

Veritas Access Appliance

8.4 Troubleshooting Guide

Access Appliance Troubleshooting Guide

Last updated: 2025-03-26

Legal Notice

Copyright © 2025 Veritas Technologies LLC. All rights reserved.

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

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

<https://www.veritas.com/about/legal/license-agreements>

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

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

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

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<https://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

https://www.veritas.com/content/support/en_US/dpp.Appliances.html

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:

APPL.docs@veritas.com

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

<https://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	8
	About general troubleshooting procedures	8
	Viewing the Access Appliance log files	8
	About event logs	10
	Setting the CIFS log level	10
	Setting the NetBackup client log levels and debugging options	11
	Retrieving and sending debugging information	12
	Collecting time-based and archived logs	16
	Collecting logs for cluster nodes	19
Chapter 3	Monitoring Access Appliance	20
	About monitoring Access Appliance operations	20
	Monitoring hardware components	20
Chapter 4	Common recovery procedures	22
	About common recovery procedures	22
	Restarting servers	23
	Restarting cluster services	23
	Bringing services online	24
	Using the services command	25
	Recovering from a non-graceful shutdown	26
	Testing the network connectivity	26
	Troubleshooting with traceroute	27
	Using the traceroute command	28
	Collecting the metasave image of a file system	28
	Replacing an Ethernet interface card (online mode)	29
	Replacing an Access Appliance node	32
	Speeding up episodic replication	33
	About synchronizing an episodic replication job	33

	Synchronizing an episodic replication job	34
	Uninstalling a patch release or software upgrade	34
Chapter 5	Troubleshooting the Access Appliance cloud as a tier feature	36
	Troubleshooting tips for cloud tiering	36
	Issues when reading or writing data from the cloud tier	36
	Log locations for checking for cloud tiering errors	37
Chapter 6	Troubleshooting Access Appliance installation and configuration issues	38
	How to find the management console IP	38
	Viewing the installation logs	39
	Installation fails and does not complete	40
Chapter 7	Troubleshooting Access Appliance CIFS issues	41
	User access is denied on a CTDB directory share	41
Chapter 8	Troubleshooting Access Appliance GUI startup issues	42
	Resolving GUI startup issues	42
Chapter 9	Troubleshooting Veritas Data Deduplication issues	45
	Log locations for the Veritas Data Deduplication server	45
Index		46

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 Access Appliance 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 `maintenance 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 *Access Appliance Release Notes*. The release notes contain a list of known issues for Access Appliance 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:

- 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 showall` command.
Set up a process 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 Access Appliance log files](#)
- [About event logs](#)
- [Setting the CIFS log level](#)
- [Setting the NetBackup client log levels and debugging options](#)
- [Retrieving and sending debugging information](#)
- [Collecting time-based and archived logs](#)
- [Collecting logs for cluster nodes](#)

About general troubleshooting procedures

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

Viewing the Access Appliance log files

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

To view the Access Appliance log files when the appliance lockdown mode is not set

- 1 Use `support elevate` to log in with root privileges.
- 2 Navigate to the `/opt/VRTSnas/log` directory.

If your appliance is in lockdown mode, you need assistance from Veritas Support to access the log files.

To view the Access Appliance log files when the appliance lockdown mode is set

- 1 Log on to the node-level CLI on any node in the cluster.
- 2 Run the `show serial-number` command to get the serial number for a specific node.
- 3 If you have permission to generate access key on the SHI portal, then go to the portal to generate access key with the serial number that you got in step 2. For more information, refer to the *System Health Insights User Guide*.

If you do not have the permission to generate the access key on the SHI portal, open a ticket with Veritas Support to generate an access key. Mention the serial number in the Support case. The Veritas Support will generate an access key with the serial number provided. Set an access passphrase which will be required later to elevate to root.

- 4 Log on to the Access Appliance shell on the same node as in step 2.
- 5 When your representative asks you to, run the `support unlock` command. You are prompted to enter the maintenance password. Enter the maintenance password and press **Enter**. You are prompted to enter an access key. Enter the access key that you got in step 3. You are prompted to enter an access passphrase. Enter the access passphrase that you set in step 3. Press **Enter** to unlock the root shell access to the current node (all other nodes remain locked).
- 6 Run the `support elevate` command. Enter the maintenance password. You are prompted to enter an access passphrase. Enter the access passphrase set in step 3. Press **Enter** to get into the root shell.
- 7 To view the logs, navigate to the `/opt/VRTSnas/log` directory.
- 8 Run the `support lock` command on a specific node to lock that node. If no manual lock is issued, the node is locked automatically after 12 hours. All the current users are removed from the root shell in a single node.

Note: The access key expires in 2 hours.

About event logs

In addition to the system log, each Access Appliance 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 Access Appliance 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 SSH to log in to the node using admin privileges.
- 2 Type `support elevate` and when prompted enter the password for the maintenance account.
- 3 Navigate to the `/opt/VRTSnas/log` directory.

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

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

Setting the CIFS log level

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

Use the `Support>man debuginfo` command to see the 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 Access Appliance 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 Access Appliance 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 Access Appliance debugging information from the Access Appliance node and send the information to a server using an external FTP or SCP server.

If a node is part of the 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

You can 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*.

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 Access Appliance product is installed.
- `install`: This module collects the install logs from the CPI install log directory `/opt/VRTS/install/logs/`.
- `nas-procstacks`: This module collects the stack trace for all the daemons that are in running state.
- `netbackup`: This module collects logs about NetBackup from nodes.
- `appliance`: This module collects the logs related to the appliance.
- `sds`: This module collects SDS logs.
- `api_gateway`: This module collects API gateway logs.
- `upgrade`: This module collect the upgrade logs for all the cluster nodes.
- `backup`: This module collects NetBackup client logs and other diagnostic information.
- `vdd`: This module collects the deduplication logs provided dedupe is configured on the cluster nodes.
- `default`: This module collects logs from all modules of the cluster nodes except the `os` module.
- `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 `/root/log` 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 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 Access Appliance product is installed.
- *install*: This module collects the install logs from the CPI install log directory `/opt/VRTS/install/logs/`.
- *nas-procstacks*: This module collects the stack trace for all the daemons that are in running state.
- *netbackup* : This module collects logs about NetBackup from nodes.
- *appliance*: This module collects the logs related to the appliance.
- *sds*: This module collects SDS logs.
- *api_gateway*: This module collects API gateway logs.
- *upgrade*: This module collect the upgrade logs for all the cluster nodes.
- *backup*: This module collects NetBackup client logs and other diagnostic information.
- *vdd*: This module collects the deduplication logs provided dedupe is configured on the cluster nodes.
- *default*: This module collects logs from all modules of the cluster nodes except the *os* module.
- *all*: This module collects information for all the above modules from the cluster nodes.

Collecting time-based and archived logs

You can collect logs for a specific duration for the **sds** and **backup** modules. Collecting logs for a specific duration reduces the time taken to collect and upload the logs to a remote server. You can also collect archived logs for all the modules.

The following command shows the options for collecting time-based and archived logs:

```
debuginfo upload nodename debug-URL module  
archivedLogs={on|off|vxdefault} startDate endDate tarName
```

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
```

<i>module</i>	<p>Specifies the values for <i>module</i>.</p> <p>The following options are supported:</p> <ul style="list-style-type: none"> ■ <code>service-status</code>: Current status of services that are running in the cluster. ■ <code>performance</code>: Performance logs for the cluster nodes. ■ <code>common-logs</code>: Logs that are common across multiple modules for the cluster nodes. ■ <code>nas</code>: Product information from the cluster nodes. ■ <code>os</code>: Logs for kernel and user dumps. ■ <code>sos-report</code>: Uses RHEL utility to collect troubleshooting data for the operating system. The <code>sosreport</code> is collected for all the loaded modules except for the <code>selinux</code> module. ■ <code>explorer</code>: The VxExplorer utility collects logs and environment data from all the servers where the Access Appliance product is installed. ■ <code>install</code>: Install logs from the CPI install log directory <code>/opt/VRTS/install/logs/</code>. ■ <code>nas-procstacks</code>: Stack trace for all the daemons that are in running state. ■ <code>appliance</code>: Logs related to the appliance. ■ <code>sds</code>: SDS logs. ■ <code>upgrade</code>: Upgrade logs for all the cluster nodes. ■ <code>backup</code>: NetBackup client logs and other diagnostic information. ■ <code>vdd</code>: Deduplication logs provided Veritas Data Deduplication server is configured for the cluster. ■ <code>default</code>: Logs from <code>service-status</code>, <code>nas</code>, <code>explorer</code>, <code>install</code>, <code>nas-procstacks</code>, <code>appliance</code>, <code>sds</code>, <code>upgrade</code>, and <code>backup</code> modules of the cluster nodes. This is the default option. ■ <code>all</code>: This module collects information for all the above modules from the cluster nodes.
<code>archivedLogs</code>	<p>Collects historical logs. Options are <code>on</code>, <code>off</code>, and <code>vxdefault</code>.</p>
<code>startDate</code>	<p>Start date in <code>mm/dd/yy</code> format if you want to collect logs for a specific duration.</p> <p>Time-based logs are collected only for <code>nas</code> and <code>sds</code> modules. Logs for <code>sds</code> and <code>nas</code> modules from the archived directory are also included.</p>
<code>endDate</code>	<p>End date in <code>mm/dd/yy</code> format if you want to collect logs for a specific duration.</p> <p>Time-based logs are collected only for <code>nas</code> and <code>sds</code> modules. Logs for <code>sds</code> and <code>nas</code> modules from the archived directory are also included.</p>

tarName Name of the tar file that includes all the collected logs.

The following example shows how to upload all debug information to an FTP server including the archived logs:

```
Support> debuginfo upload node1_1  
ftp://admin@ftp.docserver.veritas.com/patches/ all archivedLogs=on
```

The following example shows how to upload all debug information to an FTP server excluding the archived logs:

```
Support> debuginfo upload node1_1  
ftp://admin@ftp.docserver.veritas.com/patches/ all archivedLogs=off
```

The following example shows how to upload all debug logs from all node for all modules for a specific duration. The logs for only **sds** and **backup** modules are collected for the specified duration; all logs are collected for all the other modules.

```
Support> debuginfo upload all file:///log/ all archivedLogs=off  
06/12/2023 06/25/2023
```

Collecting logs for cluster nodes

You can collect logs for an individual node in the cluster for error analysis and troubleshooting. You can choose to collect logs for all the components to get a comprehensive view of the system, or you can collect logs for specific components that have an issue.

To select logs individually

- 1 Login to <https://managementconsoleIP:14161/> using admin user credentials where *managementconsoleIP* is the management console IP address.
- 2 Go to **Settings > Diagnostics**.
- 3 Click **Generate log package**.
- 4 On the **Generate log package** page, select **Advanced** tab where you can select individual log files to include in the log package.
- 5 On the **Advanced** tab, under **Log settings**, expand the nodes and select the log files.
- 6 Under **Package options**, specify a name for the log package and optionally the case number if provided by Veritas Support.
- 7 Click **Generate**.

Monitoring Access Appliance

This chapter includes the following topics:

- [About monitoring Access Appliance operations](#)
- [Monitoring hardware components](#)

About monitoring Access Appliance operations

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

As you work with Access Appliance, 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 hardware components

Use the following commands to monitor the hardware components:

- `show hardware-health appliance component=`
Use this command to view the health of the hardware components and the primary and expansion shelves. The options for the component are **All/Product/Fan/Power/Temperature/CPU/Network/PCI/Firmware/Partition/RAID/Disk/DIMM/Certificate/CMOSBattery/DIMMPopulation/StorageStatus/Connection)[all]**
To view the health of the primary shelf, use the `show hardware-health primary-shelf component=` command. For example:

```
show hardware-health primary-shelf component=FAN
```

To view the health of the expansion shelf use the `show hardware-health expansion-shelf component= tray-id=` command where **tray-id** is the ID of the expansion shelf. For example:

```
show hardware-health expansion-shelf component=FAN tray-id=1
```

- `show hardware-errors`

Use this command to view the errors that are related to the hardware components.

Common recovery procedures

This chapter includes the following topics:

- [About common recovery procedures](#)
- [Restarting servers](#)
- [Restarting cluster services](#)
- [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 Access Appliance node](#)
- [Speeding up episodic 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 Access Appliance.

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 Access Appliance 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.
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.

Restarting cluster services

There might be situations where you need to stop both the cluster nodes for maintenance tasks. When you stop both the cluster nodes, all the cluster services are stopped. You can stop each node individually by using the `cluster stop nodename` command, or you can stop both the cluster nodes using the `cluster stop all` command. When you stop both the nodes, either successively or

simultaneously, all the cluster services are stopped and you don't have access to the command-line interface. To start the cluster services, you must use the following procedure.

To start the cluster services:

- 1 Use SSH to log on to one of the nodes using the admin credentials.
- 2 Elevate to root by using the `support elevate` command.
- 3 Run the following command to start the cluster services:

```
/opt/VRTSnas/scripts/cluster/cluster_config.sh start all
```

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
- 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
iSCSITarget     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
iSCSITarget     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 <i>packets</i> field is omitted, five packets are sent to the destination by default. The <i>packets</i> 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 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 to create a consistent metasave image, if the file system is not offline on all the cluster nodes, a warning appears which has to be confirmed before you proceed. You can 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. You can run other Access Appliance 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 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 Access Appliance 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 Access Appliance node

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

It is not recommended to delete the node where the management console is running. however, if you want to delete this node, you should first switch the management console on to other node and then attempt to the delete node.

To switch the management console node:

- 1 Go to the Access command-line interface.
- 2 Run the `network ip addr show` command and note the console IP address.
- 3 Run the `network ip addr online console_ip target_node` command.

You will exit the Access command-line interface. You need to log in again.

To replace an Access Appliance 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 another node.
- 2 Power-off the node to delete. Once the node is powered-off, run the `cluster del` command for the node that is to be replaced.

```
rvclus>cluster del node-name
```

- 3 Reimage the deleted node with the same ISO as before. Upgrade the node if required to bring the deleted node to the same version as the cluster node.
- 4 Run the `cluster add` command to add reimaged node to the cluster.

```
rvclus>cluster add node-ip
```

Speeding up episodic replication

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

About synchronizing an episodic replication job

The first time an episodic replication job is run, Access Appliance 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 episodic replication cannot make the incremental changes. Examples of this situation include:

- When episodic 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 an episodic 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 episodic 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 episodic 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 episodic 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 episodic job sync` command to leverage the existing data at the destination and avoid requiring a full copy. The `replication episodic 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 episodic job sync` command to recover from this state, follow these steps. First, disable the job, then enable the job again. Then, use the `replication episodic job sync` command to synchronize the job.

Note: Synchronization cannot be performed on a paused episodic 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
```

job_name 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 Access Appliance, refer to the *Veritas Access Installation Guide*.

Troubleshooting the Access Appliance 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 Access Appliance server.
- 2 Make sure that the system time is accurate on the Access Appliance 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/storage_tier_X.log
- /var/log/messages

Troubleshooting Access Appliance 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](#)

How to find the management console IP

To find the management console IP:

- 1 Identify on which node the management console is running:

```
hastatus -sum | grep Management
```

The output displays the node on which the management console is online. The management console is online on only one of the nodes.

- 2 On the node where the management console is online, run the following command to get the IP address:

```
ifconfig -a eth1
```

- 3 To access the Access command-line interface, use Secure Shell (ssh) and connect to the management console IP with admin user credentials:

```
ssh admin@managementconsoleIP
```

- 4 To access the Access Appliance GUI, use the following URL:

```
https://managementconsoleIP:14161/
```

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 Access Appliance installation logs

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

Access Appliance 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 Access Appliance specific installation logs.
 For example:
`/opt/VRTS/install/logs/installaccess-201602021544AsJ`

Access Appliance service group configuration logs `/opt/VRTSsnas/log/Install.log`
 This directory contains the Access Appliance specific configuration logs.
 For example:
`/opt/VRTSsnas/log/Install.log.201407030655`

Access Appliance network installation and configuration logs `/opt/VRTSsnas/log/install_network.log`
 This directory contains the Access Appliance network configuration logs.
 For example:
`/opt/VRTSsnas/log/install_network.log.201407030655`

Installation fails and does not complete

Some common reasons for installation failures include:

- Gateway access
The Access Appliance node must be able to reach the default gateway using the public network. Verify with your network administrator that the gateway is reachable.

Troubleshooting Access Appliance 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 Access Appliance GUI startup issues

This chapter includes the following topics:

- [Resolving GUI startup issues](#)

Resolving GUI startup issues

Access Appliance 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 Access Appliance GUI after installing Access Appliance

- 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 Access Appliance GUI.

Example:

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

- 3 Log on to the Access Appliance GUI using the `admin` 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
```

Troubleshooting Veritas Data Deduplication issues

This chapter includes the following topics:

- [Log locations for the Veritas Data Deduplication server](#)

Log locations for the Veritas Data Deduplication server

If the VDD server goes into unhealthy state, you can check the HA log files for the reason of failure.

- `/log/VRTSnas/log/vacontainersha_monitor.log`
- `/log/VRTSnas/log/msdpwormha_monitor_<config_name>.log`

The log files for the operations that are performed on the Veritas Data Deduplication server are present in the following format:

`/log/VRTSnas/log/dedupe_<operation>.log`

Index

A

- a node
 - replacing 32
- about
 - common recovery procedures 22
 - event logs 10
 - job resynchronization 33
 - monitoring commands 20
 - services command 24
- Access Appliance log files
 - viewing 8
- an Ethernet interface card (online mode)
 - replacing 29

C

- CIFS
 - setting the log level 10
- cloud tiering
 - log locations 37
 - troubleshooting 36
- common recovery procedures
 - about 22
- configuring
 - job resynchronization 34

D

- debugging information
 - retrieving and sending 12
- debugging options
 - setting for NetBackup 11

E

- event logs
 - about 10

G

- general techniques
 - troubleshooting 7
- general tips
 - troubleshooting process 6

I

- installation
 - common failures 40
- installation logs
 - viewing 39

J

- job resynchronization
 - about 33
 - configuring 34

L

- log locations
 - cloud tiering 37

M

- monitoring
 - installation logs 39
- monitoring commands
 - about 20

N

- NetBackup client log levels
 - setting 11
- NetBackup debugging options
 - setting 11
- network
 - testing connectivity 26

P

- patch release
 - uninstalling 34

R

- reading or writing data from the cloud tier
 - troubleshooting 36
- recovering
 - from a non-graceful shutdown 26

- replacing
 - a node 32
 - an Ethernet interface card (online mode) 29
- replication
 - speeding up 33
- restarting
 - servers 23
- retrieving
 - debugging information 12

S

- sending
 - debugging information 12
- servers
 - restarting 23
- services command
 - about 24
 - using 25
- setting
 - CIFS log level 10
 - NetBackup client log levels 11
- shutdown
 - recovering from a non-graceful 26
- software upgrade
 - uninstalling 34

T

- testing
 - network connectivity 26
- traceroute
 - troubleshooting with 27
- traceroute command
 - using 28
- troubleshooting
 - about 6
 - general procedures 8
 - issues when reading or writing data from the
 - cloud tier 36
- troubleshooting process
 - general techniques 7
 - general tips 6

U

- uninstalling
 - patch release or software upgrade 34
- using
 - services command 25
 - traceroute command 28

V

- viewing
 - Access Appliance log files 8
 - installation logs 39