

Veritas Access Troubleshooting Guide

7.4.2 Linux

To be used with the Veritas Access
Appliance v7.4.2 (3340)
documentation

Veritas Access Troubleshooting Guide

Last updated: 2021-02-14

Document version: 7.4.2 Rev 0

Legal Notice

Copyright © 2018 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Veritas InfoScale, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-Party Programs"). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/licensing/process>

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

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

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

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

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

accessdocs@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	6
	General techniques for the troubleshooting process	7
Chapter 2	General troubleshooting procedures	9
	About general troubleshooting procedures	9
	Viewing the Veritas Access log files	9
	About event logs	10
	About shell-activity logs	10
	Setting the CIFS log level	10
	Setting the NetBackup client log levels and debugging options	11
	Retrieving and sending debugging information	12
	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	24
	Restarting servers	24
	Bringing services online	25
	Using the services command	26
	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 (online mode)	30

	Replacing an Ethernet interface card (offline mode)	33
	Replacing a Veritas Access node	34
	Replacing a disk	47
	Speeding up replication	55
	About synchronizing a replication job	55
	Synchronizing an episodic replication job	56
	Uninstalling a patch release or software upgrade	56
Chapter 5	Troubleshooting the Veritas Access cloud as a tier feature	58
	Troubleshooting tips for cloud tiering	58
	Issues when reading or writing data from the cloud tier	58
	Log locations for checking for cloud tiering errors	59
Chapter 6	Troubleshooting Veritas Access installation and configuration issues	60
	How to find the management console IP	60
	Viewing the installation logs	61
	Installation fails and does not complete	61
	Excluding PCI IDs from the cluster	62
	Cannot recover from root file system corruption	63
	The storage disk list command returns nothing	63
Chapter 7	Troubleshooting the LTR upgrade	64
	Locating the log files for troubleshooting the LTR upgrade	64
	Troubleshooting pre-upgrade issues for LTR	64
	Troubleshooting post-upgrade issues for LTR	65
Chapter 8	Troubleshooting Veritas Access CIFS issues	66
	User access is denied on a CTDB directory share	66
Chapter 9	Troubleshooting Veritas Access GUI startup issues	67
	Resolving GUI startup issues	67
Index		70

Introduction

This chapter includes the following topics:

- [About troubleshooting](#)
- [General tips for the troubleshooting process](#)
- [General techniques for the troubleshooting process](#)

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 on as the `support user`.

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.

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 12.

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 12.

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.

If a node is a part of the Access cluster, you can use the Veritas Access CLI to run the `debuginfo` command to collect and upload debug information to an SCP or an FTP server. If the node is not a part of the Access cluster, the Veritas Access CLI is not accessible, and you must instead use the `collect_debuginfo.sh` script to collect and upload the debug information.

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

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

Uploading debug information from a cluster node

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
```

<i>nodename</i>	Specifies the <i>nodename</i> from which to collect the debugging information.
-----------------	--

<i>debug-URL</i>	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*.

The following options are supported:

- **generic**: This module collects generic information such as details about the operating system and the CIFS module and SOS reports. This does not collect product information.
- **nas**: This module collects product information from the cluster nodes.
- **os**: This module collects SOS reports from all the cluster nodes.
- **explorer**: The VxExplorer utility collects logs and environment data from all the servers where the Veritas Access product is installed.
- **install**: This module collects the install logs from the CPI install log directory `/opt/VRTSnas/install/logs/`.
- **all**: This module collects information for all the above modules from the cluster nodes.

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.

Uploading debug information from a node that is not a part of the Access cluster

Use the `/opt/VRTSnas/scripts/support/collect_debuginfo.sh` script to collect and upload debug information when a node is not a part of the Access cluster. The script uses the following format:

```
/opt/VRTSnas/scripts/support/collect_debuginfo.sh -o upload -n
nodename -d debug-URL -m module
```

where:

upload

Specifies the operation to perform. Use the `upload` option to upload the debug information to the location specified by the *debug-URL* parameter.

nodename

Specifies the name of the node from where you want to collect the debug information.

debug-URL

Specifies the location for uploading the debug information. The location can be an SCP server, an FTP server, or a remote file or a directory. Use one of the following formats based on location where you want to upload the debug information:

```
ftp://url
```

For example:

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

```
scp://url
```

For example:

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

If you specify a remote file, the debug information is saved in the specified file. If you specify a remote directory, the debug information is saved in the directory as a `tar.gz` file in the following format:

```
nas_debuginfo_nodename_modulename_timestamp.tar.gz
```

For example, the `sh`

```
/opt/VRTSnas/scripts/support/collect_debuginfo.sh
```

```
-o upload -n node_01 -d
```

```
file:///root/log/ -m
```

`nas-procstacks` uploads the information to the

```
nas_debuginfo_node_01_nas-procstacks_201808270145.tar.gz
```

file in the `logs` directory.

modules

Specifies the modules. The following options are supported:

- `generic` This module collects generic information such as details about the operating system and the CIFS module, and the SOS reports. This does not collect product information.
- `nas` This module collects product information from the cluster nodes.
- `os` This module collects SOS reports from all the cluster nodes.
- `explorer` The VxExplorer utility collects logs and environment data from all the servers where the Veritas product is installed.
- `install` This module collects the install logs from the CPI install log directory `/opt/VRTS/install/logs/`.
- `all` This module collects information for all the above modules from the cluster nodes.

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]
```

<i>nodename</i>	Displays the resources and processes at a given time for the specified node.
<i>iterations</i>	Specifies the number of iterations you want to run. The default is three.
<i>delay</i>	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 \(online mode\)](#)
- [Replacing an Ethernet interface card \(offline mode\)](#)
- [Replacing a Veritas Access node](#)
- [Replacing a disk](#)
- [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
```

	access	
Service	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 online.

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.

To recover a node

- 1 Login to the Management console (Veritas command-line interface) prompt.
- 2 Run the following command:

```
storage>fs fsck <fsname>
```

where *fsname* is the name of the file system.

Once the file check is complete, you can continue to use the file system.

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 (online mode)

In some cases, you may need to replace an existing Ethernet interface card on a node. This section describes the steps to replace the NIC card. When you replace the NIC card, there should not be any mismatch with the number of NICs in the cluster. All the nodes in the cluster should have an equal number of disks after you replace the NIC card.

You need to provide an accurate and permanent MAC address (in case of bonded NICs) before you proceed with the NIC replacement. High availability services of service groups are temporarily disabled during the NIC replacement operation.

Note: This procedure does not work for adding an Ethernet interface card to the cluster and in VLAN environments. After the successful replacement operation, remove the faulty NIC card. Before you install the Ethernet interface card on the node, install the required device driver for the Ethernet interface card.

To replace an online Ethernet interface card (NIC)

- 1 Add a new NIC card on the server.

Note: The new NIC card should be online and searchable by the server.

- 2 Run the `# ip link show` command to get the MAC address of the new NIC card.
- 3 Run the following command on the Veritas Access node to replace the NIC card.

```
# /opt/VRTSnas/scripts/net/net_device_add.sh -r  
<old_mac_address> -w <new_mac_address>
```

- 4 To replace the NIC card in the bonded interface, you need to find a permanent hardware address by using one of the following commands.

```
# ethtool -P <interfacename>
```

or

```
# cat /proc/net/bonding/<bondname>
```

For details, see the following examples.

Example: To replace the "pubeth2" interface in the bond with the new NIC "eth0"

Bonding details:

```
# cat /proc/net/bonding/bond0  
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)  
  
Bonding Mode: load balancing (round-robin)  
MII Status: up  
MII Polling Interval (ms): 100
```

```
Up Delay (ms): 0
Down Delay (ms): 0
```

```
Slave Interface: pubeth1
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 00:50:56:05:0a:ea
Slave queue ID: 0
```

```
Slave Interface: pubeth2
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 00:50:56:05:e0:45
Slave queue ID: 0
```

```
# ip link show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN mode DEFAULT qlen 1
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: priveth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode
DEFAULT qlen 1000
link/ether 00:50:56:05:a3:1d brd ff:ff:ff:ff:ff:ff
3: pubeth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq
state UP mode DEFAULT qlen 1000
link/ether 00:50:56:05:e0:44 brd ff:ff:ff:ff:ff:ff
4: pubeth1: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500
qdisc mq master bond0 state UP mode DEFAULT qlen 1000
link/ether 00:50:56:05:0a:ea brd ff:ff:ff:ff:ff:ff
5: pubeth2: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500
qdisc mq master bond0 state UP mode DEFAULT qlen 1000
link/ether 00:50:56:05:0a:ea brd ff:ff:ff:ff:ff:ff
6: eth0: <BROADCAST,MULTICAST> mtu 1500 qdisc mq state DOWN mode DEFAULT qlen 1000
link/ether 00:50:56:05:41:53 brd ff:ff:ff:ff:ff:ff
7: priveth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode
DEFAULT qlen 1000
link/ether 00:50:56:05:e0:41 brd ff:ff:ff:ff:ff:ff
8: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc
noqueue state UP mode DEFAULT qlen 1000
link/ether 00:50:56:05:0a:ea brd ff:ff:ff:ff:ff:ff
```


NIC replacement operation

```
# /opt/VRTSnas/scripts/net/net_device_add.sh -r 00:50:56:05:e0:45  
-w 00:50:56:05:41:53  
100% [#] Success: Device replace successful.
```

After NIC replacement

```
# cat /proc/net/bonding/bond0  
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)  
  
Bonding Mode: load balancing (round-robin)  
MII Status: up  
MII Polling Interval (ms): 100  
Up Delay (ms): 0  
Down Delay (ms): 0  
  
Slave Interface: pubeth1  
MII Status: up  
Speed: 10000 Mbps  
Duplex: full  
Link Failure Count: 0  
Permanent HW addr: 00:50:56:05:0a:ea  
Slave queue ID: 0  
  
Slave Interface: pubeth2  
MII Status: up  
Speed: 10000 Mbps  
Duplex: full  
Link Failure Count: 0  
Permanent HW addr: 00:50:56:05:41:53  
Slave queue ID: 0
```

Replacing an Ethernet interface card (offline mode)

In some cases, you may need to replace an existing Ethernet interface card (NIC) on a node. This section describes the steps to replace the NIC card. When you replace the NIC, there should not be any mismatch with the number of NICs in the cluster. All the nodes in cluster should have an equal number of disks after you replace the NIC card.

You need to provide a correct and permanent MAC address (in case of bonded NICs) before you proceed with the NIC replacement. High availability services of service groups are temporarily disabled when you replace the NIC. The same PCI slot needs to be used when you replace an Ethernet card.

In the VLAN configured environment, when you add nodes, the IPs/netmasks that are assigned to the devices may not work correctly. To avoid these issues, modify the IPs by using the following command:

```
clish> network ip addr modify 192.168.10.21 192.168.10.27 255.255.240.0
```

After you add a node in the cluster, it triggers the recovery for the detached disks.

Note: These steps do not work for adding an Ethernet interface card to the cluster. After the successful replacement operation, remove the faulty NIC. Before you install the Ethernet interface card on the node, install the device driver that is required for the Ethernet interface card.

To replace an offline Ethernet interface card

- 1 Run the `Cluster> del` command to delete the node from the cluster.
- 2 Install the Ethernet interface card on the node that you want replace with the existing NIC card and restart the node.

Note: Make sure that the Ethernet interface card and node are online and searchable.

- 3 Run the `Cluster> add` command to add the node into the cluster.

For example:

```
Cluster> add 172.16.113.118
```

Replacing a Veritas Access node

In some cases, you may need to replace a Veritas Access node. This section describes the steps to replace a Veritas Access node.

To replace a Veritas Access node

- 1 Before you delete the node from the cluster, make sure that you do not remove the CVM master node. To remove the CVM master node, you need to switch the CVM master node by switching the Management Console to other node.
- 2 If you do not want to trigger the Hot Relocation, set the following tunables to -1 from the CVM master node.

```
#vxtune node_reloc_timeout -1
```

Note: After you set the `node_reloc_timeout`, the `storage_reloc_timeout` is automatically set to -1.

- 3 Run the `cluster del` command for the node that is to be replaced.

```
fss7310>cluster del fss7310_02
```

- 4 Verify that all the plexes are in the NODEVICE/DISABLED state. You can use the `#vxprint -p` command to check the plex states.
- 5 Run the following command to detach the plexes of the volumes:

```
# vxplex -f -g <dg-name> -o rm dis <plex-name>
```

- 6 Remove all the disks that are in `failed was:` state from the disk group by using the `vx dg rmdisk` command. This command needs to be run from the CVM master node.

```
#vx dg -g <dg-name> rmdisk <disk-name>
```

- 7 Run the `vx disk rm` command for the removed disks from all the nodes in the cluster.

```
#vx disk rm <disk-name>
```

Note: This command needs to be run for all the disks from all the nodes in the cluster.

- 8 After all the plexes are disabled, add the new node in the cluster by using the following command:

```
fss7310>cluster add <node-ip>
```

- 9** Run the `storage disk format` command from the Veritas Access management console node for all the disks from the newly added node.

```
fss7310>storage disk format <list-of-disks>
```

- 10** Add all the disks from the newly added node to the Veritas Access pool created by using the `storage pool adddisk` command.

```
fss7310> storage pool adddisk pool1 <list-of-devices>
```

- 11** Run the `storage fs addmirror` command to mirror the file system.

```
fss7310> storage fs addmirror <fs-name> <pool-name>
```

- 12** Run the `vxassist` command to mirror the `_nlm_` volume as well.

```
#vxassist -b -g <dg-name> mirror _nlm_
```

Example: Replacing a Veritas Access node

To replace a Veritas Access node

- 1** Change the value of the vxtune tunable to disable hot relocation:

```
# vxtune node_reloc_timeout -1
```

2 Run the following command to remove the node from the cluster.

```
fss7310> cluster del fss7310_02
```

Veritas Access 7.4.2 Delete Node Program

```
fss7310_02
```

Copyright (c) 2017 Veritas Technologies LLC. All rights reserved. Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The Licensed Software and Documentation are deemed to be "commercial computer software" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Logs are being written to /var/tmp/installaccess-201803130635kXW while installaccess is in progress.

Veritas Access 7.4.2 Delete Node Program

```
fss7310_02
```

```
Checking communication on fss7310_01 ..... Done
Checking communication on fss7310_02 ..... Done
Checking communication on fss7310_03 ..... Done
Checking communication on fss7310_04 ..... Done
Checking VCS running state on fss7310_01 ..... Done
Checking VCS running state on fss7310_02 ..... Done
Checking VCS running state on fss7310_03 ..... Done
Checking VCS running state on fss7310_04 ..... Done
```

The following changes will be made on the cluster:

Failover service group VIPgroup4 will be switched to fss7310_01

```
Switching failover service group(s) ..... Done
Waiting for service group(s) to come online on the other sub-cluster ..... Done
All the online failover service group(s) that can be switched have been switched to
the other sub-cluster.
```

The following parallel service group(s) in the sub-cluster will be offline:

```
fss7310_02: CanHostConsole CanHostNLM Phantomgroup_pubeth0 ReconfigGroup cvm iSCSI_INIT
vrts_vea_cfs_int_cfsmount1 vrts_vea_cfs_int_cfsmount2 vrts_vea_cfs_int_cfsmount3
vrts_vea_cfs_int_cfsmount4 vrts_vea_cfs_int_cfsmount5 vrts_vea_cfs_int_cfsmount6
Offline parallel service group(s) ..... Done
Waiting for service group(s) to be taken offline on the sub-cluster ..... Done
```

```
Stopping VCS on fss7310_02 ..... Done
Stopping Fencing on fss7310_02 ..... Done
Stopping gab on fss7310_02 ..... Done
Stopping llt on fss7310_02 ..... Done
Clean up deleted nodes information on the cluster ..... Done
Clean up deleted nodes ..... Done
Delete node completed successfully
installaccess log files and summary file are saved at:
/opt/VRTS/install/logs/installaccess-201803130635kXW
```

3 Verify that the plex states are set to NODEVICE/DISABLED.

```
[root@fss7310_01 ~]# vxclustadm nidmap
```

Name	CVM Nid	CM Nid	State
fss7310_01	2	0	Joined: Master
fss7310_03	3	2	Joined: Slave
fss7310_04	1	3	Joined: Slave

```
[root@fss7310_01 ~]# vxprint -p | grep -i nodevice
```

pl _nlm_-02	_nlm_	DISABLED	2097152	- NODEVICE	- -
pl _nlm_dcl-02	_nlm_dcl	DISABLED	67840	- NODEVICE	- -
pl test1_tier1-P02	test1_tier1-L01	DISABLED	699392	- NODEVICE	- -
pl test1_tier1-P04	test1_tier1-L02	DISABLED	699392	- NODEVICE	- -
pl test1_tier1-P06	test1_tier1-L03	DISABLED	699392	- NODEVICE	- -
pl test1_tier1_dcl-02	test1_tier1_dcl	DISABLED	67840	- NODEVICE	- -
pl test2_tier1-P02	test2_tier1-L01	DISABLED	699392	- NODEVICE	- -
pl test2_tier1-P04	test2_tier1-L02	DISABLED	699392	- NODEVICE	- -
pl test2_tier1-P06	test2_tier1-L03	DISABLED	699392	- NODEVICE	- -
pl test2_tier1_dcl-02	test2_tier1_dcl	DISABLED	67840	- NODEVICE	- -
pl test3_tier1-P02	test3_tier1-L01	DISABLED	699392	- NODEVICE	- -
pl test3_tier1-P04	test3_tier1-L02	DISABLED	699392	- NODEVICE	- -
pl test3_tier1-P06	test3_tier1-L03	DISABLED	699392	- NODEVICE	- -
pl test3_tier1_dcl-02	test3_tier1_dcl	DISABLED	67840	- NODEVICE	- -
pl test4_tier1-P02	test4_tier1-L01	DISABLED	699392	- NODEVICE	- -
pl test4_tier1-P04	test4_tier1-L02	DISABLED	699392	- NODEVICE	- -
pl test4_tier1-P06	test4_tier1-L03	DISABLED	699392	- NODEVICE	- -
pl test4_tier1_dcl-02	test4_tier1_dcl	DISABLED	67840	- NODEVICE	- -
pl test5_tier1-P02	test5_tier1-L01	DISABLED	699392	- NODEVICE	- -
pl test5_tier1-P04	test5_tier1-L02	DISABLED	699392	- NODEVICE	- -
pl test5_tier1-P06	test5_tier1-L03	DISABLED	699392	- NODEVICE	- -
pl test5_tier1_dcl-02	test5_tier1_dcl	DISABLED	67840	- NODEVICE	- -

```
[root@fss7310_01 ~]# vxdisk list | grep "failed was:"
```

```
- - emc0_2256 sfsdg failed was:emc0_2256
- - emc0_2264 sfsdg failed was:emc0_2264
- - emc0_2272 sfsdg failed was:emc0_2272
- - emc0_2280 sfsdg failed was:emc0_2280
- - emc0_2288 sfsdg failed was:emc0_2288
- - emc0_2296 sfsdg failed was:emc0_2296
- - emc0_2304 sfsdg failed was:emc0_2304
- - emc0_2312 sfsdg failed was:emc0_2312
- - emc0_2320 sfsdg failed was:emc0_2320
- - emc0_2328 sfsdg failed was:emc0_2328
```



```
- - emc0_2336 sfsdg failed was:emc0_2336
- - emc0_2344 sfsdg failed was:emc0_2344
- - emc0_2352 sfsdg failed was:emc0_2352
- - emc0_2360 sfsdg failed was:emc0_2360
```

4 Remove the affected mirrors for the volumes that are present on the system.

```
[root@fss7310_01 ~]# vxplex -f -g sfsdg -o rm dis test1_tier1-P02
[root@fss7310_01 ~]# for i in `vxprint -p | grep -i NODEVICE | awk '{print $2}`
> do
> echo "vxplex -f -g sfsdg -o rm dis $i"
> vxplex -f -g sfsdg -o rm dis $i
> done
vxplex -f -g sfsdg -o rm dis _nlm_-02
vxplex -f -g sfsdg -o rm dis _nlm__dcl-02
vxplex -f -g sfsdg -o rm dis test1_tier1-P04
vxplex -f -g sfsdg -o rm dis test1_tier1-P06
vxplex -f -g sfsdg -o rm dis test1_tier1_dcl-02
vxplex -f -g sfsdg -o rm dis test2_tier1-P02
vxplex -f -g sfsdg -o rm dis test2_tier1-P04
vxplex -f -g sfsdg -o rm dis test2_tier1-P06
vxplex -f -g sfsdg -o rm dis test2_tier1_dcl-02
vxplex -f -g sfsdg -o rm dis test3_tier1-P02
vxplex -f -g sfsdg -o rm dis test3_tier1-P04
vxplex -f -g sfsdg -o rm dis test3_tier1-P06
vxplex -f -g sfsdg -o rm dis test3_tier1_dcl-02
vxplex -f -g sfsdg -o rm dis test4_tier1-P02
vxplex -f -g sfsdg -o rm dis test4_tier1-P04
vxplex -f -g sfsdg -o rm dis test4_tier1-P06
vxplex -f -g sfsdg -o rm dis test4_tier1_dcl-02
vxplex -f -g sfsdg -o rm dis test5_tier1-P02
vxplex -f -g sfsdg -o rm dis test5_tier1-P04
vxplex -f -g sfsdg -o rm dis test5_tier1-P06
vxplex -f -g sfsdg -o rm dis test5_tier1_dcl-02
```

```
[root@fss7310_01 ~]# vxprint -p
Disk group: sfsdg
```

TY NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE
TUTIL0 PUTIL0					
pl _nlm_-01	_nlm_	ENABLED	2097152	-	ACTIVE
- -					
pl _nlm__dcl-01	_nlm__dcl	ENABLED	67840	-	ACTIVE
- -					
pl test1_tier1-P01	test1_tier1-L01	ENABLED	699392	-	ACTIVE
- -					
pl test1_tier1-P03	test1_tier1-L02	ENABLED	699392	-	ACTIVE
- -					

```

pl test1_tier1-P05      test1_tier1-L03 ENABLED 699392 - ACTIVE
- -
pl test1_tier1-03       test1_tier1      ENABLED 2098176 - ACTIVE
- -
pl test1_tier1_dcl-01   test1_tier1_dcl ENABLED 67840 - ACTIVE
- -
pl test2_tier1-P01      test2_tier1-L01 ENABLED 699392 - ACTIVE
- -
pl test2_tier1-P03      test2_tier1-L02 ENABLED 699392 - ACTIVE
- -
pl test2_tier1-P05      test2_tier1-L03 ENABLED 699392 - ACTIVE
- -
pl test2_tier1-03       test2_tier1      ENABLED 2098176 - ACTIVE
- -
pl test2_tier1_dcl-01   test2_tier1_dcl ENABLED 67840 - ACTIVE
- -
pl test3_tier1-P01      test3_tier1-L01 ENABLED 699392 - ACTIVE
- -
pl test3_tier1-P03      test3_tier1-L02 ENABLED 699392 - ACTIVE
- -
pl test3_tier1-P05      test3_tier1-L03 ENABLED 699392 - ACTIVE
- -
pl test3_tier1-03       test3_tier1      ENABLED 2098176 - ACTIVE
- -
pl test3_tier1_dcl-01   test3_tier1_dcl ENABLED 67840 - ACTIVE
- -
pl test4_tier1-P01      test4_tier1-L01 ENABLED 699392 - ACTIVE
- -
pl test4_tier1-P03      test4_tier1-L02 ENABLED 699392 - ACTIVE
- -
pl test4_tier1-P05      test4_tier1-L03 ENABLED 699392 - ACTIVE
- -
pl test4_tier1-03       test4_tier1      ENABLED 2098176 - ACTIVE
- -
pl test4_tier1_dcl-01   test4_tier1_dcl ENABLED 67840 - ACTIVE
- -
pl test5_tier1-P01      test5_tier1-L01 ENABLED 699392 - ACTIVE
- -
pl test5_tier1-P03      test5_tier1-L02 ENABLED 699392 - ACTIVE
- -
pl test5_tier1-P05      test5_tier1-L03 ENABLED 699392 - ACTIVE
- -
pl test5_tier1-03       test5_tier1      ENABLED 2098176 - ACTIVE

```

```
--  
pl test5_tier1_dcl-01    test5_tier1_dcl ENABLED 67840    -    ACTIVE  
--
```

- 5 Remove the affected disks from the disk group by using the `vxdg rmdisk` command and from all the nodes in the cluster by using the `vxdisk rm` command.

```
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2288  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2272  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2280  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2296  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2304  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2312  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2320  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2328  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2336  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2344  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2352  
[root@fss7310_01 bin]# vxdg -g sfsdg rmdisk emc0_2360  
[root@fss7310_01 bin]# for i in `vxdisk list | grep -i error | awk '{print $1}'`;  
do vxdisk rm $i; done  
[root@fss7310_03 ~]# for i in `vxdisk list | grep -i error | awk '{print $1}'`;  
do vxdisk rm $i; done  
[root@fss7310_04 ~]# for i in `vxdisk list | grep -i error | awk '{print $1}'`;  
do vxdisk rm $i; done
```

- 6 Run the `addnode` command for the cluster by using IP.

7 Add the disks from the newly added node in the pool that is already present.

```
[root@fss7310_01 scripts]# /opt/VRTSnas/clish/bin/clish -u admin -c  
"storage disk format emc0_2257,emc0_2265,emc0_2273,emc0_2281,emc0_2289,emc0_2297,emc0_2305,  
emc0_2313,emc0_2321,emc0_2329,emc0_2337,emc0_2345,emc0_2353,emc0_2361"
```

```
You may lose all the data on the disk, do you want to continue (y/n, the default is n):y  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2257 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2265 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2273 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2281 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2289 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2297 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2305 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2313 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2321 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2329 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2337 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2345 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2353 has been formatted successfully.  
ACCESS Disk SUCCESS V-493-10-4 disk format: emc0_2361 has been formatted successfully.
```

```
[root@fss7310_01 scripts]# /opt/VRTSnas/clish/bin/clish -u admin -c "storage pool  
adddisk pool1 emc0_2257,emc0_2265,emc0_2273,emc0_2281,emc0_2289,emc0_2297,emc0_2305,  
emc0_2313,emc0_2321,emc0_2329,emc0_2337,emc0_2345,emc0_2353,emc0_2361"
```

```
ACCESS Pool SUCCESS V-493-10-2914 Successfully added disks to pool
```

8 Mirror the volume by using the `storage addmirror` command.

```
fss7310> storage fs list
```

FS	STATUS	SIZE	LAYOUT	MIRRORS	COLUMNS	USE%	USED	NFS SHARED	CIFS SHARED	FTP SHARED	SECONDARY TIER
test1	online	1.00G	striped	1	3	10%	103M	no	no	no	no
test2	online	1.00G	striped	1	3	10%	103M	no	no	no	no
test3	online	1.00G	striped	1	3	10%	103M	no	no	no	no
test4	online	1.00G	striped	1	3	10%	103M	no	no	no	no
test5	online	1.00G	striped	1	3	10%	103M	no	no	no	no

```
fss7310> storage fs addmirror test1 pool1
100% [#] Adding mirror to filesystem
ACCESS fs SUCCESS V-493-10-2131 Added mirror for fs test1
fss7310> storage fs addmirror test2 pool1
100% [#] Adding mirror to filesystem
ACCESS fs SUCCESS V-493-10-2131 Added mirror for fs test2
fss7310> storage fs addmirror test3 pool1
100% [#] Adding mirror to filesystem
ACCESS fs SUCCESS V-493-10-2131 Added mirror for fs test3
fss7310> storage fs addmirror test4 pool1
100% [#] Adding mirror to filesystem
ACCESS fs SUCCESS V-493-10-2131 Added mirror for fs test4
```

9 Mirror the _nlm_ volume by using the vxassist mirror command.

```
[root@fss7310_01 bin]# vxassist -b -g sfsdg mirror _nlm_
```

```
[root@fss7310_01 bin]# vxprint _nlm_
```

Disk group: sfsdg

TY	NAME	ASSOC	KSTATE	LENGTH	PLOFFS	STATE
TUTIL0	PUTIL0					
v	_nlm_	fsgen	ENABLED	2097152	-	ACTIVE
ATT1	-					
pl	_nlm_-01	_nlm_	ENABLED	2097152	-	ACTIVE
-	-					
sd	emc0_2255-01	_nlm_-01	ENABLED	2097152	0	-
-	-					
pl	_nlm_-02	_nlm_	ENABLED	2097152	-	
TEMPRMSD	ATT	-				
sd	emc0_2257-01	_nlm_-02	ENABLED	2097152	0	-
-	-					
dc	_nlm__dco	_nlm_	-	-	-	-
-	-					
v	_nlm__dcl	gen	ENABLED	67840	-	ACTIVE
-	-					
pl	_nlm__dcl-01	_nlm__dcl	ENABLED	67840	-	ACTIVE
-	-					
sd	emc0_2255-02	_nlm__dcl-01	ENABLED	67840	0	-
-	-					
sp	_nlm__cpmap	_nlm_	-	-	-	-
-	-					

Replacing a disk

In some cases, you may need to replace an existing disk. This section describes the steps for replacing a disk.

Note: You can replace a disk in a storage pool from a master or a subordinate node. From a Cluster Volume Manager (CVM) master node, you can replace DAS and SAN disks. From a subordinate node, you can replace only DAS disks.

To replace a disk

- 1 Remove the disk that needs to be replaced from the array side.
- 2 Add the new disk to the system from the array side.
- 3 Run the following command on all the nodes in the cluster to eliminate the old disk from the Veritas Volume Manager (VxVM) view.

```
#vxdisk rm <old-disk-name>
```

- 4 Run the following command on the node for which you want to replace the disk.

```
# vxdisk scandisks
```

- 5 Initialize the new disk that has been added to the cluster by using the `vxdisksetup` command.

```
#!/etc/vx/bin/vxdisksetup -fi <new-disk-name>
```

- 6 Apply the tag to the newly added device that is similar to the Veritas Access pool name on the node where the failed disk resides.

```
# vxdisk settag site=<pool-name> <new-disk-name>
```

- 7 Run the `vxdiskadm` command and choose option #5 to replace a failed disk with a new disk on the node where the failed disk resides.

```
#vxdiskadm
```

Note: If the disk replacement is triggered from a subordinate node, the `vxrecover` command fails.

- 8 If the disk replacement is triggered from a subordinate node, run the following command from the slave node for all the affected volumes.

```
#vxrecover -b -c -s <vol-name>
```

- 9 Rename the newly added disk to the disk access name.

```
#vxedit -g <dg-name> rename <old-disk-name> <new-disk-name>
```

- 10 Rename the subdisks as per the newly added disk name.

```
#vxedit -g <dg-name> rename <old-subdisk-name> <new-subdisk-name>
```


Example: Disk replacement from a CVM master node

This example describes the procedure for replacing the `emc0_2255` disk with `emc0_2263` disk. The `emc0_2263` disk has been excluded and it is added later to simulate the disk addition.

Replacing a disk from a CVM master node

- 1 Run the `vxddmpadm exclude` command to remove the `emc0_2255` disk.

```
# vxddmpadm exclude dmpnodename=emc0_2255
```

- 2 Run the the `vxddmpadm include` command to include the `emc0_2263` disk.

```
# vxddmpadm include dmpnodename=emc0_2263
```

Note: You can run the `vxdisk scandisks` command to scan the disks.

- 3 Run the `settag` command to apply the tag to the underlying disk.

```
# vxdisk settag emc0_2263 tag=pool1
```

4 Run the `vxdiskadm` command and choose option #5 to replace the failed disk.

```
[root@fss7310_01 ~]# vxdiskadm
Volume Manager Support Operations
Menu:: VolumeManager/Disk
1 Add or initialize one or more disks
2 Encapsulate one or more disks
3 Remove a disk
4 Remove a disk for replacement 5 Replace a failed or removed disk
6 Mirror volumes on a disk
7 Move volumes from a disk
8 Enable access to (import) a disk group
9 Remove access to (deport) a disk group
10 Enable (online) a disk device
11 Disable (offline) a disk device
12 Mark a disk as a spare for a disk group
13 Turn off the spare flag on a disk
14 Unrelocate subdisks back to a disk
15 Exclude a disk from hot-relocation use
16 Make a disk available for hot-relocation use
17 Prevent multipathing/Suppress devices from VxVM's view
18 Allow multipathing/Unsuppress devices from VxVM's view
19 List currently suppressed/non-multipathed devices
20 Change the disk naming scheme
21 Change/Display the default disk layouts
22 Dynamic Reconfiguration Operations
list List disk information
```

Select an operation to perform: 5

```
Replace a failed or removed disk
Menu:: VolumeManager/Disk/ReplaceDisk
```

Use this menu operation to specify a replacement disk for a disk that you removed with the "Remove a disk for replacement" menu operation, or that failed during use. You will be prompted for a disk name to replace and a disk device to use as a replacement. You can choose an uninitialized disk, in which case the disk will be initialized, or you can choose a disk that you have already initialized using the Add or initialize a disk menu operation.

Select a removed or failed disk [<disk>,list,q,?] list

```

Disk group: sfsdg
DM NAME DEVICE TYPE PRIVLEN PUBLLEN STATE
dm emc0_2255 - - - - NODEVICE

Select a removed or failed disk [<disk>,list,q,?] emc0_2255

The following devices are available as replacements:
emc0_2263
You can choose one of these devices to replace emc0_2255.
Choose "none" to initialize another device to replace emc0_2255.
Choose a device, or select "none" [<device>,none,q,?]
(default: emc0_2263) emc0_2263
VxVM INFO V-5-2-382
The requested operation is to use the initialized device emc0_2263
to replace the removed or failed disk emc0_2255 in disk group sfsdg.
Continue with operation? [y,n,q,?] (default: y) y
Use FMR for plex resync? [y,n,q,?] (default: n)
VxVM INFO V-5-2-282 Replacement of disk emc0_2255 in group
sfsdg with disk device emc0_2263 completed successfully.
Replace another disk? [y,n,q,?] (default: n)

```

5 Rename the disk as per the disk access name to avoid the vx dg issue.

```

# vxedit -g sfsdg rename emc0_2255 emc0_2263

# vxdisk list | grep emc0_2263 emc0_2263 auto:cdsdisk emc0_2263
sfsdg online shared

```

6 Rename the subdisks as per disk access name.

```

# vxedit -g sfsdg rename emc0_2255-03 emc0_2263-03
# vxedit -g sfsdg rename emc0_2255-02 emc0_2263-02

[root@fss7310_01 ~]# vxprint -pvs | grep -i 2263
sd emc0_2263-02 vol1-P01 ENABLED 699136 0 - - -
sd emc0_2263-03 vol1_dcl-01 ENABLED 67840 0 - - -
[root@fss7310_01 ~]# vxprint -pvs | grep -i 2255
[root@fss7310_01 ~]#

```

Example: Disk replacement from a subordinate node

This example describes the procedure for replacing the emc0_2273 disk with emc0_2305 disk. The emc0_2263 disk has been excluded and it is added later to simulate the disk addition.

To replace a disk from a subordinate node

- 1 Run the `vxddmpadm exclude` command to remove the `emc0_2273` disk.

```
# vxddmpadm exclude dmpnodename=emc0_2273
```

- 2 Run the `vxddmpadm include` command to include the `emc0_2305` disk.

```
# vxddmpadm include dmpnodename=emc0_2305
```

Note: You can run the `vxdisk scandisks` command to scan the disks.

- 3 Run the `vxdisk rm` command from the remaining nodes in the cluster:

```
[root@fss7310_02 ~]# vxdisk rm emc0_2273  
[root@fss7310_01 ~]# vxdisk rm emc0_2273
```

Note: If you encounter an error, wait for a few minutes and run the command again.

- 4 Run the `settag` command to apply the tags to the underlying disk:

```
# vxdisk settag emc0_2305 tag=pool1
```

5 Run the `vxdiskadm` command and choose option #5 to replace the failed disk.

```
[root@fss7310_01 ~]# vxdiskadm
```

```
Volume Manager Support Operations
```

```
Menu:: VolumeManager/Disk
```

```
1 Add or initialize one or more disks
2 Encapsulate one or more disks
3 Remove a disk
4 Remove a disk for replacement 5 Replace a failed or removed disk
6 Mirror volumes on a disk
7 Move volumes from a disk
8 Enable access to (import) a disk group
9 Remove access to (deport) a disk group
10 Enable (online) a disk device
11 Disable (offline) a disk device
12 Mark a disk as a spare for a disk group
13 Turn off the spare flag on a disk
14 Unrelocate subdisks back to a disk
15 Exclude a disk from hot-relocation use
16 Make a disk available for hot-relocation use
17 Prevent multipathing/Suppress devices from VxVM's view
18 Allow multipathing/Unsuppress devices from VxVM's view
19 List currently suppressed/non-multipathed devices
20 Change the disk naming scheme
21 Change/Display the default disk layouts
22 Dynamic Reconfiguration Operations
list List disk information
```

```
? Display help about menu
```

```
?? Display help about the menuing system
```

```
q Exit from menus
```

```
Select an operation to perform: 5
```

```
Replace a failed or removed disk
```

```
Menu:: VolumeManager/Disk/ReplaceDisk
```

Use this menu operation to specify a replacement disk for a disk that you removed with the "Remove a disk for replacement" menu operation, or that failed during use. You will be prompted for a disk name to replace and a disk device to use as a replacement.

You can choose an uninitialized disk, in which case the disk will be initialized, or you can choose a disk that you have already initialized using the Add or initialize a disk menu operation. Select a removed or failed disk [<disk>,list,q,?] list
Disk group: sfsdg

DM NAME	DEVICE	TYPE	PRIVLEN	PUBLLEN	STATE
dm emc0_2273	-	-	-	-	NODEVICE

Select a removed or failed disk [<disk>,list,q,?] emc0_2273

The following devices are available as replacements:

emc0_2305

You can choose one of these devices to replace emc0_2255. Choose "none" to initialize another device to replace emc0_2255.

Choose a device, or select "none" [<device>,none,q,?]

(default: emc0_2305) emc0_2305

VxVM INFO V-5-2-382

The requested operation is to use the initialized device emc0_2305 to replace the removed or failed disk emc0_2273 in disk group sfsdg.

Continue with operation? [y,n,q,?] (default: y)

Use FMR for plex resync? [y,n,q,?] (default: n) VxVM vxrecover

ERROR V-5-1-16084 Disk group: sfsdg is shared. The command can be executed only on the CVM master. Use -c option to recover all the shared disk groups from slaves.

VxVM INFO V-5-2-282 Replacement of disk emc0_2273 in group sfsdg with disk device emc0_2305 completed successfully.

Replace another disk? [y,n,q,?] (default: n)

6 Run the following command to trigger a recovery for the affected volumes.

```
# vxrecover -b -c -s voll
```

7 Rename the disk to the disk access name to avoid the vxdbg issue.

```
# vxedit -g sfsdg rename emc0_2273 emc0_2305

# vxdisk list | grep emc0_2305
emc0_2305 auto:cdsdisk emc0_2305 sfsdg online shared
```

8 Rename the subdisks as per the naming convention that is followed for the newly added disk.

```
# vxedit -g sfsdg rename emc0_2273-02 emc0_2305-02
# vxedit -g sfsdg rename emc0_2273-03 emc0_2305-03

# vxprint -pvs | grep -i emc0_2305
sd emc0_2305-02 vol1-P02 ENABLED 699136 0 - - -
sd emc0_2305-03 vol1_dcl-02 ENABLED 67840 0 - - -
# vxprint -pvs | grep -i emc0_2273
```

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 an episodic replication job

To synchronize an enabled episodic replication job

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

```
Replication> episodic job sync job_name
```

<i>job_name</i>	Specify the name of the episodic 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 Log on to the Veritas Access management console IP using the admin user credentials.

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

Note: The PCI ID feature is deprecated in this release.

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 the LTR upgrade

This chapter includes the following topics:

- [Locating the log files for troubleshooting the LTR upgrade](#)
- [Troubleshooting pre-upgrade issues for LTR](#)
- [Troubleshooting post-upgrade issues for LTR](#)

Locating the log files for troubleshooting the LTR upgrade

Log files for the LTR upgrade process are located at:

- `/opt/VRTSnas/log/ltr_preupgrade.log`
- `/opt/VRTSnas/log/ltr_post_upgrade.log`

Troubleshooting pre-upgrade issues for LTR

Following are the pre-upgrade scenarios with the recommended workaround:

- Failure in provisioning of `odd_cache_fs`
Workaround: Ensure that sufficient storage space is available in default pool(s) that is configured for ObjectAccess.
- Failure in backup of configuration file
Workaround: Ensure that the `odd_vipgrp_map.conf` file is successfully copied to `/shared/openedup/`.
- Failure in backup of XML configuration files

Workaround: Ensure that all the OpenDedup volume configuration XML files are successfully copied to `/shared/opendedup/sdfs/`.

- Failure in shutdown of OpenDedup volume(s)

Workaround:

- Ensure that the volume is unmounted by using the `mount -t fuse` command.
- Ensure all the jsvc processes are stopped. If the processes are not stopped, use the `kill` command to stop them.
- Shutdown of the OpenDedup volume may result in java stack trace to be printed on console. This can be ignored as long as the volume is unmounted.

- Failure in backup of the OpenDedup cache data

Workaround: Ensure that the size of the local cache location, that is, `/opt/sdfs/volumes/<vol-name>` is equivalent to the shared cache location, that is, `/vx/odd_cache_fs/volumes/<vol-name>`.

Troubleshooting post-upgrade issues for LTR

Following are the post-upgrade scenarios with the recommended workaround:

- Failure in creation of soft-links

Workaround: Ensure that the soft links are created for the following paths:

- `/etc/sdfs -> /shared/opendedup/sdfs`
- `/opt/sdfs -> /vx/odd_cache_fs/sdfs`

- Failure in configuration restore

Workaround: Ensure that the `odd_vipgrp_map.conf` file has been successfully copied from `/shared/opendedup/` to `/opt/VRTSnas/conf/`.

- Failure in online of OpenDedup volumes

Workaround:

- Log locations of the OpenDedup volumes:
 - `/opt/VRTSnas/log/odd.log`
 - `/var/log/sdfs/<vol-name>.log`
- Manual mount of OpenDedup volume can be verified by using the following commands from the Veritas Access command-line interface:
 - `Opendedup volume list`
 - `Opendedup volume online <vol-name>`

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/VRTSnas/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 firewalld 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.

```
# /opt/VRTSnas/pysnas/bin/vamgmt -h
# /opt/VRTSnas/pysnas/bin/vamgmt status

# ps -ef | grep node
```

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

```
/opt/VRTSnas/isagui/ext_modules/node /opt/VRTSnas/isagui/application/server.js production
```

- 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/VRTSnas/pysnas/bin/isagui_cluster_perf.py --full
```

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

```
/opt/VRTSnas/log/isagui_cluster_perf.log
```

Index

A

- a disk
 - replacing 47
- a node
 - replacing 34
- about
 - common recovery procedures 24
 - event logs 10
 - job resynchronization 55
 - monitoring commands 17
 - services command 25
 - shell-activity logs 10
- an Ethernet interface card (offline mode)
 - replacing 33
- an Ethernet interface card (online mode)
 - replacing 30

C

- CIFS
 - setting the log level 10
- cloud tiering
 - log locations 59
 - troubleshooting 58
- Cluster
 - Excluding PCI IDs 62
- common recovery procedures
 - about 24
- configuring
 - job resynchronization 56
- CPU utilization report
 - generating 19

D

- debugging information
 - retrieving and sending 12
- debugging options
 - setting for NetBackup 11
- device utilization report
 - generating 19

- displaying
 - node-specific network traffic details 21

E

- event logs
 - about 10
- Excluding PCI IDs
 - cluster 62
- exporting
 - network traffic details 21

F

- fsck
 - unable to run if operating system partition is corrupted 63

G

- general techniques
 - troubleshooting 7
- general tips
 - troubleshooting process 6
- generating
 - CPU utilization report 19
 - device utilization report 19

I

- installation
 - common failures 61
- installation logs
 - viewing 61

J

- job resynchronization
 - about 55
 - configuring 56

L

- locating the log files
 - troubleshooting the LTR upgrade 64

- log locations
 - cloud tiering 59

M

- monitoring
 - installation logs 61
 - processor activity 17
- monitoring commands
 - about 17

N

- NetBackup client log levels
 - setting 11
- NetBackup debugging options
 - setting 11
- network
 - testing connectivity 27
- network traffic details
 - exporting 21
- node-specific network traffic details
 - displaying 21

O

- operating system
 - unable to repair using fsck 63

P

- patch release
 - uninstalling 56
- processor activity
 - monitoring 17

R

- reading or writing data from the cloud tier
 - troubleshooting 58
- recovering
 - from a non-graceful shutdown 27
- replacing
 - a disk 47
 - a node 34
 - an Ethernet interface card (offline mode) 33
 - an Ethernet interface card (online mode) 30
- replication
 - speeding up 55
- restarting
 - servers 24

- retrieving
 - debugging information 12

S

- sending
 - debugging information 12
- servers
 - restarting 24
- services command
 - about 25
 - using 26
- setting
 - CIFS log level 10
 - NetBackup client log levels 11
- shell-activity logs
 - about 10
- shutdown
 - recovering from a non-graceful 27
- software upgrade
 - uninstalling 56

T

- testing
 - network connectivity 27
- traceroute
 - troubleshooting with 28
- traceroute command
 - using 29
- troubleshooting
 - about 6
 - general procedures 9
 - issues when reading or writing data from the cloud tier 58
- troubleshooting process
 - general techniques 7
 - general tips 6
- troubleshooting the LTR upgrade
 - locating the log files 64

U

- uninstalling
 - patch release or software upgrade 56
- using
 - services command 26
 - traceroute command 29

V

Veritas Access log files

viewing 9

viewing

installation logs 61

Veritas Access log files 9