

Veritas Access 7.2 Release Notes

Linux

Veritas Access Release Notes

Last updated: 2016-11-02

Document version: 7.2 Rev 0

Legal Notice

Copyright © 2016 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/about/legal/license-agreements>

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

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

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

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The 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	Overview of Veritas Access	6
	About this release	6
	Important release information	6
	Changes in this release	7
	Changes to the cloud as a tier feature	7
	Installer performance improvements	8
	NetBackup integration	8
	New web-based Graphical User Interface	9
	Removal of the Storage> fss commands and the ability to include DAS disks in addition to SAN disks	9
	Configuring a CIFS share as secondary storage for an Enterprise Vault store	9
Chapter 2	Fixed issues	10
	Fixed issues since the last release	10
Chapter 3	Software limitations	12
	Flexible Storage Sharing limitations	12
	If your cluster has DAS disks, you must limit the cluster name to ten characters at installation time	12
	Limitations related to installation and upgrade	13
	Licensing messages displayed for Beta version of Veritas Access	13
	Limitations in the Backup mode	13
	Limitations in the Veritas InfoScale Operations Management Server console	13
	Veritas Access IPv6 limitations	14
	FTP create_homedirs limitation	14
	Samba ACL performance-related issues	14
	Veritas Access language support	14
	Veritas Access does not support non-English characters when using the CLISH (3595280)	15
	Limitations on using InfiniBand NICs in the Veritas Access cluster	15

	Limitation on using Veritas Access in a virtual machine environment	15
	File system limitation	16
	Any direct NLM operations from CLISH can lead to system instability (IA-1640)	16
Chapter 4	Known issues	17
	Veritas Access known issues	17
	Backup issues	17
	CIFS issues	17
	Deduplication issues	18
	FTP issues	19
	InfoScale Access Insight Add-on issues	19
	Installation and configuration issues	20
	Networking issues	24
	NFS issues	26
	ObjectAccess issues	28
	OpenStack issues	29
	Replication issues	29
	SmartIO issues	32
	Storage issues	32
Chapter 5	Getting help	41
	Displaying the online Help	41
	Displaying the man pages	41
	Using the Veritas Access product documentation	41

Overview of Veritas Access

This chapter includes the following topics:

- [About this release](#)
- [Important release information](#)
- [Changes in this release](#)

About this release

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 cloud based on policies.

This document provides release information about the Veritas Access product, including changes in this release.

Important release information

Review these Release Notes (this document) for the latest information before you install the product.

The hardware compatibility list contains information about supported hardware and is updated regularly. You can use any commodity hardware that is certified and mentioned in the hardware compatibility list. For the latest information on supported hardware, visit the following URL:

<http://www.veritas.com/docs/000019707>

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

<http://www.veritas.com/docs/000116052>

Changes in this release

This section shows the major new features and enhancements added in the 7.2 version of Veritas Access.

Changes to the cloud as a tier feature

The cloud as a tier feature is supported in this release. This feature was introduced as a Technical Preview feature in Veritas Access 7.1.

The following sections describe the changes to this feature:

Changes to the `Storage> tier move` commands

This release introduces new syntax for the `Storage> tier move` commands.

The `Storage> tier move` commands are used to move files to or from the cloud tier. The syntax has changed for the `Storage> tier move` commands to support the following operations:

- Start the move operation.
- Perform a dry run of the move operation.
- Specify the files to move according to the following criteria:
 - file or directory name pattern matching
 - last accessed time (`atime`)
 - last modified time (`mtime`)
- List the move operations in progress.
- Abort a move operation.

Controlling data movement to and from the cloud with policies

This release adds the ability to control data movement to and from the cloud with policies.

The `Storage> fs policy` commands are used to add a set of rules to control data movement between the primary tier and the cloud tier for a scale-out file system.

The following operations are supported:

- Add a policy to a file system to specify what files to delete or move between tiers.
- Add a schedule to the policy to automatically run the policy at specified times.

Obtaining statistics on data usage in the cloud tier in scale-out file systems

This release includes the following new commands for displaying and monitoring data usage in the cloud:

- `Storage> tier stats show` - displays the number of GET, PUT, and DELETE requests in the cloud tier
- `Storage> tier stats monitor` - monitors the usage of data in the cloud tier
- `Storage> tier stats usage` - displays the total data usage in the cloud tier

Installer performance improvements

This release includes the following installer improvements:

- Includes the `-precheck` option when you run the `installaccess -h` command.
- Provides a user interface to set up the Red Hat Subscription Manager (RHSM), and automatically enables the `rhel-6-server-debug-rpms` repository to install `kernel-debuginfo` and `kernel-debuginfo-common-x86_64` RPMs.
- Supports installing new required operating system RPMs, like `hal`, `sysstat`, `python-paramiko`, and new package updates for `glibc`, `OpenSSL`, `httpd`, which includes critical security fixes.
- Supports installing new third-party RPMs like `libuv`, `psutils`, and `python-crontab`.
- Supports installing the NetBackup client version 7.7.
- Supports installing the new Graphical User Interface (GUI).
- Performance improvements to reduce the Veritas Access installation and configuration time.
- Supports `ssh` multiplexing to improve `ssh` connection performance.

NetBackup integration

This release includes a built-in NetBackup client that is preinstalled on all the nodes in the Veritas Access cluster. Once the NetBackup domain information is configured

in Veritas Access, the NetBackup administrator can back up your Veritas Access file systems and retain the data as per your company policy. Once data is backed up, a storage administrator can delete unwanted data from Veritas Access to free up expensive primary storage for more data.

New web-based Graphical User Interface

The Veritas Access 7.2 release introduces a new Graphical User Interface (GUI). The GUI provides a dashboard view 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 Veritas Access cluster and monitor the results.

In this release, the GUI is part of Veritas Access.

Removal of the `Storage> fss` commands and the ability to include DAS disks in addition to SAN disks

The following `Storage> fss` commands have been removed from the CLISH:

- `Storage> fss disk format`
- `Storage> fss disk list`
- `Storage> fss pool adddisk`
- `Storage> fss pool create`
- `Storage> fss pool destroy`
- `Storage> fss pool free`
- `Storage> fss pool list`
- `Storage> fss pool rmdisk`
- `Storage> fss pool set`

The functionality for the `Storage> fss` commands has been merged with the `Storage> disk` and `Storage> pool` commands so that you can use DAS disks and SAN disks (LUNs) in any storage pool that you define. Multiple storage pools can have DAS disks, and any storage pool can have a mix of DAS and SAN disks.

The `fss layout` was removed from the `Storage> fs create` command.

Configuring a CIFS share as secondary storage for an Enterprise Vault store

This release includes functionality for configuring a CIFS share as secondary storage with Enterprise Vault 12.0 by exporting the file system over the CIFS protocol.

Fixed issues

This chapter includes the following topics:

- [Fixed issues since the last release](#)

Fixed issues since the last release

This section includes the issues fixed since the last release.

Table 2-1 Fixed issues since the last release

Fixed issues	Description
IA-905	Network DNS set nameserver accepts the same IP address multiple times.
IA-1191	RAID LEVEL displays in the Assert Summary.
IA-1225	Open request during the lock reclaim frequently fails with NFS4ERR_NO_GRACE in the node restart process.
IA-1235	The disk used for I/O fencing cannot be distinguished in the Management Server console.
IA-1335	Extra columns show in the exported sheet.
IA-1484	In the Storage Pools tab, there is no column to show whether the created pool is isolated or not.
IA-1497	The input description in the Create Pool wizard does not show in the Management Server console.
IA-1632	Host consuming shares values are missing for the file system.
IA-1633	Management Server console storage pool capacity.
IA-1635	Management Server console file system capacity.

Table 2-1 Fixed issues since the last release (*continued*)

Fixed issues	Description
IA-1788	Data delays to sync up in the Management Server console if you perform any operation in CLISH.
IA-1810	There is no validation when you type an unavailable user name in the CIFS: Add Permission window.
IA-1828	Enclosure does not recover to the Healthy status after it goes into the At Risk status.
IA-1840	InfoScale Access Insight Add-on should not be installed if the Storage Insight Add-on of the required version is not installed on the Veritas InfoScale Operations Manager server.
IA-1848	If you run the installer from one of the cluster nodes, the installer log is placed in <code>/var/tmp/</code> , instead of <code>/opt/VRTS/install/logs/</code> .
IA-1863	Response file generated by the installer has an incorrect variable.
IA-1893	You cannot create a snapshot with the same name of an existing snapshot.
IA-1894	You cannot grow or shrink the file system when it is offline.
IA-1897	Error message failed to update <code>/boot/initramfs-2.6.32-504.el6.x86_64.img</code> shows in the install log of package <code>VRTSvxvm</code> on RHEL 6.6.
IA-1905	You may not see all the disk names in the Summary page.
IA-1944	Object access service fails to assemble the uploaded chunks after failover.
IA-2810, IA-2928	Installer should install <code>sysstat</code> , <code>hal</code> , and <code>python-paramiko</code> as required operating system RPMs to enable CLISH commands to operate correctly.
IA-3226	<code>Cluster> add node</code> fails and the system gives an error.
3680071	Creating or deleting NIC bonds stops the replication services.
3790781	Creating network interface bond when other nodes of the cluster are down results in inconsistent network configuration state.
3811579	Storage fencing off does not work.
3845085	If the FC cables are pulled out from the node hosting the Management Server console, the CTDB server goes into the FAULTED state. CTDB service cannot be restored to ONLINE state when you use the "support-service-autofix" command.

Software limitations

This chapter includes the following topics:

- [Flexible Storage Sharing limitations](#)
- [Limitations related to installation and upgrade](#)
- [Limitations in the Backup mode](#)
- [Limitations in the Veritas InfoScale Operations Management Server console](#)
- [Veritas Access IPv6 limitations](#)
- [FTP create_homedirs limitation](#)
- [Samba ACL performance-related issues](#)
- [Veritas Access language support](#)
- [Limitations on using InfiniBand NICs in the Veritas Access cluster](#)
- [Limitation on using Veritas Access in a virtual machine environment](#)
- [File system limitation](#)

Flexible Storage Sharing limitations

The following issues relate to Veritas Access Flexible Storage Sharing (FSS).

If your cluster has DAS disks, you must limit the cluster name to ten characters at installation time

When formatting the DAS disks, the disks are given unique names. The names include the embedded cluster name. There is a limit of 25 characters for a DAS

disk name. When choosing the cluster name for a cluster that has DAS disks, you must limit the cluster name to ten characters.

Limitations related to installation and upgrade

The following limitations are related to installation and upgrade.

Licensing messages displayed for Beta version of Veritas Access

The Beta version of Veritas Access has a temporary license key. If you install a Beta version of the product, after 60 days you start seeing licensing messages such as the following:

```
WARNING V-365-1-1 This host is not entitled to run Veritas Storage
Foundation/Veritas Cluster Server.
As set forth in the End User License Agreement (EULA) you must complete one of the
two options set forth below. To comply with this condition of
the EULA and stop logging of this message, you have 0 days to either:
- make this host managed by a Management Server
(see http://go.veritas.com/sfhakeyless for details and free download), or
- add a valid license key matching the functionality in use on this host using the
command 'vxlicinst' and validate using the command 'vxkeyless set NONE'.
```

For Beta, these licensing messages cannot be avoided. The two options described do not apply for Veritas Access.

The GA version of Veritas Access has a permanent key.

Limitations in the Backup mode

If the backup group is online while performing a `cluster> del` operation, the `cluster> del` operation fails with the following error message:

```
CPI WARNING V-9-40-6450 Active backup jobs are running on access_01.
Deleting this node from the cluster may cause the backup to fail.
```

Limitations in the Veritas InfoScale Operations Management Server console

The InfoScale Access Insight Add-on is supported only on the Linux platform of Veritas InfoScale Operations Management Server. The InfoScale Access Insight Add-on cannot be deployed on a Windows server.

Veritas Access IPv6 limitations

The following Veritas Access modules are not supported for IPv6:

- NIS

The following IPv6 functionality is not supported for CIFS:

- CIFS does not support IPv4/IPv6 mixed mode for the domain controller. The IPv4 DNS entry needs to be removed from the DNS server.
- CIFS does not accept IPv6 addresses for the domain controller in the Veritas Access CLI. Only hostnames are allowed for the domain controller entry.

FTP create_homedirs limitation

Due to a limitation, you must manually create the user's logon directory even if the `create_homedirs` option is set to `yes`.

See the *Veritas Access Command-Line Administrator's Guide* for more information.

Samba ACL performance-related issues

For the ACL improvements to be effective (fewer number of attr nodes), the default mask for creating files and directories is set to 775. Previously, the create mask was set to 744.

If the mask is changed from 775, the ACL improvements may not be effective since the POSIX ACL's calculation changes significantly when the mask changes.

The performance improvements also depend on the file open mode. The current implementation considers normal file open using Windows Explorer or the command window. Samba may calculate a different open mode, depending on the permissions of the parent directory and the actual open request that is issued from the Windows client. These considerations impact the actual performance improvement.

Veritas Access language support

Veritas Access supports only English.

Veritas Access does not support non-English characters when using the CLISH (3595280)

The Veritas Access CLISH supports only English characters. File names such as CIFS shares must not include non-English characters. For example, the following command is not supported:

```
access> cifs share add sample "simfs01/サンプル"
```

Limitations on using InfiniBand NICs in the Veritas Access cluster

- InfiniBand NICs are preferred as private NICs, unless the NICs are connected to a public network or excluded.
- NIC bond function may not be supported on InfiniBand NICs when the PCI IDs are identical for the NICs on the same network card.

Note: The case is observed on Mellanox card.

- NIC exclusion function is supported on InfiniBand NICs, but all the NICs with the same PCI ID are excluded during the exclusion operation.

Note: The case is observed on Mellanox card.

- Newly added node should share the same configuration of InfiniBand NICs. For example, if the Veritas Access cluster uses LLT over RDMA, the newly added node should have RDMA NICs connected as a private NIC.
- Veritas Access does not support mixed LLT connections, which means all the nodes in the cluster nodes should have InfiniBand NICs if you plan to use LLT over RDMA. Otherwise, use NIC exclusion to exclude InfiniBand NICs during the Veritas Access installation.

Limitation on using Veritas Access in a virtual machine environment

Veritas Access must be installed on physical machines. Veritas Access is not supported in a virtual machine environment.

File system limitation

The following issue relates to the Veritas Access file system.

Any direct NLM operations from CLISH can lead to system instability (IA-1640)

Do not perform any file-system related operations by CLISH on the Network Lock Manager (NLM), as it is used for internal purposes. If NLM is used, then Veritas Access cannot guarantee the stability of the cluster.

Known issues

This chapter includes the following topics:

- [Veritas Access known issues](#)

Veritas Access known issues

The following known issues relate to the Veritas Access commands.

Backup issues

This section describes known issues related to backup.

Backup or restore status may show invalid status after the BackupGrp is switched or failed over to the other node when the SAN client is enabled (3606322)

When a backup job or a restore job is in progress over the SAN, and the BackupGrp is switched or failed over to the other node, the status option of the backup job in the CLISH may show the wrong status.

Workaround:

There is no workaround.

CIFS issues

This section describes known issues related to CIFS.

Cannot enable the quota on a file system that is appended or added to the list of homedir (3853674)

After enabling the `Storage> quota cifshomedir` command, if you set the additional file system as `cifshomedir`, the quota is not enabled on it by default. To enable the quota, if you use the `Storage> quota cifshomedir enable` command, it may or may not succeed, depending on the order in which you have specified the file systems as `cifshomedir`.

The `Storage> quota cifshomedir enable` command checks only for the first file system in the `cifshomedir` list. If the quota is already enabled on that file system, a quota on the rest of the file system in the list is not enabled.

Workaround:

To solve this issue, follow these steps:

- 1 Run the `Storage> quota cifshomedir disable` command. This disables the quota on all the homedir file systems.
- 2 Run the `Storage> quota cifshomedir enable` command. This enables the quota on all the homedir file systems.

Deleting a CIFS share resets the default owner and group permissions for other CIFS shares on the same file system (3824576, 3836861)

When you delete a CIFS share, the owner and the group on the file system revert to the default permissions. The default values for both the owner and the group are set to root. This behavior may be an issue if you have more than one CIFS share on the same file system. Deleting any of the shares also resets the owner and the group for the other shares on the file system.

If you previously set owner permissions or group permissions for the CIFS shares that remain, you must set the permissions again.

Workaround:

If you previously set owner permissions or group permissions for the CIFS shares that remain, you must set the owner or group permissions for the CIFS shares on the file system again, using the following command:

```
CIFS> share modify
```

Deduplication issues

This section describes known issues related to deduplication.

Removing lost+found files for a mount point that has deduplication enabled may cause issues with deduplication (3472414)

For a mount point that has deduplication enabled, the `lost+found` directory includes some files that are related to deduplication. If you remove the `lost+found` files, deduplication jobs may not work properly.

Workaround:

If you accidentally delete the deduplication files in the `lost+found` directory, perform the following steps to enable deduplication.

To enable the deduplication job:

- 1 Disable the deduplication job.
- 2 Enable the deduplication job.

FTP issues

The following issues relate to the Veritas Access FTP commands.

If a file system is used as `homedir` or `anonymous_login_dir` for FTP, this file system cannot be destroyed (IA-1876)

There is no `unset` command in FTP to change `homedir` or `anonymous_login_dir` to empty its value. You can use the FTP `set` commands to empty the values of the above two fields. Once all or any of the above fields are updated, either to point to some other file system or to be made empty, the original file system can be destroyed.

Workaround:

Use the `FTP> set` command to unset the values for `homedir` and/or `anonymous_login_dir`.

```
# isa> ftp set homedir_path
```

InfoScale Access Insight Add-on issues

The following known issues relate to the InfoScale Access Insight Add-on.

CLISH and the Management Server console do not allow decimal values for the file system size

When you create a file system, you are not allowed to type decimal values both on the CLISH and the Management Server console.

Workaround:

It is by design.

Ascending or descending does not work correctly as per disk name in the Physical Disks table (IA-1340)

If the disks are named as *words_integer*, and the *integer* is larger than 10, ascending or descending does not work correctly in the **Physical Disks** table.

Workaround:

There is no workaround.

Some special operations may make the enclosure-related tabs appear in the Properties windows (IA-1338)

Some operations may make the enclosure-related tabs appear in the **Properties** window when you leave the **Properties** window open.

Workaround:

There is no workaround.

When you create CIFS shares, some options in Export Options may conflict among themselves (IA-1899)

When you create CIFS shares, some options in **Export Options** may conflict among themselves.

Workaround:

There is no workaround.

You are not allowed to perform multiple operations in parallel

When you start multiple operations in parallel, only the first operation may succeed, and others fail.

Workaround:

There is no workaround.

Installation and configuration issues

The following issues relate to Veritas Access installation and configuration.

After you restart a node that uses RDMA LLT, LLT does not work, or the `gabconifg -a` command shows the jeopardy state (IA-1796)

The iptables are enabled by default on the Veritas Access cluster nodes. The iptables can affect the LLT function for the RDMA network.

Because LLT uses UDP to communicate in an RDMA network, you should add rules into the iptables to allow the LLT connection.

The iptable rules take effect before the LLT module is loaded. The iptables rules are managed by the Veritas Access script, which is executed after VCS comes up (it is started when the VCS Service Group comes online). When LLT is loaded, the iptables are in the default state, and the LLT connection through UDP is blocked.

Workaround:

For a fresh configuration of Veritas Access in an RDMA LLT environment:

- 1 After all the configurations are finished, log on to each node and disable the iptables by entering:

```
# chkconfig --level 123456 iptables off
```

- 2 Restart all the nodes. If the restart process cannot unload the OPENIB module, reset the node from the power management.

For adding a Veritas Access node in an RDMA LLT environment:

- 1 After completing the adding node, log on to each node (including the newly added one) and disable the iptables by entering:

```
# chkconfig --level 123456 iptables off
```

- 2 Restart all the nodes. If the restart process cannot unload the OPENIB module, reset the node from the power management.

Running individual Veritas Access scripts may return inconsistent return codes (3796864)

Individual scripts in Veritas Access are not intended to be run independently. The CLISH is the only supported interface for any operations in Veritas Access. If you run the Veritas Access scripts independently, then the return codes may not be consistent with the results in some cases.

Configuring Veritas Access with the installer fails when the SSH connection is lost (3794964)

When you install and configure Veritas Access with the installer, you may see the following error message:

```
CPI ERROR V-9-20-1073 Failed to copy /opt/SYMCsnas/conf/conf.tar
```

This message occurs in the rare case when the installer cannot copy the configuration file to the nodes in the cluster because the SSH connection is lost.

Workaround:

To work around this issue:

- 1 Recover the SSH connection manually.
- 2 Uninstall Veritas Access.
- 3 Reinstall Veritas Access.

Excluding PCI's from the configuration fails when you configure Veritas Access using a response file (3686704)

If you configure Veritas Access using a response file, Veritas Access does not exclude the PCI's that are marked for exclusion. During the configuration, the installer skips the NICs that need to be excluded.

Workaround:

Use the standard configuration method, or configure the NIC bonding and exclusion at the same time in the response file.

Installer does not list the initialized disks immediately after initializing the disks during I/O fencing configuration (3659716)

When you choose to configure I/O fencing after the installer starts the processes, you should have at least three initialized shared disks. If you do not have three shared disks, the installer can initialize the shared disks. After the installer initializes the disks, the installer does not list the initialized disks immediately.

Workaround:

After you initialize the disks, if you do not see the new disks in the installer list, wait for several seconds. Then select **y** to continue to configure I/O fencing. The installer lists the initialized disks.

Installer does not reboot the Veritas Access cluster nodes automatically (IA-3330)

After installation, the Veritas Access cluster nodes do not reboot automatically. After Veritas Access is installed, the cluster nodes should automatically reboot for dump to take effect. The installer does not reboot the system.

Workaround:

Reboot all the nodes of the cluster manually after the installation is complete.

Enter the following command as the root user to reboot each node:

```
# /sbin/shutdown -r now
```

If the same driver node is used for two installations at the same time, then the second installation shows the status of progress of the first installation (IA-3446)

The Veritas Access installer does not support multiple installations from the same driver node at the same time. This is by design. If you start two installations from the same driver node, then the second installation also shows the progress status of the first installation.

Workaround:

Do not perform multiple installations at the same time on the same driver node.

If the same driver node is used for two or more installations at the same time, then the first installation session is terminated (IA-3436)

The Veritas Access installer does not support multiple installations from the same driver node at the same time. This is by design. If you start two installations from the same driver node, then the first installation is terminated.

Workaround:

Do not perform multiple installations at the same time on the same driver node.

If you run the Cluster> show command when a slave node is in the restart, shutdown, or crash state, the slave node throws an exception (IA-900)

In a particular flow, if the node that is in the restart, shutdown, or crash state is running, the system calculates the running node list. It turns unreachable on SSH when the command starts to calculate the CPU or network statistics. The internal library throws an exception.

Once the state of the node is in shutdown, restart, or crash state, the slave node changes from RUNNING to FAULTED in Veritas Cluster Server (VCS). The `Cluster> show` command resumes its normal behavior. That is, it does not show any exception and gives an expected output.

Workaround:

There is no workaround for this issue. The system recovers itself. You need to wait for some time and run the `Cluster> show` command once again.

If duplicate PCI IDs are added for the PCI exclusion, the Cluster> add node name command fails (IA-1850)

To add a new node that has unique PCI IDs to be excluded, you need to add these unique PCI IDs through CLISH by using the `Network> pciexclusion add` command. If these unique PCI IDs already exist in the PCI exclusion configuration of Veritas Access, the resulting configuration has duplicate entries. After the resulting configuration for the PCI exclusion, if you proceed with the added node, the operation fails. The `Cluster> add node` operation cannot handle the duplicate entries in the PCI exclusion configuration.

Workaround:

Contact Technical Support to remove the duplicated PCI IDs from the Veritas Access PCI exclusion configuration files. Then you can run the `Cluster> add node` command.

If installing using a response file is started from the cluster node, then the installation session gets terminated after the configuring NICs section (IA-3570)

If you install Veritas Access using a response file from the cluster node, the installer does not provide a warning message to connect back to the installation after configuring the NICs.

Workaround:

- 1 Log on to Veritas Access with a new public IP address.
- 2 Execute the following command to proceed with the installation:

```
# /opt/VRTS/install/bin/tmux attach-session -t VA_INSTALL
```

Networking issues

This section describes known issues related to networking.

CVM service group goes into faulted state unexpectedly (3793413)

This issue occurs when the connectivity of storage is interrupted and brought back to a normal state. Veritas Volume Manager (VxVM) cannot join the cluster on that node if it hits the "minor number mismatch" issue.

Workaround:

Reboot the node on which this issue occurs.

In a mixed IPv4 and IPv6 VIP network setup, the IP balancing does not consider IP type (3616561)

In a mixed IPv4 and IPv6 setup, the IP balancing does not consider IP type. This behavior means that a node in the cluster might end up with no IPv6 VIP on it. IP balancing should consider the type of IP.

Workaround:

If required, manually bring online a VIP of the appropriate IP type on the node.

The netgroup search does not continue to search in NIS if the entry is not found in LDAP (3559219)

If the netgroups lookup order in the nsswitch settings is LDAP followed by NIS, a netgroup search does not continue to search in NIS if the netgroup entry is not found in LDAP. In this case, if the share is exported using netgroup, the NFS mount on the NFS client fails.

Workaround:

Change the netgroups lookup order so that NIS is before LDAP:

```
Network> nsswitch conf netgroups nis ldap
```

VIP and PIP hosted on an interface that is not the current IPv6 default gateway interface are not reachable outside the current IPv6 subnet (3596284)

IPv6 addresses configured on a non-default gateway interface are not reachable from outside the current subnet. That is, it is unable to use the current default gateway. Only IPv6 addresses that are hosted on the current default IPv6 gateway interface are reachable using the gateway.

Workaround:

Do not use VIPs that are currently not online on the default gateway interface for cluster communication outside the current subnet.

NFS issues

This section describes NFS issues.

Slow performance with Solaris 10 clients with NFS-Ganesha version 4 (IA-1302)

For the NFS-Ganesha server directory operations `mkdir`, `rmdir`, and `open`, the operations are slow when performed from the Solaris clients.

Workaround:

For performance-critical workloads using the Solaris platform, use the kernel-based NFS version 3 server.

Random-write performance drop of NFS-Ganesha with Linux clients (IA-1304)

There is a drop in the random-write performance for NFS-Ganesha with Linux clients. There is no drop in performance with Solaris clients.

Workaround:

For high-performance random-write workloads, use the kernel-based NFS server.

Latest directory content of server is not visible to the client if time is not synchronized across the nodes (IA-1002)

If the share is updated from multiple nodes, the actual server directory content may not be immediately visible on the client and will take some time. The cache invalidation of directory content is based on the modification time of the directory. Since the time is not in synchronization on the nodes of the cluster, this cache invalidation displays.

Workaround:

Configure NTP on the server to synchronize the time of all the nodes.

NFS> share show may list the shares as faulted for some time if you restart the cluster node (IA-1838)

This may occur when the NFS-Ganesha server is restarted across the cluster. It does not affect any ongoing NFS loads.

Workaround:

Wait for some time for the NFS-Ganesha shares to display as online.

NFS-Ganesha shares faults after the NFS configuration is imported(IA-849)

If you use the `System> config import` command to import any NFS configuration, then all the existing NFS shares go into the faulted state.

Workaround:

Restart the NFS service.

NFS-Ganesha shares may not come online when the number of shares are more than 500 (IA-1844)

The NFS-Ganesha shares may not come online, or take more time to come online, during the restart process if the number of NFS-Ganesha shares are about 500 or more.

Workaround:

Use netgroups or Kerberos instead of creating a large number of individual shares.

Exporting a single path to multiple clients through multiple exports does not work with NFS-Ganesha (3816074, 3819836)

Due to certain limitations of NFS-Ganesha, exporting a path to multiple clients (with the same or different permissions) through multiple exports does not work in Veritas Access.

Workaround:

Use netgroups to export the same path to multiple clients with the same permissions. Exporting the same path to multiple clients with different permissions is not supported.

For the NFS-Ganesha server, bringing a large number of shares online or offline takes a long time (3847271)

The NFS-Ganesha server has reduced performance when a large number of resources (that is, exported file system paths) are present. This behavior may result in slow recovery after a server failure. Starting or stopping the NFS server may also take a long time.

Workaround:

Use netgroups with the NFS-Ganesha server. If you encounter this issue, reduce the number of shares. This issue is only observed with a large number of shares.

NFS client application may fail with the stale file handle error on node reboot (3828442)

When a node restarts, all of the virtual IPs of the node are switched back to the restarted node. To preserve the lock information, the NFS-Ganesha server is restarted on this node. The VIP may be available for a short time before the shares are added back to the NFS-Ganesha server. This behavior causes applications to fail with a stale file handle error.

Workaround:

If this error is encountered, the client should retry the operation.

NFS> share show command does not distinguish offline versus online shares (IA-2758)

The `NFS> share show` command does not distinguish between offline and online shares. Shares that are faulted are listed correctly. You cannot determine the status of the share, Online or Offline, using only the CLISH commands.

Workaround

You can use the output of the Linux `showmount -e` command to get the list of exported shares from that specific cluster node.

Difference in output between NFS> share show and Linux showmount commands (IA-1938)

When using the `NFS> share show` command, you see the host name of the exported NFS client. When using the Linux `showmount` command, you see the IP address of the exported NFS client.

The NFS-Ganesha server always resolves the given host name to an IP address and exports the NFS share to that IP address. Unlike the kernel-based NFS server, the Linux `showmount` command returns IP addresses instead of host names provided in the export command. This does not affect any functionality, but the output is different between the two commands.

Workaround:

You can verify the given IP addresses by using DNS.

ObjectAccess issues

This section describes ObjectAccess issues.

Object access server goes in to faulted state while doing multi-part upload of a 10-GB file with a chunk size of 5 MB (IA-1943)

For large files, if the chunk size is small (5 MB), then while doing a multi-part upload, the object access server crashes while joining the large number of parts.

Workaround:

Veritas Access supports chunk sizes from 5 MB to 100 MB, so while uploading large files, it is recommended to use large chunk sizes up to 100 MB.

OpenStack issues

The following issues are related to OpenStack.

Cinder and Manila shares cannot be distinguished from the CLISH (3763836)

Any file system exported through NFS using the `OPENSTACK> cinder share` command, and any file system that is exported through NFS from OpenStack Manila cannot be distinguished through CLISH.

Workaround:

Use the `OPENSTACK> manila resource list` command to see only the shares that have been exported through Manila. There is no way to see Cinder shares exclusively.

Replication issues

This section describes known issues related to replication.

Running replication and dedup over the same source, the replication file system fails in certain scenarios (3804751)

The replication job may fail when the following situations occur on the same source replication file system:

1. NFS has a heavy I/O workload.
2. Deduplication that is running in parallel creates several shared extents.

Workaround:

There is no workaround.

The System> config import command does not import replication keys and jobs (3822515)

The `System> config import` command imports the configuration that is exported by the `System> config export` command. In the importing process, the replication repunits and schedules are imported correctly. The command fails to import the keys and jobs.

Workaround:

First run the `Replication> config import` command, and then perform the following steps.

- 1 Make sure the new target binds the replication IP, because the replication IP is not changed on the new source.
- 2 Run the `Replication> config import_keys` command on the source and the target.
- 3 Run the `Replication> config auth` command on the source and the target.
- 4 Delete the job directory from the new source `/shared/replication/jobs # rm -rf jobname/.`
- 5 Create the job from the new source.

Replication job with encryption fails after job remove and add link with SSL certificate error (3839319)

When you remove the link from an already configured job with encryption and again add the new link to the same job, the next replication cycle fails with the error:

```
SSL certificate error.
```

Workaround:

Follow these steps to solve this issue:

- 1 Execute the `Replication> job remove_link` command and exit the CLISH prompt on the source and the target.
- 2 Create a link `ln -s /shared/replication/SSL/cluster_cert /opt/VRTSfsadv/cert` on both cluster nodes of the source and the target.
- 3 Execute the `Replication> job add_link` command to add the link back to the job, and enable or sync the replication job.

Replication job status shows the entry for a link that was removed (3797560)

If a replication target in a multi-target job is removed, and you use the `Replication> job remove_link` command, then it is simply marked for removal. The actual removal of the link occurs during the next replication iteration.

Until the link is completely removed, the `Replication> job show` command displays the previous status of the removed link.

Workaround:

Use the `Replication> job show` command to verify when the link is completely removed.

The job uses the schedule on the target after replication failover (3668957)

This issue occurs if the schedules on the source cluster and the target cluster have the same name but different intervals. After replication fails over to a target, the job uses the schedule on the target.

Workaround:

Do not use the same schedule name on the source cluster and the target cluster.

Replication fails with error “connection reset by peer” if the target node fails over (IA-3290)

Replication creates a connection between the source and the target to replicate data. Replication uses one of the nodes from the target to access the file system to replicate data. In case the connection to this node breaks due to some error like a reboot, replication fails with an error message. If there is a scheduled replication job, the next iteration continues this failed replication session, possibly with a new node from the target.

Workaround:

If there is no scheduled replication job, you need to issue the `Replication> job sync` command to start the replication job once the target node is up.

Replication job modification fails (IA-3356)

Replication has a facility to have a multiple recovery point objective (RPO) report on the target side. The `Replication> job modify rep_dest_ckpt_cnt` command controls RPO. The default value is 10. Having RPO on the target side consumes some space on the target side, and hence replication can fail with an ENOSPC error. In this case, any replication job modification command fails.

Workaround:

Grow the target file system to make some more space. Modify the replication job to set the appropriate `rep_dest_ckpt_cnt` value. This modified value is not effective until the current replication session completes successfully. Once the modified value is applied, the existing RPO is adjusted as per the new value.

SmartIO issues

The following issue relates to the Veritas Access SmartIO commands.

SmartIO writeback cachemode for a file system changes to read mode after taking the file system offline and then online (IA-3423)

The SmartIO features lets you set writeback or read cache modes on a file system. Once the cachemode is set on a file system, it persists while the file system remains online. If the file system goes offline and is brought online again, the earlier cachemode does not persist and is reset to read cache mode.

Workaround:

Manually set the cachemode again once the file system comes online.

Storage issues

The following issues relate to the Veritas Access Storage commands.

Snapshot mount can fail if the snapshot quota is set (IA-1542)

If the snapshot quota is set, and the snapshot disk usage hits the quota hard limit, the checkpoint mount might fail, even when the removable snapshots exist. The snapshot operations can trigger snapshot removal to free some disk space if the file system runs out of space or the snapshot quota is exceeded. However, the snapshot mount cannot trigger this space-cleaning operation, so in some rare cases, the snapshot mount can fail.

Workaround:

Remove the oldest checkpoint and retry.

Sometimes the `Storage> pool rmdisk` command does not print a message (IA-1733)

A rare condition exists where the `Storage> pool rmdisk` command does not print either an error message or a success message due to a problem with output redirection.

Workaround:

Use the `history` command to check the status of the command. You can also use the `Storage> pool list` command to verify whether the disk was removed from the pool.

The `Storage> Pool rmdisk` command sometimes can give an error where the file system name is not printed (IA-1639)

If the disk being removed has NLM on it, the `Storage> pool rmdisk` command handles it differently, and no file system name is printed. Whether this error occurs depends on multiple factors, such as the pool size, how NLM uses disks, and the spread across disks.

Workaround:

There is no workaround.

Not able to enable quota for file system that is newly added in the list of CIFS home directories (IA-1851)

If you add a new file system as the CIFS home directory, then the quota is not enabled by default.

Workaround:

Run the following commands from CLISH:

```
Storage> quota cifshomedir disable
```

```
Storage> quota cifshomedir enable
```

Destroying the file system may not remove the `/etc/mtab` entry for the mount point (3801216)

When you destroy a file system, the `/etc/mtab` entry should be removed. If the file system `umount` command hangs during the destroy operation, the `/etc/mtab` entry might not be removed. The file system is destroyed but you cannot create a new file system with the same name.

Workaround:

Reboot the cluster nodes.

The Storage> fs online command returns an error, but the file system is online after several minutes (3650635)

The Storage> fs online command returns the following error:

```
access.Storage> fs online fs1  
  
ACCESS fs ERROR V-288-1873 filesystem fs1 not mounted on nodes  
access_01 access_02.
```

When you mount a file system with many checkpoints, the Veritas Cluster Server (VCS) resource might not respond for more than 100 seconds. . This causes the CFS command to timeout.

Workaround:

Even though the online failure is reported, the file system will be online.

Removing disks from the pool fails if a DCO exists (3452098)

If you specify disks on the command line when you create a file system, Veritas Access might create a data change object (DCO) on disks other than those specified. If free disks are available in the pool, Veritas Access prefers those for the DCO. The DCO is required to handle synchronization between the mirror and the original volume. The DCO is used when a disk that contains the data volume fails.

If you try to remove the disk from the pool, the following error displays because the disk is in use by the DCO.

```
SFS pool ERROR V-288-2891 Disk(s) sde are used by the following:  
DCO of primary tier of fs_mirror, Primary tier of filesystem fs_mirror
```

Workaround:

There is no workaround.

Scale-out file system returns an ENOSPC error even if the df command shows there is space available in the file system (IA-3545)

A scale-out file system returns an ENOSPC error even if the Linux `df` command shows there is space available in the file system.

This situation can happen in one of the following cases:

- A scale-out file system uses a hashing algorithm to distribute data between the storage containers. The algorithm makes sure that data is evenly distributed between all the containers, and depending on the type of the data, one of the storage containers is used more often than the other containers. A scale-out file system can reach 100% usage early. In this scenario, any allocation going to the 100% full container returns an ENOSPC error.
- A scale-out file system constitutes a metadata container and multiple data containers. Space for the metadata container is allocated at the time of creation of the file system. If the data containers are all full and the metadata container has available space, then the file system does not use the space in the metadata container. Because of this, the Linux `df` command can show there is still available space, but applications see an ENOSPC when writing to the file system.

Workaround:

Grow the file system.

Rollback refresh fails when running it after running Storage> fs growby or growto commands (3588248)

A rollback refresh fails if you run the rollback after running the `Storage> fs growby` or `Storage> fs growto` commands.

You create a rollback of a file system. After creating a rollback of a file system, you use the `Storage> fs growby` or `Storage> fs growto` commands to increase the size of the file system. If you perform a `Storage> rollback refresh` on the previously created rollback, the operation fails.

Currently the `Storage> rollback` command is designed to allow only using the same size in the `Storage> rollback refresh` command as that of the source file system. Automatically resizing snapshots before performing a rollback refresh is complicated, especially when a storage pool does not have enough space. The ability to automatically resize a snapshot is not implemented yet.

Workaround:

There is no workaround.

If an exported DAS disk is in error state, it shows ERR on the local node and NOT_CONN on the remote nodes in Storage> list (IA-3269)

If an exported DAS disk goes to an error state, its properties are not available on the remote nodes. The `Storage> disk list` command shows `NOT_CONN` on the remote nodes.

Workaround:

No workaround is necessary. If the disk goes online on the local node, it goes online on all the nodes.

Storage> disk remove for fusion-io disks is unable to remove fusion-io disks (IA-3217)

The `Storage> disk remove` command is not able to remove fusion-IO disks because some properties of the fusion-io disks are not available.

Workaround:

There is no workaround. Make sure that the fusion-io disks are not present in any storage pools before removing them from the array.

Inconsistent cluster state with management service down when disabling I/O fencing (IA-3427)

Disabling I/O fencing when one of the nodes is down results in the Veritas Access cluster being in an inconsistent state.

Workaround:

There is no workaround. Ensure that all the nodes in the cluster are up when disabling I/O fencing.

Storage> scanbus force command results in error (IA-3293)

If the management console node has any remote DAS disks, then the `Storage> scanbus force` command fails with an error.

Workaround:

There is no workaround. Ensure that no remote DAS disks are exported on the management node.

Storage> tier move command failover of node is not working (IA-3091)

The `Storage> tier move` command does not failover to another node if the node where it is running goes down.

Workaround:

Run the `Storage> tier move` command again from the CLISH.

File system creation fails for a pool that contains all SSDs (IA-3358)

This issue occurs with any DAS disks. The issue is when creating a common file system on DAS disks, Veritas Access tries to create the file system with mirrored disks across the hosts. If the pool, in which the file system is created, contains only disks from a single node, the common shared file system cannot be mirrored. For high availability, Veritas Access mandates that the common shared file system be mirrored on DAS disks across the hosts.

Workaround:

Add disks from more than one node in the pool in which the first file system was created.

Storage> scanbus operation hangs at the time of I/O fencing operation (IA-3257)

`Storage> scanbus` operation hangs during I/O fencing operation.

Workaround:

There is no workaround. Contact Veritas Technical Support.

Rollback service group goes in faulted state when respective cache object is full and there is no way to clear the state (IA-3251)

This issue relates to I/O errors after cache objects get full. In cases of cache-backed rollbacks, having cache full due to heavy I/O creates I/O errors in snapshots, and snapshots are automatically detached from the main file system. Snapshots go in to a faulted state. The fix for this requires clearing the faulty rollback state and doing rollback refreshes. There is no CLISH command to handle these cases. Manual intervention by Veritas Technical Support is required to preserve the rollback.

Workaround:

There is no workaround.

Rollback cache grow option missing in CLISH (IA-3240)

This issue relates to cases where the cache gets full and a cache grow operation can avoid the rollback going in to a faulted state. There is an enhancement request for adding a command in CLISH for rollback cache grow.

Workaround:

There is no workaround.

Event messages are not generated when cache objects get full (IA-3239)

This issue is related to customer visible events for rollback cache full scenarios.

Workaround:

There is no workaround.

Veritas Access CLISH interface should not allow uncompress and compress operations to run on the same file at the same time (IA-2995)

The Veritas Access CLISH interface does not block compress or uncompress operations while one of the other operations is running. This is a legacy behavior and should be fixed in a future release.

Workaround:

Do not initiate compress or uncompress operations on the same file at the same time while there are other compress or uncompress operations running on the same file.

Storage device fails with SIGBUS signal causing the abnormal termination of the scale-out file system daemon (IA-2915)

When a storage device fails and sends out a SIGBUS signal (bus error), it causes the abnormal termination of the scale-out file system daemon. The recovery process does not migrate the scale-out file system and the associated virtual IP of the file system's NFS share to the same claimed node. The output of the Linux `df` command on the NFS client shows incorrect sizes and usages (`Size Used, Avail, and Use%`) of the mounted scale-out file system's NFS share.

When this situation occurs, applications should stop using the NFS share of the scale-out file system before the issue resolves.

Workaround:

Re-export the scale-out file system's NFS share by logging on to the Veritas Access management console, and run the CLISH commands to delete and then add the NFS share again. If necessary, re-mount the NFS share on the NFS client for the applications as well.

Storage> tier move list command fails if one of the cluster nodes is rebooted (IA-3241)

The `Storage> tier move list` command fails until the cluster node is back up and running.

Workaround:

There is no workaround.

Pattern given as filter criteria to Storage> fs policy add sometimes erroneously transfers files that do not fit the criteria (IA-3432)

This issue was observed when the `**/*.txt` pattern was given as filter criteria when using the `Storage> fs policy add` command. When the policy was run, some of the files inside a `txt` directory, which did not have the file extension `.txt`, were selected for transfer or deletion. The expectation is that none of the files that do not have `.txt` as their extension should be selected for transfer or deletion.

Workaround:

There is no workaround.

Storage> pool destroy displays a false result if one of the nodes containing DAS disks is down (IA-3365)

This issue appears when you try to destroy a pool that contains DAS disks coming from a node that is down. The `Storage> pool destroy` operation erroneously reports success. The pool-related configuration is only partially removed. This is because Veritas Access cannot remove the pool configuration for the disks that contain DAS disks for a node that is down. This creates a further problem if the node comes up. The `Storage> pool list` command shows the pool information for the pool that was destroyed. If you try to create a pool with the same name while the node is down, after the node is back up, the `Storage> pool list` command shows even older disks as part of the new pool, even though the pool was not created using those disks.

Workaround:

There is no workaround. It requires clearing the configuration using the `support` account. If you encounter this issue, avoid creating pools by the same name as the destroyed pools. To bring the cluster to a clean state, you should contact Veritas Technical Support.

When a policy run completes after issuing `Storage> fs policy resume`, the total data and total files count might not match the moved data and files count as shown in `Storage> fs policy status (IA-3398)`

The `Storage> fs policy pause` command immediately stops the policy execution. If any files are transferred when this command is executed, the command does not stop for the transfer to be completed. While reporting the status of the `Storage> policy run` command, Veritas Access does not account for the data size and file count of the files that were in transit when the `Storage> fs policy pause` command executed.

Workaround:

You should perform a `Storage> fs policy dryrun` of the same policy again to check if there are any files that were missed in the transfer. You can also use the `Storage> tier mapfiles` and `Storage> tier listfile` commands to verify the location of the files.

Getting help

This chapter includes the following topics:

- [Displaying the online Help](#)
- [Displaying the man pages](#)
- [Using the Veritas Access product documentation](#)

Displaying the online Help

You can access the online Help through the Management Server console of Veritas InfoScale Operations Manager by clicking **Help**.

Displaying the man pages

You can enter Veritas Access commands on the system console or from any host that can access Veritas Access through a session using Secure Socket Shell (SSH).

Veritas Access provides the following features to help you when you enter commands on the command line:

- Command-line help by typing a command and then a question mark (?)
- Command-line man pages by typing `man` and the name of the command
- To exit a man page, type `q` (for quit).

Using the Veritas Access product documentation

The latest version of the Veritas Access product documentation is available on the Veritas Services and Operations Readiness Tools (SORT) website.

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

You need to specify the product and the platform and apply other filters for finding the appropriate document.

Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The documents are updated periodically for errors or corrections.

The following documents are available on the SORT site:

- *Veritas Access Command-Line Administrator's Guide*
- *Veritas Access Installation Guide*
- *Veritas Access Quick Start Guide*
- *Veritas Access Release Notes*
- *Veritas Access Third-Party License Agreements*
- *Veritas Access Troubleshooting Guide*