

Veritas Access 7.3 Troubleshooting Guide

Linux

7.3

Veritas Access Troubleshooting Guide

Last updated: 2019-04-04

Document version: 7.3 Rev 1

Legal Notice

Copyright © 2019 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	Introduction	6
	About troubleshooting	6
	General tips for the troubleshooting process	7
	General techniques for the troubleshooting process	7
	About the support user account	8
	Configuring the support user account	8
	Using the support login	9
Chapter 2	General troubleshooting procedures	11
	About general troubleshooting procedures	11
	Viewing the Veritas Access log files	11
	About event logs	12
	About shell-activity logs	12
	Setting the CIFS log level	12
	Setting the NetBackup client log levels and debugging options	13
	Retrieving and sending debugging information	14
	Insufficient delay between two successive OpenStack commands may result in failure	16
Chapter 3	Monitoring Veritas Access	17
	About monitoring Veritas Access operations	17
	Monitoring processor activity	17
	Generating CPU and device utilization reports	19
	Monitoring network traffic	20
	Exporting and displaying the network traffic details	21
Chapter 4	Common recovery procedures	23
	About common recovery procedures	23
	Restarting servers	24
	Bringing services online	25
	Using the services command	25
	Recovering from a non-graceful shutdown	27
	Testing the network connectivity	27
	Troubleshooting with traceroute	28

	Using the traceroute command	29
	Collecting the metasave image of a file system	29
	Replacing an Ethernet interface card	30
	Speeding up replication	31
	About synchronizing a replication job	31
	Synchronizing a replication job	32
	Uninstalling a patch release or software upgrade	33
Chapter 5	Troubleshooting the Veritas Access cloud as a tier feature	34
	Troubleshooting tips for cloud tiering	34
	Issues when reading or writing data from the cloud tier	34
	Log locations for checking for cloud tiering errors	35
Chapter 6	Troubleshooting Veritas Access installation and configuration issues	36
	How to find the management console IP	36
	Viewing the installation logs	37
	Installation fails and does not complete	37
	Excluding PCI IDs from the cluster	38
	Cannot recover from root file system corruption	39
	The storage disk list command returns nothing	39
Chapter 7	Troubleshooting Veritas Access CIFS issues	40
	User access is denied on a CTDB directory share	40
Chapter 8	Troubleshooting Veritas Access GUI startup issues	41
	Resolving GUI startup issues	41
Index		44

Introduction

This chapter includes the following topics:

- [About troubleshooting](#)
- [General tips for the troubleshooting process](#)
- [General techniques for the troubleshooting process](#)
- [About the support user account](#)
- [Configuring the support user account](#)
- [Using the support login](#)

About troubleshooting

Troubleshooting procedures for Veritas Access include the following types of procedures:

- Alert and log message review
- Routine maintenance tasks
- Commonly used recovery procedures
- Feature-specific problems and resolutions

Each of these procedures are described in the remaining chapters of this guide.

Some of the troubleshooting procedures in this guide require that you log in as the `support user`.

See [“About the support user account”](#) on page 8.

General tips for the troubleshooting process

To troubleshoot a problem, it helps to consider the following:

- Check for previous occurrence.
Check existing troubleshooting information to see if the problem has occurred before. For this type of information, a good source is the *Veritas Access Release Notes*. The release notes contain a list of known issues for Veritas Access and possible workarounds.
- Consider recent alterations.
If a system has problems immediately after some kind of maintenance, software upgrade, or other change, the problems might be linked to those changes.
- Determine what works.
If a system does not produce the desired end result, look for what operates properly. Identify where the problem is not and focus your efforts in other areas. Whatever components or subsystems necessary for the properly working parts to function are probably okay.
- Use your experience.
Based on your knowledge of how a system works, think of various failures that might cause this problem to occur. Check for those failures. Start with the most likely failures based on circumstances, history, or knowledge of existing feature weaknesses.

General techniques for the troubleshooting process

After applying some general troubleshooting tips to narrow the scope of a problem, here are some techniques to further isolate the problem:

- Swap identical parts.
In a system with identical or parallel parts and subsystems, it is a good idea to swap components between those subsystems and see whether or not the problem moves with the swapped component. For example, if you experience Veritas Access network connection problems on one node in a cluster, you could swap Ethernet Interface cards to determine if the problem moves to the new node.
- Remove parallel components.
If a system is composed of several parallel or redundant components that can be removed without crippling the whole system, start removing these components (one at a time) and see if things start to work. For example, in a cluster, shutdown the nodes one-by-one to see if the problem still persists.

- Divide the system into sections.
In a system with multiple sections or stages, carefully measure the variables going in and out of each stage until you find a stage where things do not look right. For example, if you run across a problem with a replication job, check to see if the job has run successfully before and try to determine the time frame when the job started to fail.
- Monitor system behavior over time (or location).
Display a list of services and their current status using the `Support> services show all` command.
Set up a process (such as the `Support> traceroute` command or a series of `Support> iostat` commands) to monitor system activity over a period of time or to monitor system activity across the network. This monitoring is especially helpful to track down intermittent problems, processor activity problems, node connection problems, and so on.

About the support user account

Generally, to access Veritas Access, you log into the management console with a Veritas Access user account. When you log in, you enter the command-line interface shell (CLISH). The command-line options depend on the role that the user account is assigned.

In some cases, the troubleshooting techniques in this guide require access to the underlying Linux command line and additional support utilities. The support user account provides access to these utilities. The support user account must be enabled (the default).

When you log in as support, you can access logs and other files that reside outside the CLISH.

Warning: Use caution when executing commands as the support user. The support commands are intended for advanced users who are familiar with Veritas Access features and functions. If you have any questions about using these commands, contact your Veritas Technical Support Representative for further guidance.

Configuring the support user account

A Veritas Access user with the `Master` role can enable, disable, change the password, or check the status of the support user.

The support user account is enabled by default.

To configure the support user account

- 1 To enable the support user, enter the following:

```
Admin> supportuser enable
Enabling support user.
support user enabled.
```

- 2 To verify that the support user is enabled:

```
Admin> supportuser status
support user status : Enabled
```

- 3 To change the support user password, enter the following:

```
Admin> supportuser password
Changing password for support.
Old password:
New password:
Re-enter new password:
Password changed
```

To disable the support user account

- 1 To disable the support user, enter the following:

```
Admin> supportuser disable
Disabling support user.
support user disabled.
```

- 2 To verify that the support user is disabled:

```
Admin> supportuser status
support user status : Disabled
```

Using the support login

When you log in as support, you can access logs and other files that reside outside the CLISH. Some of the troubleshooting techniques in this guide require that you log in as the support user.

The support user account must be enabled by an administrator with the `master` role.

See [“Configuring the support user account”](#) on page 8.

Note: The `support` account is intended for Technical Support and advanced users only.

To use the support login

- 1 Log in to the physical IP address of the cluster using the support account by entering:

```
support
```

Then enter the password. The default password is:

```
veritas
```

For example,

```
login as: support
support@<ip_address>'s password:
Last login: Tue Apr 26 14:53:32 2016 from ip_address
*****
*                               Veritas Access                               *
*                                                                           *
*                               Enterprise Edition                           *
*      Warning: Only Veritas Access distributed                            *
*      patches & RPMs can be installed on this system!                    *
*      Do not delete contents of lost+found directory of                  *
*      filesystems as it may contain critical temporary                   *
*      Veritas Access configuration data!                                  *
*****
```

```
WARNING: System configured with default password. It's recommended to
change password now. Please proceed with changing the password :
```

```
Changing password for support.
```

```
New password:
```

```
Re-enter new password:
```

```
Password changed
```

```
Default password is changed successfully on all the nodes.
```

```
ACCESSRC2_01:~ #
```

- 2 If you need to access the CLISH, you can use the following command:

```
su - master
```

General troubleshooting procedures

This chapter includes the following topics:

- [About general troubleshooting procedures](#)
- [Viewing the Veritas Access log files](#)
- [About event logs](#)
- [About shell-activity logs](#)
- [Setting the CIFS log level](#)
- [Setting the NetBackup client log levels and debugging options](#)
- [Retrieving and sending debugging information](#)
- [Insufficient delay between two successive OpenStack commands may result in failure](#)

About general troubleshooting procedures

This chapter provides an overview of general troubleshooting procedures you can use to help discover and fix problems.

Viewing the Veritas Access log files

In addition to the Alerts panel on the Veritas Access Operations Manager console dashboard, the Veritas Access `/opt/VRTSnas/log` directory is a good place to find out more about problems that may occur.

To view the Veritas Access log files

- 1 Use the support account to login.
- 2 Navigate to the `/opt/VRTSnas/log` directory.

About event logs

In addition to the system log, each Veritas Access feature has an associated event log. When a problem occurs, one of the quickest ways to learn more about what occurred is to examine these log files. Event logs for Veritas Access features are stored in the `/opt/VRTSnas/log` directory.

Note: You should not delete or alter log files while troubleshooting, as it may hamper further investigation by Veritas Technical Support.

To view the event logs:

- 1 Use the support account to login.
- 2 Navigate to the `/opt/VRTSnas/log` directory.

Event logs for Veritas Access features are stored in this directory.

For example, the `cifs.log` contains CIFS event logs.

About shell-activity logs

You can use the shell-activity logs to capture any command-line operations performed by the end user or the customer. The shell-activity logs help you to understand any unwanted operations done by the end user either intentionally or unintentionally.

You can find the shell-activity logs for the following at:

- Support account - `/var/log/shell_activity_log`
- CLI commands - `/opt/VRTSnas/log/command.log`

Setting the CIFS log level

You can set the CIFS log level for the Veritas Access cluster, and then upload the debugging information to an external server for troubleshooting.

See the `support_debug.1` man page.

See [“Retrieving and sending debugging information”](#) on page 14.

To set the CIFS log level

- ◆ Set the CIFS-related log level for the Veritas Access cluster.

```
Support> debuginfo setlog cifs loglevel
```

A valid `loglevel` ranges from 0 to 10, 10 being the most detailed log level. It is recommended to increase the CIFS log level, reproduce the CIFS issue, and then upload debugging information for the CIFS issue.

The default log level is 2.

For example, to set the CIFS log level to 10 for the Veritas Access cluster:

```
Support> debuginfo setlog cifs 10
```

Setting the NetBackup client log levels and debugging options

You can set NetBackup client log levels as well as different debugging options, and then upload the information to an external FTP or SCP server. You can use this debugging information to send to Veritas Technical Support.

See [“Retrieving and sending debugging information”](#) on page 14.

You can find NetBackup log information by using the `Backup> show` command.

See the `backup_show(1)` man page.

You can see what NetBackup log levels and debugging options have been enabled by executing the `Backup> show` command.

See the *Veritas NetBackup Administrator's Guide, Volume 1* for more information on NetBackup logging.

Valid log level values range from 1 to 5, 5 being the most detailed. See the `support_debuginfo(1)` man page.

To set the NetBackup client log levels

- 1 Set the NetBackup database log level:

```
Support> debuginfo setlog nbu database loglevel
```

- 2 Set the NetBackup global debugging log level:

```
Support> debuginfo setlog nbu global loglevel
```

Global logging controls the logging level for the processes that are set in the **Logging** dialog box in the NetBackup Administration Console.

To set the NetBackup debugging options

- 1 Enable the NetBackup client to perform robust logging in the cluster.

```
Support> debuginfo setlog nbu enable robust
```

Robust logging limits the amount of disk space that a log directory consumes.

- 2 Enable the NetBackup client to perform critical process logging in the cluster.

```
Support> debuginfo setlog nbu enable critical
```

The enable critical process option lets you automatically log critical NetBackup processes. Log directories for the critical processes are created and logging begins when this option is enabled in the **Logging** host properties in the NetBackup Administration Console.

Retrieving and sending debugging information

You can retrieve Veritas Access debugging information from a Veritas Access node and send the information to a server using an external FTP or SCP server.

See the following article for more information on how to provide data for Veritas Technical Support:

<http://www.veritas.com/docs/000097935>

To upload debugging information from a specified node to an external server

- ◆ Upload debugging information from a specified node to an external server.

```
Support> debuginfo upload nodename debug-URL module
```

For example, to upload all debugging information to an FTP server:

```
Support> debuginfo upload node1_1  
ftp://admin@ftp.docserver.company.com/patches/ all
```

For example, to upload CIFS-related debugging information to an SCP server:

```
Support> debuginfo upload node1_1
scp://root@server.company.com:/tmp/node1_1-cifs-debuginfo.tar.gz
```

nodename Specifies the *nodename* from which to collect the debugging information.

debug-URL Specifies the remote file or directory for uploading debugging information.

Depending on the type of server from which you upload debugging information, use one of the following example URL formats:

```
ftp://admin@ftp.docserver.company.com/
patches/
```

```
scp://root@server.company.com:/tmp/
```

If *debug-URL* specifies a remote file, the debuginfo file is saved by that name. If *debug-URL* specifies a remote directory, the debuginfo file name displays as the following:

```
nas_debuginfo_nodename_modulename_timestamp.tar.gz
```

module Specifies the values for *module*.

Supported module values are the following:

- all - use to collect all information for debugging
- generic - use to collect all debugging information except for information related to Veritas products
- cifs - use to collect CIFS-related debugging information
- nas - use to collect Veritas Access related debugging information
- netbackup - use to collect NetBackup client-related debugging information

The `Support> debuginfo` command also collects information for the `sosreport` command for Red Hat Enterprise Linux (RHEL). The `sosreport` is collected for all the loaded modules except for the `selinux` module.

Insufficient delay between two successive OpenStack commands may result in failure

If sufficient delay is not there between two successive OpenStack commands, it may lead to failure of the operation in some cases.

For example:

```
# for i in {01..32}; do cinder create --image zesty-server-cloudimg-  
amd64 --volume-type vrts_vol_type --name zesty-vol-$i 10; done
```

The command may fail to create the new volumes.

To avoid this kind of failure, introduce sufficient delay between the commands by adding some sleep time.

For example:

```
# for i in {01..32}; do cinder create --image zesty-server-cloudimg-amd64  
--volume-type vrts_vol_type --name zesty-vol-$i 10;sleep 20; done
```


Monitoring Veritas Access

This chapter includes the following topics:

- [About monitoring Veritas Access operations](#)
- [Monitoring processor activity](#)
- [Generating CPU and device utilization reports](#)
- [Monitoring network traffic](#)
- [Exporting and displaying the network traffic details](#)

About monitoring Veritas Access operations

This chapter describes several support tasks that are useful for monitoring Veritas Access operations. Perform these monitoring tasks periodically to ensure that Veritas Access is running smoothly.

As you work with Veritas Access, keep an ongoing record of the output created by monitoring commands. This process gives you a baseline for judging normal operations and helps you to flag potential problems before they become serious.

Monitoring processor activity

The `support> top` command displays the dynamic real-time view of currently running tasks. It shows the resources that users and processes consume at a given time for a specified node.

To use the top command

- ◆ To use the `Support> top` command, enter the following:

```
Support> top [nodename] [iterations] [delay]
```

nodename Displays the resources and processes at a given time for the specified node.

iterations Specifies the number of iterations you want to run. The default is three.

delay Specifies the delay between screen updates. The default is five seconds.

For example, to show the dynamic real-time view of tasks running on the node `access_01`, enter the following:

```
Support> top access_01 1 1
```

```
top - 16:28:27 up 1 day, 3:32, 4 users, load average: 1.00, 1.00, 1.00
Tasks: 336 total, 1 running, 335 sleeping, 0 stopped, 0 zombie
Cpu(s): 0.1% us, 0.1% sy, 0.0% ni, 99.7% id, 0.0% wa, 0.0% hi, 0.0% si
Mem: 16405964k total, 1110288k used, 15295676k free, 183908k buffers
Swap: 1052248k total, 0k used, 1052248k free, 344468k cached
```

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
6314	root	15	0	5340	1296	792	R	3.9	0.0	0:00.02	top
1	root	16	0	640	260	216	S	0.0	0.0	0:04.86	init

Generating CPU and device utilization reports

To use the iostat command

- ◆ To use the `Support> iostat cpu` command, enter the following:

```
Support> iostat cpu [nodename] [interval] [count]
```

nodename The name of the node from where the report is generated. The default is `console` for the Management Console.

interval The duration between each report in seconds. The default is 2 seconds.

count The number of reports generated at the `interval` entered in seconds. The default is one report.

where the *nodename* option asks for the name of the node from where the report is generated. The default is `console` for the Veritas Access Operations Manager console.

For example, to generate the CPU utilization report of the console node, enter the following:

```
Support> iostat cpu access_01
```

```
Linux 2.6.16.60-0.21-smp (access_01)            02/09/16
```

avg-cpu:	%user	%nice	%system	%iowait	%steal	%idle
	1.86	0.07	4.53	0.13	0.00	93.40

To use the iostat device command

- ◆ To use the `Support> iostat device` command, enter the following:

```
Support> iostat device [nodename] [dataunit]
[interval] [count]
```

<i>nodename</i>	The <i>nodename</i> option asks for the name of the node from where the report is generated. The default is <code>console</code> for the Management Console.
<i>dataunit</i>	The <i>dataunit</i> option lets you generate the report in block(s) or kilobytes(s). The default is block(s).
<i>interval</i>	The duration between each report in seconds. The default is two seconds.
<i>count</i>	The number of reports generated at the <i>interval</i> entered in seconds. The default is one report.

For example, to generate a device utilization report of a node, enter the following:

```
Support> iostat device access_01 Blk
Linux 2.6.16.60-0.21-smp (access_01)          02/09/16

Device:      tps    Blk_read/s    Blk_wrtn/s    Blk_read    Blk_wrtn
hda          4.82        97.81         86.11    1410626    1241992
sda          1.95        16.83          4.05     242712      58342
hdc          0.00         0.01          0.00        136         0
```

Monitoring network traffic

Tethereal is a command-line version of Ethereal, a network protocol analyzer supported by the Linux operating system. It lets you capture packet data from a live network or read packets from a previously-saved capture file.

To help you monitor network traffic, Veritas Access provides a `Support> tethereal` command that lets you display and export network traffic data.

- The `Support> tethereal show` command displays packed data captured from a live network.
- The `Support> tethereal export` command lets you export network traffic details for further analysis.

Exporting and displaying the network traffic details

To use the tethereal command

- ◆ To use the `Support> tethereal export` command, enter the following:

```
Support> tethereal export url [nodename] [interface] [count]
[source]
```

<i>url</i>	Provides the location to export the network traffic details. The default file name <code>tethereal.log</code> is used if a file name is not specified in the url.
<i>nodename</i>	The name of the node from where the traffic details are generated.
<i>interface</i>	Specifies the network interface for the packet capture.
<i>count</i>	Specifies the maximum number of packets to read. The maximum allowed file size to capture the network traffic details is 128 MB. For a very large "count" value, if the file size exceeds 128 MB, then the command stops capturing the network traffic details.
<i>source</i>	Specifies the node to filter the packets.

For example, to export the network traffic details, enter the following:

```
Support> tethereal export scp://user1@172.31.168.140:~/
Password: *****
Capturing on pubeth0 ...
Uploading network traffic details to scp://user1@172.31.168.140:~/
is completed.
```

When you export network traffic details, press the **Ctrl + C** keys to stop the capture process and upload traffic details to the URL site.

To use the `tethereal show` command

- ◆ To use the `Support> tethereal show` command, enter the following:

```
Support> tethereal show [nodename] [interface] [count]  
[source]
```

nodename The name of the node from where the traffic details are displayed.

interface Specifies the network interface for the packet capture.

count Specifies the maximum number of packets to read.

If you do not specify a count value, the network traffic continues to be displayed until you interrupt it.

source Specifies the node to filter the packets.

For example, the traffic details for five packets are:

```
Support> tethereal show access_01 pubeth0 5 172.31.168.140  
0.000000 172.31.168.140 -> 10.209.105.147 ICMP Echo (ping) request  
0.000276 10.209.105.147 -> 172.31.168.140 ICMP Echo (ping) reply  
0.000473 10.209.105.147 -> 172.31.168.140 SSH Encrypted response  
packet len=112  
0.000492 10.209.105.147 -> 172.31.168.140 SSH Encrypted response  
packet len=112
```

Common recovery procedures

This chapter includes the following topics:

- [About common recovery procedures](#)
- [Restarting servers](#)
- [Bringing services online](#)
- [Recovering from a non-graceful shutdown](#)
- [Testing the network connectivity](#)
- [Troubleshooting with traceroute](#)
- [Using the traceroute command](#)
- [Collecting the metasave image of a file system](#)
- [Replacing an Ethernet interface card](#)
- [Speeding up replication](#)
- [Uninstalling a patch release or software upgrade](#)

About common recovery procedures

This chapter provides some of the most-common recovery procedures you can use to troubleshoot a problem with Veritas Access.

Restarting servers

Some configuration changes do not take effect until the associated server is restarted. Therefore, some configuration problems can be solved by stopping and restarting the associated server. For example, when you change AD Domain settings, you need to restart the CIFS server.

Table 4-1 shows commands you can use to start and stop Veritas Access servers.

Table 4-1 Commands to start and stop servers

Command	Definition
Backup> start	Starts all configured backup services.
Backup> stop	Stops all configured backup services.
CIFS> server start	Starts the CIFS server.
CIFS> server stop	Stops the CIFS server.
FTP> server start	Starts the FTP server.
FTP> server stop	Stops the FTP server.
NFS> server start	Starts the NFS server.
NFS> server stop	Stops the NFS server.
Storage> iscsi start	Starts the iSCSI initiator service.
Storage> iscsi stop	Stops the iSCSI initiator service.

Note: Some commands include the `server` argument and some do not. Also, some `Support>` commands use a `service` (instead of `server`) argument.

Bringing services online

The `Support> services` command lets you bring services that are OFFLINE or FAULTED back to the ONLINE state.

Note: After you use the `Support> services` command, if a service is still offline or faulted, you need to contact Technical Support.

These services include:

- Backup
- Console service
- CIFS server
- FTP
- FS manager
- GUI
- IP addresses
- NIC information
- NFS server

Using the services command

To display the state of the services

- ◆ To display the important services running on the nodes, enter the following:

```
Support> services show
```

Service	access	
	01	02
nfs	ONLINE	ONLINE
cifs	ONLINE	ONLINE
ftp	ONLINE	ONLINE
iSCSIInitiator	OFFLINE	OFFLINE
gui	ONLINE	ONLINE
console	ONLINE	ONLINE
nic_pubeth0	ONLINE	ONLINE
nic_pubeth1	ONLINE	ONLINE
fs_manager	ONLINE	ONLINE

To display the state of all of the services

- ◆ To display all of the services running on the nodes, enter the following:

```
Support> services showall
```

	access	
Service	01	02
-----	-----	-----
nfs	ONLINE	ONLINE
cifs	ONLINE	ONLINE
ftp	ONLINE	ONLINE
iSCSIInitiator	OFFLINE	OFFLINE
console	ONLINE	ONLINE
gui	ONLINE	ONLINE
nic_pubeth0	ONLINE	ONLINE
nic_pubeth1	ONLINE	ONLINE
fs_manager	ONLINE	ONLINE
10.182.107.201	ONLINE	ONLINE
10.182.107.202	ONLINE	ONLINE
10.182.107.203	ONLINE	ONLINE
10.182.107.204	ONLINE	ONLINE
/vx/fs1	ONLINE	ONLINE

To fix any service fault

- ◆ To fix any service fault, enter the following:

```
Support> services autofix
Attempting to fix service faults.....done
```

To bring a service online

- ◆ To bring a service online on the nodes, enter the following:

```
Support> services online servicename
```

where *servicename* is the name of the service you want to bring online.

For example:

```
Support> services online 10.182.107.203
```

This IP address is the virtual IP address that can be onlined.

Recovering from a non-graceful shutdown

In some cases, when a non-graceful shutdown of a node occurs (for example, during an unexpected system halt or power failure), you may receive an error message on the local node asking you to use the Linux `fsck` (file system check) command to repair files on the node.

Attempting to use the `fsck` command to repair the node is not recommended (and may not be possible). Instead, use a healthy node in the cluster to reinstall Veritas Access software on the damaged node.

To recover a node

- 1 Use the `master` account to log into Veritas Access.
- 2 Delete the failed node from the cluster. To delete the node, enter the following:

```
Cluster> del nodename
```

where *nodename* is the name of the failed node.

For example:

```
Cluster > del access_01
```

Note: The failed node information is deleted from the cluster. When the failed node is rebooting, it will detect that it has been deleted and clean itself up.

- 3 After the node is deleted from the cluster, reboot the deleted node and then it is reachable using its original physical IP address (before the node had been added to the cluster).
- 4 Add the node back by entering the following:

```
Cluster> add nodeip
```

where *nodeip* is the reachable IP address of the deleted node.

For example:

```
Cluster > add 172.16.113.118
```

Testing the network connectivity

You can test whether a particular host or gateway is reachable across an IP network.

To use the ping command

- ◆ To use the ping command, enter the following:

```
Network> ping destination [nodename]
[devicename] [packets]
```

For example, you can ping host1 using node1:

```
Network> ping host1 node1
```

<i>destination</i>	Specifies the host or gateway to send the information to. The destination field can contain either a DNS name or an IP address.
<i>nodename</i>	Specifies the <i>nodename</i> to ping from. To ping from any node, use <i>any</i> in the <i>nodename</i> field. The <i>nodename</i> field is an optional field. If <i>nodename</i> is omitted, any node is chosen to ping from.
<i>devicename</i>	Specifies the device through which you ping. To ping from any device in the cluster, use the <i>any</i> variable in the <i>devicename</i> field.
<i>packets</i>	Specifies the number of packets that should be sent to the destination. If the packets field is omitted, five packets are sent to the destination by default. The packets field must contain an unsigned integer.

Troubleshooting with traceroute

Traceroute is a widely-available utility supported by the Linux operating system. Much like ping, traceroute is a valuable tool to determine connectivity in a network. The Veritas Access `Support> ping` command enables you to discover connections between two systems. The `Support> traceroute` command checks system connections as well, but also lists the intermediate hosts between the two systems. Users can see the routes that packets can take from one system to another. Use the `Support > traceroute` command to find the route to a remote host.. For example, you might use the `Support> traceroute` command to verify the connectivity of each node in your cluster.

Using the traceroute command

The `Support> traceroute` command displays all of the intermediate nodes on a route between two nodes.

To use the traceroute command

- ◆ To use the `Support> traceroute` command, enter the following:

```
Support> traceroute destination [source]
[maxttl]
```

destination The target node. To display all of the intermediate nodes that are located between two nodes on a network, enter the *destination* node.

You need to specify either an IPv4 address for an IPv4 installation or an IPv6 address for an IPv6 installation.

The accepted range for an IPv6 prefix is 0-128 integers.

source Specifies the *source* node name from where you want to begin the trace.

maxttl Specifies the maximum number of hops. The default is seven hops.

For example, to trace the route to the network host, enter the following:

```
Support> traceroute www.veritas.com fssClus_01 10
traceroute to www.veritas.com (23.5.150.79), 10 hops max, 60 byte packets
 1 puna-sli-core-b01-vlan329.net.symantec.com (10.209.192.2) 0.356 ms 0.354 ms 0.376 ms
 2 punb-vfpi-eng-1-aggregate2-104.net.veritas.com (10.209.186.14) 0.298 ms 0.322 ms 0.379 ms
 3 puna-spi-core-b02-vlan105.net.symantec.com (143.127.185.130) 1.851 ms 1.964 ms 1.940 ms
 4 bnrcatcore01-teng6-2.net.symantec.com (143.127.185.205) 1.902 ms 1.903 ms 1.932 ms
 5 puna-vfpi-main-1-vip.net.veritas.com (10.212.252.50) 1.886 ms 1.945 ms 1.922 ms
```

Collecting the metasave image of a file system

You can collect a metasave image of a regular or a scale-out file system for troubleshooting file system issues. Metadata is a data structure that contains attributes about the data within a file system, but does not contain the actual data itself. You can use metadata images for tracking file system trends, such as the file size, age, and type of information in the file system.

Note: When using the `Support> metasave` command, the file system must be offline on all the cluster nodes to create a consistent metasave image. Bring the file system offline using the `Storage> fs offline` command before collecting the metasave image. Metasave image collection is a time consuming operation. The total time that is required depends on the amount of metadata information present in the file system. If you have a scale-out file system, it can take significantly longer to collect a metasave image. You can run other Veritas Access operations from a separate terminal while running the metasave operation.

To collect the metasave image of a file system

- ◆ To use the `Support> metasave` command, enter the following:

```
Support> metasave [fsname] [output_location]
```

`fsname` Specifies the name of the file system for which you want to collect a metasave image of the file system.

`output_location` Specifies the directory location of the metasave image.

For a regular file system, a single metasave image is stored at the directory location specified by `output_location`.

For a scale-out file system, multiple metasave images are produced depending on the number of container file systems inside the scale-out file system. For scale-out file systems, the namespace mapping is also included in the metasave image.

For example, to collect the metasave image of file system `testfs`, and store it under `/tmp/meta_out_dir`, enter the following:

```
Support> metasave testfs /tmp/meta_out_dir
Collecting metasave image of file system testfs. This may take some time...
SUCCESS: Metasave image of testfs collected successfully. Image is stored at /tmp/meta_out_dir.
```

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.

Speeding up replication

In some cases, a replication job may not proceed as fast as expected. In this situation, you may need to resynchronize the replication job.

About synchronizing a replication job

The first time a replication job is run, Veritas Access makes a full copy of the data from the source location to the destination. Subsequent jobs (triggered manually or through a schedule) only copy incremental changes.

In certain rare cases, data is already present at the destination, but the replication cannot make the incremental changes. Examples of this situation include:

- When replication has not been run for several days or weeks, and the changes that are tracked by the VxFS file change log have been overwritten (or possibly corrupted). This log is required for replication.

- When a replication job is temporarily disabled and started again, the next job run triggers a full copy of the data.
- When some changes have been made to the replication definition. For example, an earlier replication consisted of `fs1/folder1`, but you want to replicate data in `fs1/folder2` also. Because `fs1/folder2` requires a full copy, `fs1/folder1` is copied once again, even though only incremental changes are needed.
- When the direction of the replication has to be reversed from destination to source. Even though most data is present at both the destination and the source, anytime you create a new job at the destination, a full copy is triggered automatically for the first replication.
- If an administrator accidentally deletes the internal database for replications and no backup is available, creating a new job (even for an existing configuration) triggers a full copy.

In these cases, instead of waiting to initiate a full copy, you can use the `Replication> job sync` command to leverage the existing data at the destination and avoid requiring a full copy. The `Replication> job sync` command returns the replication job to a well-defined state and incremental replication can be used.

After you sync a job, the job is re-enabled, and you can use the standard job trigger or set the replication frequency to trigger incremental replication.

Note: Synchronization is only supported on enabled jobs. If you are not able to resume from a failed job, and you want to use the `Replication> job sync` command to recover from this state, follow these steps. First, disable the job, then enable the job again. Then, use the `Replication> job sync` command to synchronize the job.

Note: Synchronization can not be performed on a paused replication job. If synchronization is performed on a paused job that has been aborted or stopped, the last recovery point objective (RPO) for the paused job is not available.

Synchronizing a replication job

To synchronize an enabled replication job

- ◆ To synchronize an enabled replication job, enter the following:

```
Replication> job sync job_name
```

job_name

Specify the name of the replication job you want to synchronize.

Uninstalling a patch release or software upgrade

Often a problem occurs because of a known product defect. Once the defect is fixed, you can install a patch release or software upgrade to fix the issue.

When you plan to install a patch release or software upgrade:

- Before you start the installation, use the `System> config export` command to save a copy of your configuration. After the upgrade, you can use the `System> config import` command to restore your configuration.

Example:

```
System> config export local 2016_July_20
```

```
System> config import local 2016_July_20 network
```

- To upgrade with minimal downtime, you need to obtain a set of temporary VIP and IP addresses to use during the upgrade. Alternatively, you can upgrade without using temporary VIP and IP addresses, but the downtime increases.

For details on upgrading Veritas Access, refer to the *Veritas Access Installation Guide*.

Troubleshooting the Veritas Access cloud as a tier feature

This chapter includes the following topics:

- [Troubleshooting tips for cloud tiering](#)
- [Issues when reading or writing data from the cloud tier](#)
- [Log locations for checking for cloud tiering errors](#)

Troubleshooting tips for cloud tiering

To troubleshoot cloud tiering

- 1 Make sure that the cloud endpoint is pingable from the Veritas Access server.
- 2 Make sure that the system time is accurate on the Veritas Access server.
- 3 If you are using any Amazon S3-compatible storage service, make sure that the service supports AWS signature version 4.

Issues when reading or writing data from the cloud tier

If you were able to successfully add the cloud service and the cloud tier, and if you encounter issues when reading or writing data from the cloud tier, perform the following troubleshooting steps.

To troubleshoot issues when trying to read or write data from the cloud tier

Test for PUT:

- ◆ Upload an object testobj to cloud tier tier1 of file system fs1.

```
[root@clus2_01 /]# /opt/VRTSnas/bin/cldiotest -w tier1 fs1 testobj
"test_object_content"
Write: Length 19 return 19
```

Test for GET:

- ◆ Download an object testobj present in cloud tier tier1 of fs1 to /testfile.

```
[root@clus2_01 /]# /opt/VRTSnas/bin/cldiotest -d tier1 fs1 testobj
/testfile
[root@testclus2_01 /]# cat /testfile
test_object_content
```

Test for HEAD:

- ◆ Run this command.

```
[root@clus2_01 /]# /opt/VRTSnas/bin/cldiotest -x tier1 fs1 testobj
Size 19 return 0
```

Test for bucket listing:

- ◆ List all the objects present in tier tier1 of fs1.

```
[root@clus2_01 /]# /opt/VRTSnas/bin/cldiotest -l tier1 fs1
testobj          19
testobj2         20
testobj3         20
```

Log locations for checking for cloud tiering errors

Check the following log locations for finding more information on cloud tiering errors:

- /opt/VRTSnas/log/tfsd.log
- /opt/VRTSnas/log/tfslib.log
- /var/log/messages

Troubleshooting Veritas Access installation and configuration issues

This chapter includes the following topics:

- [How to find the management console IP](#)
- [Viewing the installation logs](#)
- [Installation fails and does not complete](#)
- [Excluding PCI IDs from the cluster](#)
- [Cannot recover from root file system corruption](#)
- [The storage disk list command returns nothing](#)

How to find the management console IP

To identify which node is the console IP (management console IP)

- 1 Identify which node is the management console IP.

```
# hares - state | grep -I console
```

- 2 Use a Secure Shell (ssh) to access the management console (only one node has the management console).
- 3 On the management console, log on to the CLISH using the following command:

```
su - master
```

Viewing the installation logs

If a problem occurs during installation, it can be helpful to view entries in the installation logs to help pinpoint problems.

To view the Veritas Access installation logs

- 1 During Veritas Access installation and configuration, you can access installer logs in a temporary folder under `/var/tmp`.
- 2 After Veritas Access installation and configuration, you can view a copy of the installation logs in the following locations:

Veritas Access post-installation logs `/opt/VRTS/install/logs/installaccess-timestamp`

This directory is located on the node from which the installer is triggered (the driver node). It contains the Veritas Access specific installation logs.

For example:

`/opt/VRTS/install/logs/installaccess-201602021544AsJ`

Veritas Access service group configuration logs `/opt/VRTSnas/log/Install.log`

This directory contains the Veritas Access specific configuration logs.

For example:

`/opt/VRTSnas/log/Install.log.201407030655`

Veritas Access network installation and configuration logs `/opt/VRTSnas/log/install_network.log`

This directory contains the Veritas Access network configuration logs.

For example:

`/opt/VRTSnas/log/install_network.log.201407030655`

Installation fails and does not complete

Some common reasons for installation failures include:

- Limited memory. You must have at least 32 GB of memory to install Veritas Access software on a node.
- Single core (single CPU)

You must have at least two nodes in a cluster (or a dual-CPU system) to install Veritas Access.

- **Missing required operating system packages**
 You can use yum to install the missing required operating system packages, or manually install the missing required packages.
 See the *Veritas Access Installation Guide* for more information.
- **Gateway access**
 The Veritas Access node must be able to reach the default gateway using the public network. Verify with your network administrator that the gateway is reachable.

Excluding PCI IDs from the cluster

During the initial Veritas Access software installation on the first node, you can exclude certain PCI IDs in your cluster to reserve them for future use. You may want to exclude additional PCI IDs when you add additional nodes to the cluster. You can add the PCI IDs to the exclusion list. The interface cards for which PCI ID's have been added in the PCI exclusion list are not used as private or public interfaces for the subsequent cluster node install. During a new node install, the remaining PCI bus interfaces are searched and added as public or private interfaces.

The `Network> pciexclusion` command can be used with different options:

- The `Network> pciexclusion show` command displays the PCI IDs that have been selected for exclusion. It also provides information about whether it has been excluded or not by displaying y(yes) or n(no) symbols corresponding to the node name. If the node is in the INSTALLED state, it displays the UUID of the node.
- The `Network> pciexclusion add pcilist` command allows an administrator to add specific PCI ID(s) for exclusion. These values must be provided before the installation. The command excludes the PCI from the second node installation.
pcilist is a comma-separated list of PCI IDs.
- The `Network> pciexclusion delete pci` command allows an administrator to delete a given PCI ID from exclusion. This command must be used before the installation for it to take effect. The command is effective for the next node install

The *PCI* ID bits format is hexadecimal (XXXX:XX:XX.X).

Cannot recover from root file system corruption

You cannot perform a `fsck` (file system check) if the Red Hat Enterprise Linux operating system partition is corrupted, and you need to restart the node. While trying to restart the node, the operating system prompts for the `root` user password to run the file system check.

Consult the Red Hat Enterprise Linux documentation for the solution to this issue.

https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/3/html/System_Administration_Guide/s1-rescuemode-boot.html

The storage disk list command returns nothing

The `storage disk list` command may display blank disk names due to the following reasons:

- In Virtual environments, Virtual Machine Disk (VMDK) disks are shared across the nodes of the cluster.

Resolution:

Ideally, VMDK disks should be attached as local disks.

- Local disks have a unique device identifier (UDID) that is not unique. For example, in a cluster having node1 and node2, the local disk that is attached to node1 may have the same UDID as the local disk attached to node2.

Resolution:

Run the following command:

```
# storage disk configure local <node_name> <vendor_id> <product_id>
[serial_num]
```

Where `<vendor_id>` is the ID of the vendor

`<product_id>` is the ID of the product

`<node_name>` is the name of the node on which the command is run.

`<serial_num>` should be given in the format `opcode/pagecode/offset/length` calculated carefully from the serial number format of the disk.

If you want to run the command on all the nodes, specify the `<node_name>` value as `all`.

For details on the `storage disk configure local` command, see the `storage_disk` manual page.

Troubleshooting Veritas Access CIFS issues

This chapter includes the following topics:

- [User access is denied on a CTDB directory share](#)

User access is denied on a CTDB directory share

In some cases, users or groups may be denied access to a CTDB directory share even though the correct ACL is set for the share. This issue can occur when the parent directory has an ACL that prevents access for these users or groups.

This behavior is expected. To enable access:

- Make sure the root-level directory (the parent directory) is added as a CIFS share.
- To allow access, apply the same ACL settings to the parent directory as you applied to the original CTDB directory share.

Troubleshooting Veritas Access GUI startup issues

This chapter includes the following topics:

- [Resolving GUI startup issues](#)

Resolving GUI startup issues

Veritas Access GUI accessibility issues occur if specific ports are inaccessible. Ports might be turned off on the node or on the network switch. Veritas selectively opens ports at the network switch.

To use the Veritas Access GUI after installing Veritas Access

- 1 Obtain the console virtual IP address by using the `Network> ip addr show` command.
- 2 Use the console IP with the port number 14161 to access the Veritas Access GUI.

Example:

```
https://console IP address:14161
```

- 3 Log on to the Veritas Access GUI using the `support` user name and password. If this does not work, verify the GUI set up.

To verify the GUI set up

- 1 Check the `/opt/SYMCsnas/log/ isagui_config.log` file to verify that the GUI is properly configured.

If there are any problems during the configuration, the problems are reported in this log file.

- 2 You need to allow ports 5634 and 14161 to be accessible remotely.
- 3 Open these ports by executing the following commands.

You must log on as the `root` user.

```
# /etc/init.d/iptables save
# /etc/init.d/iptables stop
```

- 4 Turn off the firewall on start up:

```
# chkconfig iptables off
```

The commands work if there is no network switch-based firewall in the environment. Otherwise you need to contact the network administrator to open these ports.

- 5 Ports must be opened before the GUI is configured. Otherwise you should rerun the GUI configuration. Before you rerun the GUI configuration, try connecting the browser to the management console.
- 6 You can verify if a port is accessible by running the following command:

```
telnet hostname/ipaddress 14161
```

If the port is not opened or not listened to, the connection waits forever. Try connecting with a random port that is not open. You see a difference in behavior.

- 7 Restart if the web server is not running.

```
service vamgmt forcestop
```

```
service vamgmt start
```

```
ps -ef | grep node
```

After running the `ps -ef | grep node` command, the results should show:

```
/opt/SYMCsnas/isagui/ext_modules/node /opt/SYMCsnas/isagui/application/server.js development
```

- 8 You should be able to connect to the GUI and be able to log on.

- 9** If data is not properly discovered or not seen in the GUI, run the following commands:

```
Export EXTRA_LOG=1
```

```
/opt/SYMCsnas/pysnas/bin/isagui_cluster_perf.py --full
```

- 10** If there are any errors, check the log file.

```
/opt/SYMCsnas/log/isagui_cluster_perf.log
```

Index

A

- about
 - common recovery procedures 23
 - event logs 12
 - job resynchronization 31
 - monitoring commands 17
 - services command 25
 - shell-activity logs 12

C

- changing
 - support user password 8
- checking
 - support user status 8
- CIFS
 - setting the log level 12
- cloud tiering
 - log locations 35
 - troubleshooting 34
- Cluster
 - Excluding PCI IDs 38
- common recovery procedures
 - about 23
- configuring
 - job resynchronization 32
- CPU utilization report
 - generating 19

D

- debugging information
 - retrieving and sending 14
- debugging options
 - setting for NetBackup 13
- device utilization report
 - generating 19
- disabling
 - support user account 8
- displaying
 - node-specific network traffic details 21

E

- enabling
 - support user account 8
- event logs
 - about 12
- Excluding PCI IDs
 - cluster 38
- exporting
 - network traffic details 21

F

- fsck
 - unable to run if operating system partition is corrupted 39

G

- general techniques
 - troubleshooting 7
- general tips
 - troubleshooting process 7
- generating
 - CPU utilization report 19
 - device utilization report 19

I

- installation
 - common failures 37
- installation logs
 - viewing 37

J

- job resynchronization
 - about 31
 - configuring 32

L

- log locations
 - cloud tiering 35

login

- support account 9
- Technical Support 9

M**monitoring**

- installation logs 37
- processor activity 17

monitoring commands

- about 17

N**NetBackup client log levels**

- setting 13

NetBackup debugging options

- setting 13

network

- testing connectivity 27

network traffic details

- exporting 21

node-specific network traffic details

- displaying 21

O**operating system**

- unable to repair using fsck 39

P**patch release**

- uninstalling 33

processor activity

- monitoring 17

R**reading or writing data from the cloud tier**

- troubleshooting 34

recovering

- from a non-graceful shutdown 27

replacing

- Ethernet interface card 30

replication

- speeding up 31

restarting

- servers 24

retrieving

- debugging information 14

S**sending**

- debugging information 14

servers

- restarting 24

services command

- about 25
- using 25

setting

- CIFS log level 12
- NetBackup client log levels 13

shell-activity logs

- about 12

shutdown

- recovering from a non-graceful 27

software upgrade

- uninstalling 33

support account

- login 9

support user account

- about 8
- disabling 8
- enabling 8

support user password

- changing 8

support user status

- checking 8

T**technical support**

- login 9

testing

- network connectivity 27

traceroute

- troubleshooting with 28

traceroute command

- using 29

troubleshooting

- about 6
- general procedures 11
- issues when reading or writing data from the cloud tier 34

troubleshooting process

- general techniques 7
- general tips 7

U

- uninstalling
 - patch release or software upgrade 33
- using
 - services command 25
 - traceroute command 29

V

- Veritas Access log files
 - viewing 11
- viewing
 - installation logs 37
 - Veritas Access log files 11