> add 10.200.114.56
```

**4** When you add nodes to a two-node cluster and writeback caching is enabled, the installer asks the following question before adding the node:

```
CPI WARNING V-9-30-2164 Adding a node to a two-node cluster
that has writeback caching enabled will change the caching
to read-only. Writeback caching is only supported for two nodes.
Do you want to continue adding new node(s)? [y,n,q] (n)
```

Enter **y** to continue adding the node. Enter **n** to exit from the add node procedure.

**5** If a cache exists on the original cluster, the installer prompts you to choose the ssd disks to create cache on the new node when CFS is mounted.

```
1) emc_clariion1_242
2) emc_clariion1_243
b) Back to previous menu
Choose disks separate by spaces to create cache on 10.198.89.164
[1-2,b,q] 1
Create cache on snas_02Done
```

- 6** If the cluster nodes have created FSS pool, and there are more than two local data disks on the new node, the installer asks you to select the disks to add into the FSS pool. Make sure that you select at least two disks for stripe volume layout. The total selected disk size should be no less than the FSS pool's capacity size.

Following storage pools need to add disk from the new node:

- 1) fsspool1
- 2) fsspool2
- 3) Skip this step

Choose a pool to add disks [1-3,q] 1

- 1) emc\_clariion0\_1570 (5.000 GB)
- 2) installres\_03\_sdc (5.000 GB)
- 3) installres\_03\_sde (5.000 GB)
- 4) sdd (5.000 GB)
- b) Back to previous menu

Choose at least 2 local disks with minimum capacity of 10 GB [1-4,b,q] 2 4

Format disk installres\_03\_sdc,sdd ..... Done

The disk name changed to installres\_03\_sdc,installres\_03\_sdd

Add disk installres\_03\_sdc,installres\_03\_sdd to storage pool fsspool1 Done

- 7** If required, add the virtual IP addresses to the cluster. Adding the node does not add new virtual IP addresses or service groups to the cluster.

To add additional virtual IP addresses, use the following command in the Network mode:

```
Network> ip addr add ipaddr virtual
```

For example:

```
Network> ip addr add 10.200.58.66 255.255.252.0 virtual
```

```
ACCESS ip addr SUCCESS V-288-1031 ip addr add successful.
```

If a problem occurs while you are adding a node to a cluster (for example, if the node is temporarily disconnected from the network), do the following to fix the problem:

To recover the node:

- Power off the node.

- Use the `Cluster> del nodename` command to delete the node from the cluster.
- Power on the node.
- Use the `Cluster> add nodeip` command to add the node to the cluster.

## Deleting a node from the cluster

This command deletes a node from the cluster. Use the node name that is displayed in the `Cluster> show` command.

---

**Note:** This command is not supported in a single-node cluster.

---

If the deleted node was in the RUNNING state prior to deletion, after you reboot the node, that node is assigned to the original IP address that can be used to add the node back to the cluster. The original IP address of the node is the IP address that the node used before it was added into the cluster.

If your cluster has configured a FSS pool, you cannot use the installer to delete nodes that would result in a single node in the node group of the FSS pool.

Deleting a node from a two-node cluster that has writeback caching enabled changes the caching to read-only. Writeback caching is only supported for two nodes.

The IP address that was used by the node before it was deleted from the cluster is still accessible until you perform a restart operation.

After the node is deleted from the cluster and you perform a reboot operation, the IP addresses associated with the node are free for use by the cluster for new nodes.

### To delete a node from the cluster

- 1 To show the current state of all nodes in the cluster, enter the following:

```
Cluster> show
```

- 2 To delete a node from a cluster, enter the following:

```
Cluster> del nodename
```

where *nodename* is the node name that appeared in the listing from the `Cluster> show` command. You cannot specify a node by its IP address.

For example:

```
Cluster> del snas_01
```

- 3 After a node is deleted from the cluster, the physical IP addresses that it used are marked as unused physical IP addresses. The IP addresses are available for use if you add new nodes. The virtual IP addresses used by a node which has been deleted are not removed. Deleting a node moves the virtual IP addresses on the deleted node to the remaining nodes in the cluster.

For example:

```
Network> ip addr show
```

| IP            | Netmask/Prefix | Device  | Node          | Type     | Status          |
|---------------|----------------|---------|---------------|----------|-----------------|
| 10.209.86.232 | 255.255.252.0  | pubeth0 | source_30a_01 | Physical |                 |
| 10.209.86.233 | 255.255.252.0  | pubeth1 | source_30a_01 | Physical |                 |
| 10.209.86.234 | 255.255.252.0  |         | ( unused )    | Physical |                 |
| 10.209.86.235 | 255.255.252.0  |         | ( unused )    | Physical |                 |
| 10.209.86.240 | 255.255.252.0  | pubeth0 | source_30a_01 | Virtual  | ONLINE (Con IP) |
| 10.209.86.236 | 255.255.252.0  | pubeth0 | source_30a_01 | Virtual  | ONLINE          |
| 10.209.86.237 | 255.255.252.0  | pubeth0 | source_30a_01 | Virtual  | ONLINE          |
| 10.209.86.238 | 255.255.252.0  | pubeth1 | source_30a_01 | Virtual  | ONLINE          |
| 10.209.86.239 | 255.255.252.0  | pubeth1 | source_30a_01 | Virtual  | ONLINE          |

If the physical or virtual IP addresses are not going to be used, they can be removed using the following command:

```
Network> ip addr del ipaddr
```

For example:

```
Network> ip addr del 10.209.86.234
```

```
ACCESS ip addr SUCCESS V-288-1031 ip addr del successful.
```

---

**Note:** If the cluster has configured NIC bonding, you also need to delete the configuration of the deleted node on the switch.

---

## Shutting down the cluster nodes

You can shut down a single node or all of the nodes in the cluster. Use the node name that is displayed in the `Cluster> show` command.

### To shut down a node or all the nodes in a cluster

- 1 To shut down a node, enter the following:

```
Cluster> shutdown nodename
```

*nodename* indicates the name of the node you want to shut down. You cannot specify a node by its IP address.

For example:

```
Cluster> shutdown snas_04
Stopping Cluster processes on snas_04
Sent shutdown command to snas_04. SSH sessions to
snas_04 may terminate.
```

- 2 To shut down all of the nodes in the cluster, enter the following:

```
Cluster> shutdown all
```

Use `all` as the *nodename* to shut down all of the nodes in the cluster.

For example:

```
Cluster> shutdown all
Stopping Cluster processes on all
SSH sessions to all nodes may terminate.
Sent shutdown command to snas_02
Sent shutdown command to snas_03
Sent shutdown command to snas_04
Sent shutdown command to snas_01
```

# Upgrading Veritas Access and operating system

This chapter includes the following topics:

- [Upgrading the operating system and Veritas Access](#)

## Upgrading the operating system and Veritas Access

Veritas Access supports the following upgrade paths for upgrades on RHEL.

**Table 9-1** Supported upgrade paths for upgrades on RHEL

| From product version | From operating system versions | To operating system versions | To product version |
|----------------------|--------------------------------|------------------------------|--------------------|
| 7.2.1.1              | RHEL 6 Update 6, 7, and 8      | RHEL 7 Update 3 and 4        | 7.3.1              |
| 7.3                  | RHEL 6 Update 6, 7, and 8      | RHEL 7 Update 3 and 4        | 7.3.1              |

### Workflow for the LTR upgrade

This process is required for upgrading the Long-term Retention (LTR) on the Veritas Access cluster. All the backup and/or restore jobs from NetBackup must be stopped before you start the upgrade process.

For the LTR upgrade scenarios, you need to use the following scripts:

- `preUpgrade_ltr_access731.py`
- `postUpgrade_ltr_access731.py`

**Before you start the upgrade for LTR:**

- You need to execute the `preUpgrade_ltr_access731.py` script where the `odd_cache_fs` file system is created to backup the cache data of OpenDedup volumes. Size of this file system is determined based on the current cache size (`/opt/sdfs`).
- The pool(s) which is configured as a default pool for Objectaccess to create this file system. Therefore, sufficient space must be available in this pool(s). After the `odd_cache_fs` file system is provisioned, all the OpenDedup volumes are made offline and configuration and the cache data are backed up.

After the upgrade of cluster is completed, you need to execute the `postUpgrade_ltr_access731.py` script where all the OpenDedup volumes are made online after restoring all the configurations.

A one-time tier policy is created for the configured cloud tiers to move the OpenDedup metadata files (ending with `.6442` extension) from the cloud tier to an on-premises storage. OpenDedup needs these metadata files to verify and restore configurations. If these files are stored on a cloud tier, performance of these operations may get degraded.

Upgrading the operating system and Veritas Access includes the following steps:

- Pre-upgrade steps only for LTR configured Veritas Access cluster
- Export the Veritas Access configurations by using the script provided by Veritas Access
- Copy the configuration file
- Install RHEL 7.3 or 7.4
- Install Veritas Access 7.3.1
- Import the Veritas Access configurations
- Verify the imported Veritas Access configurations
- Post-upgrade steps only for LTR configured Veritas Access cluster

**Pre-upgrade steps only for the LTR configured Veritas Access cluster**

---

**Note:** These steps are required when OpenDedup volumes are provisioned on the Veritas Access cluster.

---

- 1 Ensure that the backup and/or restore jobs from NetBackup are stopped.
- 2 From the ISO, copy the `upgrade_scripts/preUpgrade_ltr_access731.py` script to "/" on each node where the OpenDedup volume is online.
- 3 Execute the `preUpgrade_ltr_access731.py` script one-by-one on each node where the OpenDedup volume is online.

See [“Troubleshooting pre-upgrade issues for LTR”](#) on page 155.

### To export the Veritas Access configurations

#### 1 Prerequisite:

Any supported RHEL 6 version is installed.

Veritas Access version 7.2.1.1 or 7.3 is installed.

Make sure that you have stopped all I/Os and services related to Veritas Access by using CLISH, such as CIFS, NFS, FTP, and so on.

- 2 From the ISO, copy the `upgrade_scripts/config_export` directory on "/" on the cluster node on which the ManagementConsole service group is online.
- 3 From the directory, run the following command on Shell (terminal) by using the `root` login to export the Veritas Access configurations:

```
/bin/bash -f export_lib.sh export local <filename>
```

### To verify the Veritas Access configuration export

- ◆ Run the following command on CLISH to see the list of available configurations:

```
system config list
```

The configuration files can be found in

```
/opt/VRTSnas/conf/backup
```

---

**Note:** You need to store these configuration files on a node, which is out of the cluster nodes to avoid any damage to the file.

---

**To install RHEL 7.3 or 7.4**

**1 Prerequisites:**

Make sure that you stop all the running modules on CLISH and no I/O is running.

Run the `network ip addr show` command and `cluster show` command on CLISH before you install RHEL 7. Make a note of these IP addresses and cluster node names. Make sure to use the same IP addresses and cluster name while installing the Veritas Access after the RHEL 7 is installed.

Examples:

```
upgrade> network ip addr show
```

| IP             | Netmask/Prefix | Device  | Node       | Type     | Status          |
|----------------|----------------|---------|------------|----------|-----------------|
| --             | -----          | -----   | ----       | ----     | -----           |
| 192.168.10.151 | 255.255.255.0  | pubeth0 | upgrade_01 | Physical |                 |
| 192.168.10.158 | 255.255.255.0  | pubeth1 | upgrade_01 | Physical |                 |
| 192.168.10.152 | 255.255.255.0  | pubeth0 | upgrade_02 | Physical |                 |
| 192.168.10.159 | 255.255.255.0  | pubeth1 | upgrade_02 | Physical |                 |
| 192.168.10.174 | 255.255.255.0  | pubeth0 | upgrade_01 | Virtual  | ONLINE (Con IP) |
| 192.168.10.160 | 255.255.255.0  | pubeth0 | upgrade_01 | Virtual  | ONLINE          |
| 192.168.10.161 | 255.255.255.0  | pubeth1 | upgrade_01 | Virtual  | ONLINE          |

```
upgrade> cluster show
```

| Node       | State   | CPU (15 min) | pubeth0 (15 min) |           | pubeth1 (15 min) |           |
|------------|---------|--------------|------------------|-----------|------------------|-----------|
|            |         |              | rx (MB/s)        | tx (MB/s) | rx (MB/s)        | tx (MB/s) |
| ----       | -----   | -----        | -----            | -----     | -----            | -----     |
| upgrade_01 | RUNNING | 11.52        | 0.67             | 0.06      | 0.60             | 0.00      |
| upgrade_02 | RUNNING | 4.19         | 0.61             | 0.05      | 0.60             | 0.00      |

---

**Note:** In this example, cluster name is `upgrade` and cluster node names are `upgrade_01` and `upgrade_02`.

---

**2 Restart all the nodes of the cluster.**

- 3 Install RHEL 7.3 or 7.4 on the desired nodes.

See [“Installing the operating system on the target Veritas Access cluster”](#) on page 73.

---

**Note:** It is recommended to select the same disk or disks for the installation on which RHEL 6 was installed. Make sure that you do not to select any other disk as it may be possible that those disks are part of a pool, and result in data loss.

---

### To install Veritas Access 7.3.1

- ◆ After a restart when the nodes are up, start the Veritas Access 7.3.1 installation by using CPI.

---

**Note:** Make sure to use the same IP addresses and cluster name which were used for the Veritas Access on RHEL 6.

---

See [“Installing Veritas Access on the target cluster nodes”](#) on page 78.

### To verify the Veritas Access installation

- 1 By using the console IP, check whether the CLISH is accessible.
- 2 Run the following command in CLISH to see whether the disks are accessible:

```
storage disk list
```

---

**Note:** If the disks are not visible in the CLISH output, run the `storage scanbus force` command in CLISH.

---

- 3 Run the following command in CLISH to see whether the pools are accessible:

```
storage pool list
```

---

**Note:** If the pools are not visible in the CLISH output, run the `storage scanbus force` command in CLISH.

---

- 4 Run the following command in CLISH to see whether the file systems are accessible:

```
storage fs list
```

---

**Note:** If the file systems are not visible in the CLISH output, run the `storage scanbus force` command in CLISH.

---

- 5 Make sure that the file systems are online. If the file systems are not online, you need to run the following command in CLISH to bring them online:

```
storage fs online <fs name>
```

## To import the Veritas Access configuration

### 1 Prerequisites:

Make sure that the file systems are online. If the file systems are not online, you need to run the following command in CLISH to bring them online:

```
storage fs online <fs name>
```

---

**Note:** Make sure that the cluster uses the same IP addresses and cluster name which were used for the Veritas Access on RHEL 6.

If VIP addresses are not added during installation, which were used for Veritas Access on RHEL 6, add them from CLISH after Veritas Access is installed on RHEL 7, and then import the configuration.

---

- 2 Copy the exported configuration files to the cluster nodes in the following location:

```
/opt/VRTSnas/conf/backup/
```

- 3 Run the following command in CLISH to see the available exported configuration:

```
system config list
```

- 4 Log in to CLISH and import the module configuration by using the following command:

```
system config import local <config-filename> <module-to-import>
```

The following modules can be imported:

```
upgrade> system config import local
```

```
system config import local <file_name> [config-type]
-- Import the configuration which is stored locally
```

```
file_name : configuration file name
config-type : input type of configuration to import (network/admin/all/report/
system/support/cluster_specific/all_except_cluster_specific/nfs/cifs/ftp/backup/
replication/storage_schedules/storage_quota/storage_fs_alert/storage_fs_policy/
compress_schedules/defrag_schedules/storage_dedup/smartio/target/object_access/
loadbalance/openedup) [all]
```

```
upgrade> system config import local
```

---

**Note:** The module names are auto-suggested in CLISH.

---

### Post-upgrade steps only for the LTR configured Veritas Access cluster

---

**Note:** These steps are required when the OpenDedup volumes are provisioned on Veritas Access cluster.

---

- 1 From the ISO, copy the `upgrade_scripts/postUpgrade_ltr_access731.py` script to `/` on the Management Console node.
- 2 Execute the `postUpgrade_ltr_access731.py` script.

See [“Troubleshooting post-upgrade issues for LTR”](#) on page 156.

# Upgrading Veritas Access using a rolling upgrade

This chapter includes the following topics:

- [About rolling upgrades](#)
- [Supported rolling upgrade paths for upgrades on RHEL](#)
- [Performing a rolling upgrade using the installer](#)

## About rolling upgrades

This release of Veritas Access supports rolling upgrades from the Veritas Access 7.2.1.1 and later versions. Rolling upgrade is supported on RHEL 6.6, 6.7, and 6.8.

A rolling upgrade minimizes the service and application downtime for highly available clusters by limiting the upgrade time to the amount of time that it takes to perform a service group failover. Nodes with different product versions can be run in one cluster.

The rolling upgrade has two main phases. The installer upgrades kernel RPMs in phase 1 and VCS agent RPMs in phase 2. The upgrade process divides the cluster into two subclusters, called the first subcluster and the second subcluster. First, the upgrade is performed on the first subcluster. The upgrade process stops all services and resources on the nodes of the first subcluster. All services (including the VIP groups) fail over to the second subcluster. During the failover process, the clients that are connected to the VIP groups of the first subcluster nodes are intermittently interrupted. For those clients that do not time out, the service is resumed after the VIP groups become online on one of the nodes of the second subcluster.

While the upgrade process is running on the nodes of the first subcluster, the nodes of the second subcluster nodes continue to serve the clients. After the first subcluster node has been upgraded, it restarts the services and resources on first stage nodes. Immediately after the first subcluster comes up, the upgrade process stops the services and resources on the remaining nodes. All services and resources are online and serve clients. Meanwhile, the rolling upgrade starts the upgrade process on the remaining nodes. After the upgrade is complete on the remaining nodes, the cluster recovers and services are balanced across the cluster.

## Workflow for the LTR upgrade

This process is required for upgrading Long-term Retention (LTR) on the Veritas Access cluster. All the backup and/or restore jobs from NetBackup must be stopped before you start the rolling upgrade process.

For the LTR upgrade scenarios, you need to use the following scripts:

- `preUpgrade_ltr_access731.py`
- `postUpgrade_ltr_access731.py`

### Before you start the rolling upgrade for LTR:

- You need to execute the `preUpgrade_ltr_access731.py` script where the `odd_cache_fs` file system is created to backup the cache data of OpenDedup volumes. Size of this file system is determined based on the current cache size (`/opt/sdfs`).
- The pool(s) which is configured as a default pool for Objectaccess to create this file system. Therefore, sufficient space must be available in this pool(s). After the `odd_cache_fs` file system is provisioned, all the OpenDedup volumes are made offline and configuration and the cache data are backed up.

After the rolling upgrade of cluster is completed, you need to execute the `postUpgrade_ltr_access731.py` script where all the OpenDedup volumes are made online after restoring all the configurations.

A one-time tier policy is created for the configured cloud tiers to move the OpenDedup metadata files (ending with `.6442` extension) from the cloud tier to an on-premises storage. OpenDedup need these metadata files to verify and restore configurations. If these files are stored on a cloud tier, performance of these operations may get degraded.

## Workflow for a rolling upgrade

A rolling upgrade has two main phases where the installer upgrades the kernel RPMs in Phase 1 and VCS agent-related non-kernel RPMs in Phase 2.

1. Use the following pre-upgrade steps only for LTR configured Veritas Access cluster.

---

**Note:** These LTR related steps are required when OpenDedup volumes are provisioned on Veritas Access cluster.

---

- Ensure that the backup and/or restore jobs from NetBackup are stopped.
- From the ISO, copy the `upgrade_scripts/preUpgrade_ltr_access731.py` script to `/` on each node where the OpenDedup volume is online.
- Execute the `preUpgrade_ltr_access731.py` script one-by-one on each node where the OpenDedup volume is online.

See [“Troubleshooting pre-upgrade issues for LTR”](#) on page 155.

2. The upgrade process divides the cluster into two subclusters, first subcluster and the second subcluster.
3. In Phase 1, the upgrade is performed on the second subcluster. The upgrade process stops all services and resources on the nodes of the second subcluster. All services (including the VIP groups) failover to the first subcluster. The parallel service groups on the second subcluster are taken offline.
4. During the failover process, the clients that are connected to the VIP groups of the second subcluster nodes are intermittently interrupted. For those clients that do not time out, the service is resumed after the VIP groups become online on one of the nodes of the first subcluster.
5. The installer upgrades the kernel RPMs on the second subcluster. The nodes of the first subcluster nodes continue to serve the clients.
6. Once Phase 1 of the rolling upgrade is complete on the second subcluster, Phase 1 of the rolling upgrade is performed on the first subcluster. The applications are failed over to the second subcluster. The parallel service groups are brought online on the second subcluster and are taken offline on the first subcluster.
7. After Phase 1 is complete, the nodes run with new RPMs but with the old protocol version.
8. During Phase 2 of the rolling upgrade, all remaining RPMs are upgraded on all the nodes of the cluster simultaneously. VCS and VCS agent packages are upgraded. The kernel drivers are upgraded to the new protocol version. Applications stay online during Phase 2. The High Availability Daemon (HAD) stops and starts again.
9. Use the following post-upgrade steps only for LTR configured Veritas Access cluster.

---

**Note:** These steps are required when the OpenDedup volumes are provisioned on Veritas Access cluster.

---

- From the ISO, copy the `upgrade_scripts/postUpgrade_ltr_access731.py` script to `/` on the Management Console node.
- Execute the `postUpgrade_ltr_access731.py` script.

See [“Performing a rolling upgrade using the installer”](#) on page 145.

See [“Supported rolling upgrade paths for upgrades on RHEL”](#) on page 145.

See [“Troubleshooting post-upgrade issues for LTR”](#) on page 156.

## Supported rolling upgrade paths for upgrades on RHEL

**Table 10-1** Supported upgrade paths for upgrades on RHEL

| From product version | Operating system versions | To product version |
|----------------------|---------------------------|--------------------|
| 7.2.1.1              | RHEL 6 Update 6, 7, and 8 | 7.3.1              |
| 7.2.0.2              | RHEL 6 Update 6, 7, and 8 | 7.3.1              |
| 7.3                  | RHEL 6 Update 6, 7, and 8 | 7.3.1              |
| 7.3.0.1              | RHEL 7 Update 3           | 7.3.1              |

---

**Note:** See the "Known issues" section of the *Veritas Access Release Notes* before starting a rolling upgrade for other product versions not shown in this table.

---

See [“Performing a rolling upgrade using the installer”](#) on page 145.

See [“About rolling upgrades”](#) on page 142.

## Performing a rolling upgrade using the installer

Before you start a rolling upgrade, make sure that Cluster Server (VCS) is running on all the nodes of the cluster.

Stop all activity for all the VxVM volumes that are not under VCS control. For example, stop any applications such as databases that access the volumes, and

unmount any file systems that have been created on the volumes. Then stop all the volumes.

Unmount all the VxFS file systems that are not under VCS control.

### To perform a rolling upgrade

- 1 Phase 1 of a rolling upgrade begins on the second subcluster. Complete the preparatory steps on the second subcluster.

Unmount all VxFS file systems not under VCS control:

```
umount mount_point
```

- 2 Complete updates to the operating system, if required.

Make sure that the existing version of Veritas Access supports the operating system update you apply. If the existing version of Veritas Access does not support the operating system update, first upgrade Veritas Access to a version that supports the operating system update.

For instructions, see the Red Hat Enterprise Linux (RHEL) operating system documentation.

Switch applications to the remaining subcluster and upgrade the operating system of the first subcluster.

The nodes are restarted after the operating system update.

- 3 If a cache area is online, you must take the cache area offline before you upgrade the VxVM RPMs. Use the following command to take the cache area offline:

```
sfcache offline cachename
```

- 4 Log on as superuser and mount the Veritas Access 7.3.1 installation media.

- 5 From root, start the installer.

```
./installaccess -rolling_upgrade
```

- 6 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. The installer asks for permission to proceed with the rolling upgrade.

```
Would you like to perform rolling upgrade on the cluster? [y,n,q] (y)
```

Type **y** to continue.

- 7 Phase 1 of the rolling upgrade begins. The installer suggests system names for the upgrade. Press **y** to upgrade the suggested systems, or enter the name of any one system in the cluster on which you want to perform a rolling upgrade.
- 8 The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type **y** to continue. If you choose to specify the nodes, type **n** and enter the names of the nodes.
- 9 The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type **y** to continue or quit the installer and address the precheck's warnings.
- 10 If the boot disk is encapsulated and mirrored, you can create a backup boot disk.

If you choose to create a backup boot disk, type **y**. Provide a backup name for the boot disk group or accept the default name. The installer then creates a backup copy of the boot disk group.

- 11 After the installer detects the online service groups, the installer prompts the user to do one of the following:
  - Manually switch service groups
  - Use the CPI to automatically switch service groups

The downtime is the time that it takes for the failover of the service group.

---

**Note:** Veritas recommends that you manually switch the service groups. Automatic switching of service groups does not resolve dependency issues.

---

- 12 The installer prompts you to stop the applicable processes. Type **y** to continue.
 

The installer evacuates all service groups to the node or nodes that are not upgraded at this time. The installer stops parallel service groups on the nodes that are to be upgraded.

The installer stops relevant processes, uninstalls old kernel RPMs, and installs the new RPMs.
- 13 The installer performs the upgrade configuration and starts the processes. If the boot disk is encapsulated before the upgrade, the installer prompts you to restart the node after performing the upgrade configuration.
- 14 Complete the preparatory steps on the nodes that you have not yet upgraded.
 

Unmount all the VxFS file systems not under VCS control on all the nodes.

```
umount mount_point
```

- 15** If operating system updates are not required, skip this step.

Go to step 3.

Else, complete updates to the operating system on the nodes that you have not yet upgraded. For instructions, see the Red Hat Enterprise Linux (RHEL) operating system documentation.

Repeat steps 3 to 13 for each node.

Phase 1 of the rolling upgrade is complete on the first subcluster. Phase 1 of the rolling upgrade begins on the second subcluster.

- 16** Offline all cache areas on the remaining node or nodes:

```
sfcache offline cachename
```

- 17** The installer begins phase 1 of the upgrade on the remaining node or nodes. Type **y** to continue the rolling upgrade. If the installer was invoked on the upgraded (restarted) nodes, you must invoke the installer again.

The installer repeats step 8 through step 13.

For clusters with larger number of nodes, this process may repeat several times. Service groups come down and are brought up to accommodate the upgrade.

- 18** When phase 1 of the rolling upgrade completes, mount all the VxFS file systems that are not under VCS control manually. Begin Phase 2 of the upgrade. phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue. Phase 2 of the rolling upgrade begins here.

- 19** The installer determines the remaining RPMs to upgrade. Press **y** to continue.

- 20** The installer stops Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.

The installer performs a prestop, uninstalls the old RPMs, and installs the new RPMs. It performs post-installation tasks, and the configuration for the upgrade.

- 21** If you have a network connection to the Internet, the installer checks for updates.

If updates are discovered, you can apply them now.

- 22** Verify the cluster's status:

```
hastatus -sum
```

See [“Supported rolling upgrade paths for upgrades on RHEL”](#) on page 145.

See [“About rolling upgrades”](#) on page 142.

# Uninstalling Veritas Access

This chapter includes the following topics:

- [Before you uninstall Veritas Access](#)
- [Uninstalling Veritas Access using the installer](#)

## Before you uninstall Veritas Access

Perform the following steps before uninstalling Veritas Access:

- Before you remove Veritas Access from any node (but not in all the nodes) in a cluster, make sure the node has already been deleted from the running cluster. You can use the `Cluster> show` command to view the cluster node state, and use the `Cluster> delete` command to delete a running node from the Veritas Access cluster.  
See the relevant man pages for more information on the `Cluster> show` and `Cluster> delete` commands.
- Stop all the applications that access the file system over NFS, CIFS, or FTP.
- Destroy all the replication jobs from the cluster.  
Use the `Replication> job show` command to list all the replication jobs on the cluster.

```
Replication> job show
Job Name Role Job Type Encryption Debug Schedule
=====
job1 SOURCE DATA OFF ON sch1
State CKPT Count Exclunit Source repunit Target repunit(s)
=====
```

```
ENABLED 1 -- scr1 trg1
Link name(s)
=====
link1
```

Use the `Replication> job destroy` command to destroy the replication jobs.

```
Replication> job destroy job1
ACCESS replication SUCCESS V-288-0 Removing bandwidth limit on the
link: link1
ACCESS replication SUCCESS V-288-0 Job 'job1' disabled successfully.
ACCESS replication SUCCESS V-288-0 Job 'job1' deleted successfully.
```

- Stop the NFS, CIFS, FTP, GUI, and the replication service on the cluster using the appropriate CLISH command.

```
CLISH> cifs server stop
Stopping CIFS Server.....Success.
CLISH>
CLISH> nfs server stop
Success.
CLISH>
CLISH> ftp server stop
Success.
CLISH>
CLISH.Support> gui server stop
GUI service is OFFLINE.
CLISH>
CLISH> replication service stop
ACCESS replication SUCCESS V-288-0 Replication service stopped
CLISH>
```

- Run the following command to stop AMF:

```
/etc/init.d/amf stop
Stopping AMF...
AMF: Module unloaded
```

- Run the following command and wait for a couple of minutes:

```
/opt/VRTS/bin/hastop -all
```

- Run the following command and verify that you only see Port a and Port b:

```
gabconfig -a
GAB Port Memberships
=====
Port a gen 7f2d0a membership 01
Port b gen 7f2d09 membership 01
```

## Uninstalling Veritas Access using the installer

You can perform an uninstallation of Veritas Access. The Veritas Access uninstall program lets you uninstall Veritas Access without requiring a reinstallation of the operating system. You can also use the uninstall program in cases where there was an incomplete installation of Veritas Access.

Before you use the uninstall program to uninstall Veritas Access on all nodes in the cluster at the same time, make sure that communication exists between the nodes. By default, Veritas Access cluster nodes can communicate with each other using ssh.

If the nodes cannot communicate with each other, then you must run the uninstall program on each node in the cluster. The uninstall program removes all Veritas Access RPMs.

### Removing Veritas Access 7.3.1 RPMs

The uninstall program stops the Veritas Access processes that are currently running during the uninstallation process.

#### To uninstall Veritas Access 7.3.1 RPMs

- 1 Log in as the support user from the node where you want to uninstall Veritas Access.
- 2 Start the uninstall program.

```
cd /opt/VRTS/install
./uninstallaccess
```

The program specifies the directory where the logs are created. The program displays a copyright notice and a description of the cluster.

- 3 Enter the IP addresses of the nodes from which you want to uninstall Veritas Access.

The program performs node verification checks and asks to stop all running Veritas Access processes.

- 4 Enter **y** to stop all the Veritas Access processes.

The program stops the Veritas Access processes and uninstalls the software.

The uninstall program does the following tasks:

- Verifies the communication between nodes.
- Checks the installations on each node to determine the RPMs to be uninstalled.
- Unloads kernel modules and removes the RPMs.

Review the output as the uninstaller stops processes.

You can make a note of the location of the summary, response, and log files that the uninstaller creates after removing all the RPMs.

## Running uninstall from the Veritas Access 7.3.1 disc

You may need to use the uninstall program on the Veritas Access 7.3.1 disc in one of the following cases:

- You need to uninstall Veritas Access after an incomplete installation.
- The uninstall program is not available in `/opt/VRTS/install`.

If you mounted the installation media to `/mnt`, access the uninstall program by changing the directory.

```
cd /mnt/
```

```
./uninstallaccess
```

# Installation reference

This appendix includes the following topics:

- [Installation script options](#)

## Installation script options

[Table A-1](#) lists the available command line options for the Veritas Access installation script. For an initial install or upgrade, options are not usually required.

**Table A-1** Available command line options

| Command Line Option | Function                                                                                                                                   |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| -configure          | Configures an unconfigured product after it is installed.                                                                                  |
| -install            | Installs the product on systems.                                                                                                           |
| -precheck           | Performs checks to confirm that systems have met the products installation requirements before installing the product.                     |
| -license            | Registers or updates product licenses on the specified systems.                                                                            |
| -licensefile        | Specifies the location of the Veritas perpetual or subscription license key file.                                                          |
| -requirements       | Displays the required operating system version, required patches, file system space, and other system requirements to install the product. |

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

| Command Line Option                       | Function                                                                                                                                                                                                                    |
|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -responsefile <i>response_file</i>        | Performs automated installations or uninstalls using information stored in a file rather than prompting for the information. <i>response_file</i> is the full path of the file that contains the configuration definitions. |
| -rolling_upgrade                          | Performs 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.                 |
| -prestop_script <i>prestop_script</i>     | Executes the customized script provided by user on each host before stop processes during the upgrade procedure.                                                                                                            |
| -poststart_script <i>poststart_script</i> | Executes the customized script provided by user on each host after start processes during the upgrade procedure.                                                                                                            |
| -uninstall                                | Uninstalls the product from systems.                                                                                                                                                                                        |
| -updateparameter                          | Updates the network parameter for a running cluster.                                                                                                                                                                        |

# Troubleshooting the LTR upgrade

This appendix includes the following topics:

- [Locating the log files for troubleshooting the LTR upgrade](#)
- [Troubleshooting pre-upgrade issues for LTR](#)
- [Troubleshooting post-upgrade issues for LTR](#)

## Locating the log files for troubleshooting the LTR upgrade

Log files for the LTR upgrade process are located at:

- `/opt/VRTSnas/log/ltr_preupgrade.log`
- `/opt/VRTSnas/log/ltr_post_upgrade.log`

## Troubleshooting pre-upgrade issues for LTR

Following are the pre-upgrade scenarios with the recommended workaround:

- Failure in provisioning of `odd_cache_fs`  
**Workaround:** Ensure that sufficient storage space is available in default pool(s) that is configured for Objectaccess.
- Failure in backup of configuration file  
**Workaround:** Ensure that the `odd_vipgrp_map.conf` file is successfully copied to `/shared/openedup/`.
- Failure in backup of XML configuration files

**Workaround:** Ensure that all the OpenDedup volume configuration XML files are successfully copied to `/shared/opendedup/sdfs/`.

- Failure in shutdown of OpenDedup volume(s)

**Workaround:**

- Ensure that the volume is unmounted by using the `mount -t fuse` command.
- Ensure all the `jsvc` processes are stopped. If the processes are not stopped, use the `kill` command to stop them.
- Shutdown of the OpenDedup volume may result in java stack trace to be printed on console. This can be ignored as long as the volume is unmounted.

- Failure in backup of the OpenDedup cache data

**Workaround:** Ensure that the size of the local cache location, that is, `/opt/sdfs/volumes/<vol-name>` is equivalent to the shared cache location, that is, `/vx/odd_cache_fs/volumes/<vol-name>`.

## Troubleshooting post-upgrade issues for LTR

Following are the post-upgrade scenarios with the recommended workaround:

- Failure in creation of soft-links

**Workaround:** Ensure that the soft links are created for the following paths:

- `/etc/sdfs -> /shared/opendedup/sdfs`
- `/opt/sdfs -> /vx/odd_cache_fs/sdfs`

- Failure in configuration restore

**Workaround:** Ensure that the `odd_vipgrp_map.conf` file has been successfully copied from `/shared/opendedup/` to `/opt/VRTSnas/conf/`.

- Failure in online of OpenDedup volumes

**Workaround:**

- Log locations of the OpenDedup volumes:

- `/opt/VRTSnas/log/odd.log`
- `/var/log/sdfs/<vol-name>.log`

- Manual mount of OpenDedup volume can be verified by using the following CLISH commands:

- `Opendedup volume list`
- `Opendedup volume online <vol-name>`

# Configuring the secure shell for communications

This appendix includes the following topics:

- [Manually configuring passwordless secure shell \(ssh\)](#)
- [Setting up ssh and rsh connections using the `pwdutil.pl` utility](#)

## Manually configuring passwordless secure shell (ssh)

The secure shell (ssh) program enables you to log into and execute commands on a remote system. ssh enables encrypted communications and an authentication process between two untrusted hosts over an insecure network.

In this procedure, you first create a DSA key pair. From the key pair, you append the public key from the source system to the `authorized_keys` file on the target systems.

**To create the DSA key pair**

- 1 On the source system (sys1), log in as **root**, and navigate to the root directory.

```
sys1 # cd /root
```

- 2 To generate a DSA key pair on the source system, type the following command:

```
sys1 # ssh-keygen -t dsa
```

System output similar to the following is displayed:

```
Generating public/private dsa key pair.
Enter file in which to save the key (/root/.ssh/id_dsa):
```

- 3 Press Enter to accept the default location of `/root/.ssh/id_dsa`.
- 4 When the program asks you to enter the pass phrase, press the Enter key twice.

```
Enter passphrase (empty for no passphrase):
```

Do not enter a pass phrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

- 5 Output similar to the following lines appears.

```
Your identification has been saved in /root/.ssh/id_dsa.
Your public key has been saved in /root/.ssh/id_dsa.pub.
The key fingerprint is:
1f:00:e0:c2:9b:4e:29:b4:0b:6e:08:f8:50:de:48:d2 root@sys1
```

**To append the public key from the source system to the authorized\_keys file on the target system using secure file transfer**

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

Use the secure file transfer program.

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

Use the following command for secure file transfer:

```
sys1 # sftp sys2
```

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

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

- 2 Enter `yes`.

Output similar to the following is displayed:

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

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

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

The following output is displayed:

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

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

```
sftp> quit
```

- 6 Add the `id_dsa.pub` keys to the `authorized_keys` file on the target system. To begin the `ssh` session on the target system (`sys2` in this example), type the following command on `sys1`:

```
sys1 # ssh sys2
```

Enter the root password of `sys2` at the prompt:

```
password:
```

Type the following commands on `sys2`:

```
sys2 # cat /root/id_dsa.pub >> /root/.ssh/authorized_keys
sys2 # rm /root/id_dsa.pub
```

- 7 Run the following commands on the source installation system. If your `ssh` session has expired or terminated, you can also run these commands to renew the session. These commands bring the private key into the shell environment and make the key globally available to the `root` user.

```
sys1 # exec /usr/bin/ssh-agent $SHELL
sys1 # ssh-add
```

```
Identity added: /root/.ssh/id_dsa
```

This shell-specific step is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

### To verify that you can connect to a target system

- 1 On the source system (`sys1`), enter the following command:

```
sys1 # ssh -l root sys2 uname -a
```

where `sys2` is the name of the target system.

- 2 The command should execute from the source system (`sys1`) to the target system (`sys2`) without the system requesting a pass phrase or password.
- 3 Repeat this procedure for each target system.

## Setting up ssh and rsh connections using the `pwdutil.pl` utility

The password utility, `pwdutil.pl`, is bundled in the 7.3 release in the `/opt/VRTS/repository/ga/images/SSNAS/7.3.0.0/scripts/pwdutil.pl`

directory. The users can run the utility in their script to set up the ssh and rsh connection automatically.

```
./pwdutil.pl -h
```

Usage:

Command syntax with simple format:

```
pwdutil.pl check|configure|unconfigure ssh|rsh <hostname|IP addr>
[<user>] [<password>] [<port>]
```

Command syntax with advanced format:

```
pwdutil.pl [--action|-a 'check|configure|unconfigure']
 [--type|-t 'ssh|rsh']
 [--user|-u '<user>']
 [--password|-p '<password>']
 [--port|-P '<port>']
 [--hostfile|-f '<hostfile>']
 [--keyfile|-k '<keyfile>']
 [-debug|-d]
 <host_URI>
```

```
pwdutil.pl -h | -?
```

**Table C-1** Options with pwdutil.pl utility

| Option                                    | Usage                                                   |
|-------------------------------------------|---------------------------------------------------------|
| --action -a 'check configure unconfigure' | Specifies action type, default is 'check'.              |
| --type -t 'ssh rsh'                       | Specifies connection type, default is 'ssh'.            |
| --user -u '<user>'                        | Specifies user id, default is the local user id.        |
| --password -p '<password>'                | Specifies user password, default is the user id.        |
| --port -P '<port>'                        | Specifies port number for ssh connection, default is 22 |
| --keyfile -k '<keyfile>'                  | Specifies the private key file.                         |
| --hostfile -f '<hostfile>'                | Specifies the file which list the hosts.                |
| -debug                                    | Prints debug information.                               |

**Table C-1** Options with `pwdutil.pl` utility (*continued*)

| Option                        | Usage                                                                                                                                                                                                                         |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>-h -?</code>            | Prints help messages.                                                                                                                                                                                                         |
| <code>&lt;host_URI&gt;</code> | Can be in the following formats:<br><code>&lt;hostname&gt;</code><br><code>&lt;user&gt;:&lt;password&gt;@&lt;hostname&gt;</code><br><code>&lt;user&gt;:&lt;password&gt;@&lt;hostname&gt;:</code><br><code>&lt;port&gt;</code> |

You can check, configure, and unconfigure ssh or rsh using the `pwdutil.pl` utility. For example:

- To check ssh connection for only one host:
 

```
pwdutil.pl check ssh hostname
```
- To configure ssh for only one host:
 

```
pwdutil.pl configure ssh hostname user password
```
- To unconfigure rsh for only one host:
 

```
pwdutil.pl unconfigure rsh hostname
```
- To configure ssh for multiple hosts with same user ID and password:
 

```
pwdutil.pl -a configure -t ssh -u user -p password hostname1
hostname2 hostname3
```
- To configure ssh or rsh for different hosts with different user ID and password:
 

```
pwdutil.pl -a configure -t ssh user1:password1@hostname1
user2:password2@hostname2
```
- To check or configure ssh or rsh for multiple hosts with one configuration file:
 

```
pwdutil.pl -a configure -t ssh --hostfile /tmp/sshrsh_hostfile
```
- To keep the host configuration file secret, you can use the 3rd party utility to encrypt and decrypt the host file with password.  
 For example:

```
run openssl to encrypt the host file in base64 format
openssl aes-256-cbc -a -salt -in /hostfile -out /hostfile.enc
enter aes-256-cbc encryption password: <password>
Verifying - enter aes-256-cbc encryption password: <password>

remove the original plain text file
rm /hostfile

run openssl to decrypt the encrypted host file
pwdutil.pl -a configure -t ssh `openssl aes-256-cbc -d -a
-in /hostfile.enc`
enter aes-256-cbc decryption password: <password>
```

- To use the ssh authentication keys which are not under the default `$HOME/.ssh` directory, you can use `--keyfile` option to specify the ssh keys. For example:

```
create a directory to host the key pairs:
mkdir /keystore

generate private and public key pair under the directory:
ssh-keygen -t rsa -f /keystore/id_rsa

setup ssh connection with the new generated key pair under
the directory:
pwdutil.pl -a configure -t ssh --keyfile /keystore/id_rsa
user:password@hostname
```

You can see the contents of the configuration file by using the following command:

```
cat /tmp/sshrsh_hostfile
user1:password1@hostname1
user2:password2@hostname2
user3:password3@hostname3
user4:password4@hostname4

all default: check ssh connection with local user
hostname5
The following exit values are returned:

0 Successful completion.
1 Command syntax error.
2 Ssh or rsh binaries do not exist.
3 Ssh or rsh service is down on the remote machine.
```

```
4 Ssh or rsh command execution is denied due to password is required.
5 Invalid password is provided.
255 Other unknown error.
```

# Index

## Symbols

/etc/lfttab  
sample 54

## A

about  
NIC bonding 85  
NIC exclusion 85  
VLAN Tagging 104  
adding  
VLAN device on a particular NIC 104

## B

bond  
creating 94  
bond interface  
creating 94

## C

calculating  
IP addresses 58  
checking  
storage configuration 62  
cluster  
adding the new node to 129  
deleting a node from 132  
displaying a list of nodes 125  
including new nodes 127  
shutting down a node or all nodes in a cluster 134  
cluster installation  
overview 70  
Configuration  
LLT 53  
configuration limits 47  
configuring  
NetBackup (NBU) 107  
Veritas Access software on the cluster 79  
configuring passwordless ssh 157  
connecting  
network hardware 55

## D

deleting  
a node from the cluster 132  
disabling  
iptables rules 45  
displaying  
list of nodes in a cluster 125  
driver node 72

## E

excluding  
NIC 86

## H

Hardware requirements  
Veritas Access 51

## I

including  
new nodes in the cluster 127  
NIC 90  
install  
silent 112  
installation  
response files 111  
response files variables 112  
installation script options 153  
installation states and conditions  
about 124  
installation time  
reducing the number of IP addresses 61  
Installer  
configure 54  
installing  
cluster 70  
configuring the Veritas Access software on the cluster 79  
operating system on each node of the cluster 72  
operating system on Veritas Access cluster 73  
Oracle Linux operating system 75

- installing (*continued*)
    - prerequisites 71
    - steps 70
    - target cluster nodes 78
  - IP addresses
    - calculate 94
    - calculating 58
    - obtain 57
  - IPv6 protocol 41
- K**
- kernel RPMs
    - OL 22
- L**
- limitations of
    - VLAN Tagging 105
  - Linux requirements
    - Veritas Access 17
  - list of nodes
    - displaying in a cluster 125
  - LLT
    - RDMA 52, 54
  - locating the log files
    - troubleshooting the LTR upgrade 155
- M**
- Management Server requirements
    - Veritas Access 40
- N**
- NetBackup (NBU)
    - configuring 107
  - network and firewall requirements
    - Veritas Access 42
  - network hardware
    - connecting 55
  - network interface card (NIC) bonding 94
  - NIC
    - excluding 86
    - including 90
  - node
    - adding to the cluster 127, 129
- O**
- obtain
    - IP addresses 57
  - obtaining
    - Red Hat Enterprise Linux compatible kernels 76
  - OL kernel RPMs 22
  - OpenDedup ports
    - disabling the iptable rules 45
  - operating system
    - installing 73
    - installing on each node of the cluster 72
    - upgrade 135
  - Oracle Linux
    - installing operating system 75
  - overview
    - Veritas Access installation 49
- R**
- RDMA
    - Hardware 53
    - InfiniBand 52
    - LLT 51
  - reconfiguring
    - Veritas Access cluster name and network 108
  - Red Hat Enterprise Linux compatible kernels
    - obtaining 76
  - reducing
    - number of IP addresses required at installation time 61
  - release information 16
  - removing
    - bond 99
    - NIC from bond list 102
  - replacing
    - Ethernet interface card 105
- S**
- sample response file 121
  - shutting down
    - node or all nodes in a cluster 134
  - silent installation and configuration 112
  - storage configuration
    - checking 62
  - supported IPv6 protocol 41
  - supported rolling upgrade paths
    - upgrades on RHEL and Oracle Linux 145
  - system requirements
    - Veritas Access 16

**T**

troubleshooting the LTR upgrade  
  locating the log files 155

**U**

uninstalling Veritas Access  
  before 149  
upgrade  
  operating system 135  
  Veritas Access application 135  
upgrades on RHEL and Oracle Linux  
  supported rolling upgrade paths 145

**V**

Veritas Access  
  about 8  
  key features 8  
  Linux requirements 17  
  network and firewall requirements 42  
  system requirements 16  
  web browser requirements 40  
Veritas Access application  
  upgrade 135  
Veritas Access cluster name and network. *See*  
  reconfigure  
Veritas Access installation  
  overview 49  
VLAN device  
  adding on a particular NIC 104  
VLAN Tagging  
  about 104  
  limitations of 105