Veritas Access Release Notes

Linux 7.3.2

(Use with the Veritas Access 3340 Appliance documentation)

August 2018

Veritas Access Release Notes

Last updated: 2018-08-03

Document version: 7.3.2 Rev 1

Legal Notice

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.

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

http://www.veritas.com

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Overview of Veritas Access

This chapter includes the following topics:

About this release

Veritas Access 3340 Appliance is a software defined network attached storage (NAS) solution for unstructured data.

Veritas Access provides resiliency, multi-protocol access, and data movement to and from the public cloud based on policies.

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

Important release information

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

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

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

Changes in this release

This section shows the major new features and enhancements added in the 7.3.2 version of Veritas Access.

IP load balancing

IP load balancing has been introduced in this release. The purpose of this feature is to reduce the number of virtual IPs required for Veritas Access. A single virtual IP is used to act as a load balancer IP which distributes the incoming requests to the different nodes in the Veritas Access cluster for the services that are run on an active-active cluster.

See the *Veritas Access Administrator's Guide* for more details on this feature.

Veritas Access as an iSCSI target for RHEL 7.3 and 7.4

Veritas Access as an iSCSI target is introduced for RHEL 7. 3 and 7.4 in this release. This feature enables a Veritas Access cluster to serve block storage. Through the use of multiple portal IPs, an iSCSI target can be served in an active-active fashion.

This also enables the block storage to be capable of supporting multi-pathing at the initiator end. Veritas Access eases provisioning of block storage with the functionality to resize, clone, and snapshot the LUNs, ACL controls such as initiator mapping and user management.

Note: Veritas Access as an iSCSI target supports VMware version 5.5.0 as an initiator.

You can perform the following functions on an iSCSI target:

- Start , stop, and check status of an iSCSI target service
- Create, destroy, check status, and list iSCSI targets and add and delete multiple portal addresses
- Add, delete, resize, manage, grow, shrink LUNs and clone LUNs snapshots
- Map and remove mapping of iSCSI initiators
- Add and delete users to set up CHAP authentication
- Support for multiple portal IPs per target makes the targets active-active

See the *Veritas Access Administrator's Guide* for more details on this feature.

Changes to the GUI

The following updates are made to the GUI:

- For the Veritas Access 3340 Appliance, the **Getting Started** page is added with a workflow for configuring a Veritas Access 3340 Appliance:

  Note: The **Getting Started** page is displayed if you have not configured the Veritas Access 3340 Appliance options.

  - Configure storage
  - Configure an S3 server
  - Activate an LTR policy
  - Generate S3 keys
  - Provision storage
- For the Veritas Access 3340 Appliance, the **Hardware Monitoring** tab is added under NAS Infrastructure. It shows hardware related data for the primary and secondary storage.

Support for operating systems

The Veritas Veritas Access 7.3.2 release supports the VxOS 2.2.3 operating system.

Episodic and continuous replication in Veritas Access

Veritas Access episodic replication lets you asynchronously replicate a file system from one node in a source cluster to another node in a destination cluster at regularly timed intervals. This allows for content sharing, replication, and distribution.

Veritas Access continuous replication lets you replicate volumes from one node in the source cluster to another node in the destination cluster. The continuous replication enables you to maintain a consistent copy of application data at one remote location. It replicates the application writes on the volumes at the source location to a remote location across any distance.

See the *Veritas Access Administrator's Guide* for more details on this feature.

Active-active support for scale-out file system

A scale-out file system can be accessed through multiple nodes in the cluster concurrently with active-active support. This is an enhancement from the past releases where a scale-out file system was supported only in active-passive mode

Replication on a scale-out file system

A scale-out file system can be configured for synchronous or asynchronous replication. Synchronous replication provides zero RPO for applications. Asynchronous replication provides non-zero RPO while providing improved performance compared to synchronous replication.

Changes to the documentation set

The following changes have been made to the Access Appliance documentation set:

- The following documents have been removed from the Access Appliance documentation set:
  - *Veritas Access Getting Started Guide*
  - *Veritas Access Appliance Release Notes*
- *Veritas Access Appliance Getting Started Guide* is now renamed as *Access Appliance Initial Configuration and Administration Guide*.

Technical preview features

The following features are available as technical preview features in this release.

Veritas Access as an iSCSI Target for RHEL 6.x

Veritas Access as an iSCSI target is available as a technical preview feature for RHEL 6.x in this release.

The following functions are available in this feature:

- Veritas Access can be configured as an iSCSI target to serve the block storage.
- The iSCSI target service is hosted in active-passive mode in the Veritas Access cluster.
- After it is configured, the cluster is available to any standard iSCSI initiator over a portal IP.
- You can perform the following functions on an iSCSI target:
  - Start and stop the iSCSI target service
  - Add and delete iSCSI targets
  - Add and delete LUNs
  - Map and remove mapping of iSCSI initiators

    - Add and delete users

- See the `target` manual pages for more information.

The following limitations are present in this feature:

- Fault injection scenarios have not been covered during testing. Hence, iSCSI functionality may not behave as expected.
    - In case of a node reboot and cable pull scenarios, the feature may not behave as expected.
    - When the LUN add or destroy operation is performed concurrently with the target service restart, it may put the cluster in an inconsistent state.
- Strong integration with the rest of the Veritas Access code is incomplete.

  For example, integration with the network bonding and VLAN feature is incomplete.
- Performance testing has not been done.

Software limitations

This chapter includes the following topics:

Limitations related to installation and upgrade

The following limitations are related to installation and upgrade.

If required VIPs are not configured, then services like NFS, CIFS, and S3 do not function properly

If required number of VIPs are not configured during installation, then services like NFS, CIFS, and S3 do not function properly. High availability is also affected if VIPs are not configured correctly.

Add the required number of VIPs per service using the following CLISH command:

```
# network ip addr add <ipaddr> <netmask> <type (virtual)> [device]
[nodename]
```

Upgrade is not supported from CLISH

In this release, upgrade is only supported using the installer.

Rolling upgrade is not supported from CLISH

Rolling upgrade is only supported using the installer.

Limitations in the Backup mode

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

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

Veritas Access IPv6 limitations

IPv6 is not supported for this release.

FTP limitations

The following limitations apply to FTP.

- You have to manually create the user’s logon directory even if the `create_homedirs` option is set to `yes`.
- Multiprotocol access of FTP with other protocols such as NFS, CIFS is not supported.

Samba ACL performance-related issues

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

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

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

Veritas Access language support

Veritas Access supports only English.

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

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

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

NFS-Ganesha limitations

The following limitations apply for NFS-Ganesha:

- Clients cannot be added dynamically. Once an export is added, you cannot add more clients to the export. The workaround is to add a netgroup when you create the share. The netgroup membership can be changed dynamically.
- The `fcntl lock failover` is not supported for NFS-Ganesha v3.
- Export options like secure_locks, insecure_locks, wdelay, no_wdelay, subtree_check, no_subtree_check, and fsid are not supported with NFS-Ganesha.
- NFS-Ganesha supports only OpenStack Cinder. It does not support OpenStack Manila.
- NFS v4 ACLs are not supported by Veritas Access.

- NFS-Ganesha does not support share reservations.
- NFS-Ganesha does not support delegation.
- NFS server does not support non-ASCII characters.

Kernel-based NFS v4 limitations

The following limitations apply for kernel-based NFS v4:

- NFS v4 ACLs are not supported by Veritas Access.
- NFS v4 share reservations are not supported.
- NFS v4 delegation is not supported.
- The Veritas Access 3340 Appliance only supports using an on-premises LTR policy.

File system limitation

The following issue relates to the Veritas Access file system.

Any direct NLM operations from CLISH can lead to system instability

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

Veritas Access S3 server limitation

For downloading an object with a size more than 100M, `Range` header should be used and the range should not exceed 100M.

The object has to be downloaded in parts.

Long-term data retention limitations

The following limitations are related to long-term data retention (LTR).

- From NetBackup, if you select CloudCatalyst (CC) for the LTR use case, then you have to modify the Data Movement policy manually.
- In case a cluster node serving the OpenDedup volume crashes, the ongoing NetBackup jobs on that particular OpenDedup volume may fail. The same NetBackup job is successful in the next retry that is triggered automatically by

NetBackup. The NetBackup job may again restart when the crashed node comes up and the IP fails back.

- It is recommended on the Veritas Access 3340 Appliance to not use any policies except the LTR policies. All non-LTR policies use striped-mirror or mirror, which is not recommended on the Veritas Access 3340 Appliance because it already uses RAID-6 to protect the data. Using these policies will result in using twice the capacity to store the same amount of data. As a workaround you can manually create basic file systems in the UI (selecting file system instead of policy) or use the CLISH for full access to all options for all use cases including adding replication, configuring Scale-Out Filesystem, etc.

  LTR Policies use simple as the default file system type (can be seen in Settings > S3 Management > default parameters for S3 buckets > File System type.

Cloud tiering limitation

Cloud Tiering is a supported feature available on Scale-Out FS, it is not supported on CFS file systems. If you want to use cloud tiering you should manually create a Scale-Out file system instead of using one of the policies that are included with Access. The reason being all policies, except the LTR policies, create CFS file systems, which does not support cloud tiering.

Limitation related to replication

The following issues relate to replication in Veritas Access.

Limitation related to episodic replication authentication

When you create an episodic replication link, you have to provide the *admin* user credentials to authenticate a different cluster for episodic replication.

Limitation related to continuous replication

- Continuous replication does not support changing the mode of replication (synchronous or asynchronous) after replication is configured.
- Continuous replication does not accept file system with erasure coded (ecoded) layout and encrypted volume when you configure replication.
- The Veritas Access file system operations such as grow, shrink, resize, addition or removal of column, mirror, or tier (except cloud tier for *largefs*) are not supported for a file system which is configured under continuous replication.

Known issues

This chapter includes the following topics:

Veritas Access known issues

The following known issues relate to the Veritas Access commands.

Backup issues

This section describes known issues related to backup.

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

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

Workaround:

There is no workaround.

CIFS issues

This section describes known issues related to CIFS.

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

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

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

Workaround:

To solve this issue, follow these steps:

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

Deleting a CIFS share resets the default owner and group permissions for other CIFS shares on the same file system

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

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

Workaround:

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

```
CIFS> share modify
```

Default CIFS share has owner other than root

If a CIFS share (*share1*) is created using a non-default owner (*CIFSuser1* who is a non-root user) with file system (*fs1*) and if another share (*share2*) is created using the same file system (*fs1*) using default settings (root as the owner), then *share2* has a non-default owner (*CIFSuser1*).

Workaround:

If you want to export the same file system as different CIFS shares, then keep the owner of CIFS shares same for all shares. Otherwise, use different file systems to create different CIFS share.

Listing of CIFS shares created on a Veritas Access cluster fails on Windows server or client

If you try to list the all the CIFS shares from a Windows client machine using Veritas Access cluster IP (\\10.209.192.85,) the listing fails with an error message from Windows Explorer saying that network share is not accessible. This happens because Samba team has added new parameter `nt pipe support = no` to address vulnerability CVE-2017-7494.

Workaround:

There is no workaround for this issue.

CIFS> mapuser command fails to map all the users from Active Directory (AD) to all the NIS/LDAP users

While mapping all the CIFS users to NIS/LDAP users, the CLISH command does not accept the special character ‘*’.

Workaround:

Use one-to-one user mapping from Active Directory (AD) user to NIS/LDAP user.

Enterprise Vault Attach known issues

The following known issues relate to Enterprise Vault Attach:

Error while setting full access permission to Enterprise Vault user for archival directory

The Veritas Access GUI provides archival policies for storage provisioning for Enterprise Vault. As part of this storage provisioning, an empty folder named `ev_archival` is created in the CIFS share. This directory is used as the location of the Enterprise vault store partition. Enterprise Vault requires full access permission and ownership on the `ev_archival` folder to configure it as a vault store partition. After the creation of the `ev_archival` folder by the archival policy, you have to explicitly change the ownership of the folder before you give full access permission to this folder from Windows.

Workaround:

Perform the following steps to change ownership:

- Run the following command from the master node of the Veritas Access cluster.

  ```
  # chown "evsc\evuser" ev_archival
  ```

  Where *evsc* is the domain name and *evuser* is the Enterprise Vault user.

- From the Enterprise Vault server, access the CIFS shared network path which lists the `ev_archival` empty folder.
- Right click **ev_archival**.
- Go to **Security** tab and select **evuser**
- Click **Edit** and grant full permission.
- Click **Finish**.

GUI issues

The following issues relate to the GUI.

When both continuous and episodic replication links are set up, provisioning of storage using High Availability and Data Protection policies does not work

Performing the following steps leads to this scenario:

- Setting up both continuous and episodic replication links.
- Activating High Availability and Data Protection policies.
- Provisioning storage using either of these policies using the Provision Storage wizard.
- Setting up the episodic replication job task fails.

This happens because even though you selected the episodic replication link, during storage provisioning, the GUI selects the continuous replication link for setting up the episodic replication job, which causes the task to fail.

Workaround:

Do not create both continuous replication link and episodic replication link when you provision storage using the two policies above. Since these two policies are using the file system replication link, only create the episodic replication link.

When a new node is added or when a new cluster is installed and configured, the GUI may not start on the console node after a failover

When node failover occurs for a console node, the GUI services are expected to auto-start on the failed-over console node. But it fails to start as the GUI is not properly configured on all the nodes. You cannot use the GUI to manage the storage cluster.

Workaround:

When a failover occurs:

- Log on to the console node and run the following command:

  ```
  # python /opt/VRTSnas/isagui/init_application.py production
  ```

- Wait for the application to complete the configuration and display the message:

  ```
  Application started on Node JS
  ```

- Kill the application by entering CTRL-C.
- Enter the following command:

  ```
  # service vamgmt start
  ```

You can access the storage cluster using the GUI.

When an earlier version of the Veritas Access cluster is upgraded, the GUI shows stale and incomplete data

If you upgrade an old cluster and launch the GUI, you can see old events and incomplete data in the GUI pages.

Workaround:

After you upgrade the cluster, run the following command from the console node:

```
# /opt/VRTSnas/pysnas/bin/isaconfig
```

Restarting the server as part of the command to add and remove certificates gives an error on RHEL 7

When the external certificates are added to Veritas Access, a web server restart is implicitly performed to start the newly provided certificates. This implicit start of the web server does not work in RHEL7 because the commands are different in RHEL6 and RHEL7.

Workaround:

Run the `System> guienable` command to start the server in CLISH.

Client certificate validation using OpenSSL ocsp does not work on RHEL7

Client certificate validation is required for 2FA. The validation of certificates is successful in RHEL6. In RHEL7, an explicit parameter called *–VAfile* and the signer certificate are required to be passed, which does not happen. Hence, the client validation using the certificate does not work on RHEL7.

Workaround:

There is no workaround for this issue.

Installation and configuration issues

The following issues relate to Veritas Access installation and configuration.

Running individual Veritas Access scripts may return inconsistent return codes

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

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

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

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

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

Workaround:

To work around this issue:

1 Recover the SSH connection manually.

2 Uninstall Veritas Access.

3 Reinstall Veritas Access.

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

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

Workaround:

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

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

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

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

Workaround:

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

Phantomgroup for the VLAN device does not come online if you create another VLAN device from CLISH after cluster configuration is done

If you create a VLAN device on bond device during CPI installer configuration, and then try to create another VLAN device from CLISH after cluster configuration is done, the phantomgroup for the VLAN device does not come online successfully.

Workaround:

If the phantomgroup for the VLAN device is in *OFFLINE* or *FAULTED* state, enter the following commands:

```
# hagrp -clear <group-name>
# hagrp -online <group-name> -any
# hagrp -state <group-name>
```

The state of phantomgroup becomes *ONLINE*.

Veritas Access installation fails if the nodes have older yum repositories and do not have Internet connectivity to reach RHN repositories

If you try to install Veritas Accessand the yum repositories present in the nodes are outdated, then the installer tries to reach the RHN repositories to update the yum repository. If you do not have Internet connectivity, then the installation fails.

Workaround:

Remove the yum configuration file for the yum repository which is present in `/etc/yum.repos.d`. Then, run the `yum clean all` command to refresh the yum repository information. Re-run the Veritas Access installer.

Networking issues

This section describes known issues related to networking.

CVM service group goes into faulted state unexpectedly

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

Workaround:

Reboot the node on which this issue occurs.

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

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

Workaround:

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

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

After network interface swapping between two private NICs or one private NIC and one public NIC, the service groups on the slave nodes are not probed

For performing a network interface swapping between two private NICs or one private NIC and one public NIC, only one node should be present in the cluster. If more than one node is present, the remaining nodes are not probed after the network interface swapping.

Workaround:

Execute the following command on all nodes where resources are not probed:

```
# hastart
```

NFS issues

This section describes NFS issues.

Slow performance with Solaris 10 clients with NFS-Ganesha version 4

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

Workaround:

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

Random-write performance drop of NFS-Ganesha with Linux clients

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

Workaround:

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

Latest directory content of server is not visible to the client if time is not synchronized across the nodes

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

Workaround:

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

NFS> share show may list the shares as faulted for some time if you restart the cluster node

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

Workaround:

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

NFS-Ganesha shares faults after the NFS configuration is imported

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

Workaround:

Restart the NFS service.

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

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

Workaround:

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

Exporting a single path to multiple clients through multiple exports does not work with NFS-Ganesha

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

Workaround:

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

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

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

Workaround:

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

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

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

Workaround:

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

NFS> share show command does not distinguish offline versus online shares

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

Workaround

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

Difference in output between NFS> share show and Linux showmount commands

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

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

Workaround:

You can verify the given IP addresses by using DNS.

NFS mount on client is stalled after you switch the NFS server

When the NFS server is switched from kernel NFS to NFS-Ganesha (or vice versa), the existing NFS mounts on the client are no longer active. This is because after the server is switched, all the exports on the server are moved to the new server and the file handling method of the kernel NFS and NFS-Ganesha servers are different. Hence, the NFS mount on the client is stalled.

Workaround:

The client can remount the exports to access the shares.

Kernel NFS v4 lock failover does not happen correctly in case of a node crash

With kernel NFS v4 shares, in case of a node crash, active locks do not failover to another node in the cluster.

Workaround:

There is no workaround for this issue.

Kernel NFS v4 export mount for Netgroup does not work correctly

The Netgroup membership cannot be changed dynamically with kernel NFS v4. Hence, the kernel KNFS v4 export mount for Netgroup does not work as expected.

Workaround:

Restart the NFS service.

ObjectAccess issues

This section describes ObjectAccess issues.

ObjectAccess server goes in to faulted state while doing multi-part upload of a 10-GB file with a chunk size of 5 MB

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

Workaround:

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

When trying to connect to the S3 server over SSLS3, the client application may give a warning like "SSL3_GET_SERVER_CERTIFICATE:certificate verify failed"

Veritas Access generates a self-signed SSL certificate. This certificate is not a part of the default trusted CAs. Hence, S3 client is not able to trust it.

Workaround:

Client should ignore the warning and continue the communication over SSL.

If the cluster name does not follow the DNS hostname restrictions, you cannot work with the ObjectAccess service in Veritas Access

A cluster name cannot contain any special symbols except for a hyphen. If the cluster name has special symbols other than the hyphen, then the S3 service does not work as the DNS hostname restrictions have not been followed.

Workaround:

There is no workaround for this issue. For valid characters for naming a Veritas Access cluster, see:

https://technet.microsoft.com/en-us/library/cc959336.aspx

ObjectAccess operations do not work correctly in virtual hosted-style addressing when SSL is enabled

When SSL is enabled, ObjectAccess operations do not work correctly in virtual hosted-style addressing

Workaround:

Use path-style access when SSL is enabled.

Bucket creation may fail with time-out error

If bucket creation takes a long time, then the bucket creation request may fail with an error message even if the bucket got created successfully.

Workaround:

You can verify if the bucket exists, even if the request fails.

Bucket deletion may fail with "No such bucket" or "No such key" error

If a client request retry happens before the completion of the previous request for bucket deletion is completed, then the subsequent retry may get stale information. The bucket deletion request fails with an error message.

Workaround:

Client needs to verify bucket deletion even if the request fails.

Temporary objects may be present in the bucket in case of multi-part upload

If an object gets uploaded to the bucket using multi-part upload, then multiple temporary objects may be present in the bucket. Temporary objects have an internal naming convention and end with a sequential number.

Workaround:

Temporary objects get removed once all the parts are uploaded and reassembly is complete.

Group configuration does not work in ObjectAccess if the group name contains a space

If the group name has a space, then even if the configuration is set for that group, user of that group is unable to create a bucket with that configuration. Instead, the bucket is created with the default configuration.

The administrator should not configure ObjectAccess for a group having a space character in its name.

OpenDedup issues

This section describes known issues related to OpenDedup.

The file system storage is not reclaimed after deletion of an OpenDedup volume

When an OpenDedup volume is deleted using CLISH commands, the content of the OpenDedup data in the bucket is not removed. Hence, space is not reclaimed for the corresponding file system.

Workaround:

Delete the OpenDedup content manually using any S3 client.

Removing or modifying the virtual IP associated to an OpenDedup volume leads to the OpenDedup volume going into an inconsistent state

When an OpenDedup volume is created, a virtual IP address is assigned to the volume. If the virtual IP is modified or removed, then the associated volume cannot be used.

Workaround:

There is no workaround for this issue.

OpenDedup port is blocked if the firewall is disabled and then enabled again

When an OpenDedup volume is created, it is mounted on a specific port. If the firewall is disabled and enabled again using the CLISH commands, the firewall does not add the rule for the port number used by the OpenDedup volume. Hence, the OpenDedup port is blocked.

Workaround:

Offline the OpenDedup volume and make it online again.

The Storage> fs online command fails with an EBUSY error

If a bucket or an OpenDedup volume is present on the scale-out file system and I/O operations are running, the `Storage> fs offline` command is successful, but the `Storage> fs online` command may fail or the S3 server may not work as expected.

Workaround:

Before performing the `Storage> fs offline` command, verify that the ObjectAccess bucket or the OpenDedup volume is present on that file system and no I/O operations are running.

The OpenDedup volume is not mounted automatically by the /etc/fstab on the media server after a restart operation

If OpenDedup is running on the NetBackup media server and it is restarted, the OpenDedup volume may not be mounted even if it has a `/etc/fstab` entry due to a timing issue.

Workaround:

Use other alternatives such as the `/etc/rc.local` script to ensure that the OpenDedup volume is mounted once the media server comes up.

Output mismatch in the df -h command for OpenDedup volumes that are backed by a single bucket and mounted on two different media servers

If two OpenDedup volumes that are backed by a single bucket are mounted on two different media servers, the `df -h` command shows different output for the volumes.

Workaround:

Ensure that the serial number entry in the OpenDedup volume XMLs is the same on both media servers.

OpenStack issues

The following issues are related to OpenStack.

Cinder and Manila shares cannot be distinguished from the CLISH

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

Workaround:

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

The Veritas Access Manila driver which is upstreamed in the OpenStack Manila repository is not compatible with Veritas Access 7.3.1

The Veritas Access Manila driver which is upstreamed in OpenStack Manila Repository is not compatible with Veritas Access 7.3.1 due to modification in Veritas Access APIs.

Workaround:

Download the Veritas Access Manila driver from the Veritas VxBeta site. Use the downloaded driver with Veritas Access 7.3.1 version.

Replication issues

This section describes known issues related to both episodic and continuous replication.

Note: Episodic replication is not supported for a scale-out file system.

When running episodic replication and dedup over the same source, the episodic replication file system fails in certain scenarios

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

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

Workaround:

There is no workaround.

The System> config import command does not import episodic replication keys and jobs

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

Workaround:

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

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

The job uses the schedule on the target after episodic replication failover

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

Workaround:

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

Episodic replication fails with error “connection reset by peer” if the target node fails over

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

Workaround:

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

Episodic replication jobs created in Veritas Access 7.2.1.1 or earlier versions are not recognized after an upgrade

If you try to access or modify the episodic replication jobs that were created in Veritas Access 7.2.1.1 or earlier releases, the commands do not work since the jobs are in an unrecognized state.

Workaround:

Destroy the job and create it again.

Setting the bandwidth through the GUI is not enabled for episodic replication

The `bwlimit show` does not show the expected output in CLISH.

```
Replication> episodic bwlimit show
ERROR V-288-0 No job is configured with current node as replication source
```

Hence, the `bwlimit show` is not supported through the GUI.

Workaround:

You can use the following command to set the bandwidth:

```
 Replication> episodic bwlimit set src_to_tgt 10
```

Episodic replication job with encryption fails after job remove and add link with SSL certificate error

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

`SSL certificate error`.

Workaround:

Follow these steps to solve this issue:

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

Episodic replication job status shows the entry for a link that was removed

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

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

Workaround:

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

Episodic replication job modification fails

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

Workaround:

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

Episodic replication failover does not work

If you try to make the target cluster as the new source cluster when the source cluster has failed, it does not work. Hence, failover of the episodic replication cluster is not successful.

Workaround:

There is no workaround for this issue.

Continuous replication fails when the 'had' daemon is restarted on the target manually

If the 'had' daemon is stopped and restarted on the target, continuous replication fails. This happens because the IP tables rules are not restored for continuous replication.

Workaround:

- On target, set the following rule.

```
# iptables -I INPUT 2 -p tcp -d <replication_ip of target>
--dport 56987 -j ACCEPT
```

- Save the rule.

```
# service iptables save
```

- Restart the IP tables.

```
# service iptables restart
```

Continuous replication is unable to come in replicating state if the Storage Replicated Log becomes full

While replicating data from the source cluster to the target cluster, if the Storage Replicated Log (SRL) becomes full, It goes into Data Change Map (DCM) mode. In DCM mode, it does not show the status as *replicating*.

```
Replication> continuous status test_fs
Name                   value
=====================  ==============================================
Replicated Data Set    rvg_test_fs
Replication Role       Primary
Replication link       link1

Primary Site Info:

Host name           10.10.2.70
RVG state           enabled for I/O

Secondary Site Info:

Host name           10.10.2.72
Configured mode     synchronous-override
Data status         inconsistent
Replication status  resync in progress (dcm resynchronization)
Current mode        asynchronous
Logging to          DCM (contains 551200 Kbytes) (SRL protection logging)
```

Workaround:

Run the following command on the source cluster for continuous data replication.

```
# vxrvg -g <dg_name> resync <rvg_name>
```

The command resynchronizes the source and the target cluster. You can check the status by entering the following command:

```
Replication> continuous status test_fs
Name                    value
=====================   =======================
Replicated Data Set     rvg_test_fs
Replication Role        Primary
Replication link        link1

Primary Site Info:

Host name               10.10.2.70
RVG state               enabled for I/O

Secondary Site Info:

Host name               10.10.2.72
Configured mode         synchronous-override
Data status             consistent, up-to-date
Replication status      replicating (connected)
Current mode            synchronous
Logging to              SRL
Timestamp Information   behind by  0h 0m 0s
```

Unplanned failover and failback in continuous replication may fail if the communication of the IPTABLE rules between the cluster nodes does not happen correctly

In case of unplanned failover and failback, the IPTABLE rules may not get restored properly. Hence, the communication between the nodes does not happen correctly.

Workaround:

Flush the IPTABLES on all the nodes in the cluster on the primary as well as the secondary site.

```
# iptables -F
```

Continuous replication configuration may fail if the continuous replication IP is not online on the master node but is online on another node

At the target site, there may be a situation wherein the management console is not online on the node on which continuous replication IP is online. In that case, the configuration of continuous replication may fail since internal commands need to run on the master node.

Workaround:

Make sure that you can access CLISH through the master node and the continuous replication IP is also online on the master node. If not, then use the following command to switch the management console position to the master node.

```
# hagrp -switch ManagementConsole -to <system_name>
```

If you restart any node in the primary or the secondary cluster, replication may go into a PAUSED state

When you restart any node in the primary or the secondary cluster, the communication of the IPTABLE rules between the cluster nodes does not happen correctly. This results in replication going into PAUSED state.

Workaround:

Flush the IPTABLES on all the nodes in the cluster on the primary as well as the secondary site.

```
# iptables -F
```

SmartIO issues

The following issue relates to the Veritas Access SmartIO commands.

SmartIO writeback cachemode for a file system changes to read mode after taking the file system offline and then online

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

Workaround:

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

Storage issues

The following issues relate to the Veritas Access Storage commands.

Snapshot mount can fail if the snapshot quota is set

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

Workaround:

Remove the oldest checkpoint and retry.

Sometimes the Storage> pool rmdisk command does not print a message

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

Workaround:

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

The Storage> pool rmdisk command sometimes can give an error where the file system name is not printed

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

Workaround:

There is no workaround.

Not able to enable quota for file system that is newly added in the list of CIFS home directories

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

Workaround:

Run the following commands from CLISH:

```
Storage> quota cifshomedir disable
```

```
Storage> quota cifshomedir enable
```

Destroying the file system may not remove the /etc/mtab entry for the mount point

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

Workaround:

Reboot the cluster nodes.

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

The `Storage> fs online` command returns the following error:

```
access.Storage> fs online fs1
```

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

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

Workaround:

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

Removing disks from the pool fails if a DCO exists

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

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

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

Workaround:

There is no workaround.

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

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

This situation can happen in one of the following cases:

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

Workaround:

Grow the file system.

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

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

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

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

Workaround:

There is no workaround.

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

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

Workaround:

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

Inconsistent cluster state with management service down when disabling I/O fencing

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

Workaround:

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

Storage> tier move command failover of node is not working

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

Workaround:

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

Storage> scanbus operation hangs at the time of I/O fencing operation

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

Workaround:

There is no workaround. Contact Veritas Technical Support.

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

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

Workaround:

There is no workaround.

Event messages are not generated when cache objects get full

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

Workaround:

There is no workaround.

Veritas Access CLISH interface should not allow uncompress and compress operations to run on the same file at the same time

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

Workaround:

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

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

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

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

Workaround:

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

Storage> tier move list command fails if one of the cluster nodes is rebooted

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

Workaround:

There is no workaround.

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

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

Workaround:

There is no workaround.

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

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

Workaround:

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

Storage> fs addcolumn operation fails but error notification is not sent

`Storage> fs addcolumn` operation fails in the background but the notification of the failure is not sent as the error message is not present in CLISH. One of the reasons for the failure is not having enough storage in the given pool.

Workaround:

If required number of columns are not added, try again after adding enough storage.

Storage> fs-growto and Storage> fs-growby commands give error with isolated disks

The `Storage> fs growto` and `Storage> fs growby` commands give a *Not enough space* error even though there is enough space. The operations fail in the following scenarios:

1. The file system is created on normal pool(s). But disks from isolated pools are given for `fs growto` and `fs growby` operations.

2. The file system is created on an isolated pool but disks from normal pool(s) or different isolated pool(s) are given for `fs growto` and `fs growby` operations.

Workaround:

If the file system is created on normal pool(s), then provide disks from normal pool(s) for `fs-growto` and `fs-growby` operations. If the file system is created on an isolated pool, then add disk(s) to the same isolated pool and provide them for `fs-growto` and `fs-growby` operations.

Unable to create space-optimized rollback when tiering is present

In a tiered file system, creation of space-optimized rollbacks fails. The failure occurs when the primary tier has `fastresync` enabled while the secondary tier does not have `fastresync` enabled.

The secondary tier has `fastresync` disabled in the following scenarios:

1. The tier is mirrored but `fastresync` is manually disabled.

2. The tier is simple or striped in which case `fastresync` cannot be enabled.

Workaround:

If the secondary tier is mirrored, enable `fastresync` on it.

If the secondary tier is simple (or striped) and primary tier is mirrored, add a mirror to the secondary tier.

Ensure that the secondary tier has `fastresync` enabled if the primary tier also has `fastresync` enabled.

Enabling I/O fencing on a set up with Volume Manager objects present fails to import the disk group

If you enable I/O fencing on a set up with Volume Manager objects present, it fails to import the disk group and you get the following error message:

```
Disk <diskname> does not support SCSI-3 PR, Skipping PGR operations
for this disk
```

If there are Volume Manager objects like volumes, and volume sets, and you enable I/O fencing, then the shared disk group is not imported as a part of the cluster join.

Even manual import of the disk group using the `vxdg -s import <dgname>` command fails with the following error message:

```
SCSI-3 PR operation failed
```

This issue is due to the export flag that is missing on the disk which has been implicitly exported using the disk map command. This happens if the disk group contains disks that do not support SCSI3 PR.

Workaround:

Explicitly export all the DAS disks from all the nodes of the cluster using the following commands before you enable majority-based fencing.

```
# vxdisk -f export <DAS disk Name>
```

You can now enable I/O fencing.

File system creation fails when the pool contains only one disk

When there is only one disk in pool, the `fs creation` command fails to create an NLM on the file system. Instead, it tries to create the file system with different options.

Workaround:

Ensure that there is more than one disk in the pool.

After starting the backup service, BackupGrp goes into FAULTED state on some nodes

BackupGrp is online on only one node. When the backup service is started, it probes the group on all the cluster nodes and tries to become online on multiple nodes. But, as this is a failover group it cannot be online on more than one node. Hence, it goes into FAULTED state on some nodes.

Workaround:

Clear the fault using the following command:

```
BacupGrp> hagrp -clear BackupGrp
```

A scale-out file system created with a simple layout using thin LUNs may show layered layout in the Storage> fs list command

If you use thin LUNs, FMR is enabled by default. DCO volumes are created when the FMR feature is enabled. When DCO volumes are present on the system, the `Storage> fs list` command incorrectly derives the layout of the scale-out file system. The command either shows incorrect volume layout or if the layout is correct, the number of mirrors are shown incorrectly. This is an issue with the display of the output, the scale-out file system has the correct layout.

Workaround:

Use the `Storage> fs list fs_name` command for finding detailed information about the file system.

A file system created with a largefs-striped or largefs-mirrored-stripe layout may show incorrect number of columns in the Storage> fs list command

If you create a file system with a largefs-striped or largefs-mirrored-stripe layout, the`Storage> fs list` command incorrectly derives the details of the layout of the file system. The command either shows the number of columns incorrectly. This is an issue with the display of the output.

Workaround:

There is no workaround.

A scale-out file system may go into faulted state after the execution of Storage> fencing off/on command

When the `Storage> fencing on` and `Storage> fencing off` operations are executed, the VCS resource of the respective scale-out file system goes into faulted state.

Workaround:

Use the `Support> service autofix` command to fix the file systems that are in faulted state. If the service groups do not become online, then restart the node on which the scale-out file system's VCS service groups are in faulted state. Use the following command to check the status of the VCS service groups for the cluster nodes.

```
Support> services showall
```

The CVM service group goes in to faulted state after you restart the management console node

When the `Cluster> reboot` command is run, sometimes the CVM service group goes into faulted state on the node that was restarted. This issue is usually caused by a minor number conflict between the CVM shared disk group objects, such as volumes, volume sets or Replicated Volume Groups (RVGs) and the private disk group objects. Confirm that the minor numbers of the private disk group objects do not overlap with the CVM disk group objects on the joining CVM slave node.

https://www.veritas.com/support/en_US/article.000107801

Workaround:

To bring the CVM service group online

1. Run the following command on the node where CVM service group is in faulted state

   ```
   # hastop -local
   ```

2. Offline all the file systems. Run the following command from another node where the management console is online.

   ```
   Storage> fs offline <file system name>
   ```

3. Deport all the disk groups using the following command:

   ```
   # vxdg -s deport <disk_group>
   ```

4 Import all the disk groups using the following command:

```
# vxdg -s import <disk_group>
```

5 Start VCS.

```
# hastart
```

If the file system does not come online, then run the following command to make all the file systems online:

```
Storage> fs online <file system name>
```

After an Azure tier is added to a scale-out file system, you cannot move files to the Azure tier and the Storage> tier stats command may fail

After you add an Azure tier is to a scale-out file system, you cannot move files to the Azure tier and the `Storage> tier stats` command may fail with the following error:

```
ACCESS tier ERROR V-493-10-2059 Failed to display access statistics of cloud
tier aztierx2 (errnum=22).
```

Workaround:

Offline and online the scale-out file system that has the Azure tier using the `Storage> fs offline fs_name` and `Storage> fs online fs_name` commands. Or, you can kill the `tfsd` daemon on all the nodes of the cluster.

System issues

The following issues relate to the Veritas Access System commands.

The System> ntp sync command without any argument does not appear to work correctly

The `System> ntp sync` command without any argument does not work as per expectations. It gives a message that the date is synchronized on all the node even if the date is not synchronized.

Workaround:

The `System> ntp sync` command should be executed with an NTP server as an explicit argument for performing a sync operation on all the nodes.

Target issues

This section describes known issues related to Target.

If a user is added on the target side, the initiator cannot see the LUNs

While discovering targets, the Veritas Access initiator does not ask for credentials, and takes the default credentials. Hence, the initiator is unable to log into the targets. As a result, the target status on the initiator becomes OFFLINE and LUNs are not visible to the initiator.

Workaround:

Log in from outside of the CLISH using the following `iscsiadm` commands. Execute this command on both nodes so that the disks (LUNS) become FSS disks.

```
# iscsiadm -m node -T <target-name> -p <portal-ip>
--op=update --name node.session.auth.authmethod --value=CHAP
iscsiadm -m node -T <target-name> -p <portal-ip>
--op=update --name node.session.auth.username --value=username

iscsiadm -m node -T <target-name> -p <portal-ip>
--op=update --name node.session.auth.password --value=password
```

LIO does not support target name in uppercase

When a target name is created with uppercase letters in Veritas Access, LIO (Linux I/O) converts it into lower case. This causes inconsistency between LIO and Veritas Access

Workaround:

Create target names with lowercase letters and use the conventions that are specified in RFC-3720 document.

Access Appliance issues

The following issues are related to the Veritas Access 3340 Appliance.

Mongo service does not start after a new node is added successfully

After you add a new node, the installer does not bring the Cluster Volume Manager (CVM) and its dependent service groups (appdb_data and appdb_svc) online on the newly added node. Hence, the Mongo service does not start on the newly added node.

Workaround:

After you add a node, log on to the Veritas Access CLISH interface using the *admin* user credentials. In the CLISH interface, execute the `cluster reboot <new_node>` command to restart the newly added node.

File systems that are already created cannot be mapped as S3 buckets for local users using the GUI

If you create a file system using the GUI and try to export it as an S3 bucket for a local user, the operation is not successful.

Workaround:

Map the file system using the Veritas Access CLISH for local users.

The Veritas Access management console is not available after a node is deleted and the remaining node is restarted

If a node is deleted and the only remaining node in the cluster is restarted, the management console's IP gets cleared up. Hence, the service group of the management console goes into a faulted state and then management console becomes unavailable.

Workaround:

Perform the following steps:

- Log on using the Appliance CLISH as an *admin* user.
- Go to **Support** view.
- Go to **Management** view.
- Elevate to access the prompt. Run the following commands:

```
# /opt/VRTS/bin/hares -clear consoleIP -sys <current node>
```

```
# /opt/VRTs/bin/hagrp -online ManagementConsole -any
```

When provisioning the Veritas Access GUI, the option to generate S3 keys is not available after the LTR policy is activated

While provisioning the Veritas Access GUI, after you click **Activate LTR policy**, the **Next** option is not enabled. Hence, you cannot proceed to **Generate S3 Keys**.

Workaround:

Click on the **Close** option. Click on **Continue** to resume the Getting Started wizard.

Unable to add an Appliance node to the cluster again after the Appliance node is turned off and removed from the Veritas Access cluster

If an Appliance node is turned off and removed from the Veritas Access cluster successfully then, you cannot to add the Appliance node to the Veritas Access cluster again. This happens because the configurations on the Appliance node are not uninstalled when the node is removed.

Workaround:

Before you add the Appliance node back to the Veritas Access cluster, do a factory reset to the Appliance node.

Setting retention on a directory path does not work from Veritas Access CLISH

If you set retention on any directory path using the `Storage> fs retention set` *`<absolute directory path>`* command from Veritas Access CLISH, the command fails.

For example :

```
Storage> fs retention set /vx/test_fs/
```

Where */vx/test_fs/* is the directory path.

This is because of an internal issue. But, it is possible to set retention individually on the files present in a particular directory.

Workaround:

Set retention on the files present in the directory by using the `Storage> fs retention set` *`<absolute file path>`* command.

```
Storage> fs retention set /vx/test_fs/file1
```

Where */vx/test_fs/file1* is the file path.

Access Appliance operational notes

This section contains the topics that explain important aspects of Veritas Access Appliance 7.3.2 operations that may not be documented elsewhere in the documentation.

The following list contains the notes and the known issues that apply to this Access Appliance release:

- If you have opened the following shares on the appliance node, you may experience a 'time-out' or 'denied' error when mounting the NFS share on your client computer:

  ```
  Main > Manage > Software > Share
  Main > Support > Logs > Share
  ```

  The NFS server for the Access cluster may be offline. After the Access cluster configuration is completed, the NFS server is offline by default. If you want to access the NFS shares on the appliance nodes, you must ensure that the NFS server is online from the Access shell menu.
  To check NFS status and start the service, log on to the Access shell menu and run the following commands:

  ```
  ltra> nfs server status
  NFS Status on ltra_01 : OFFLINE
  NFS Status on ltra_02 : OFFLINE
  ltra> nfs server start
  Success.
  ltra> nfs server status
  NFS Status on ltra_01 : ONLINE
  NFS Status on ltra_02 : ONLINE
  ```

- **Storage Connection** appears twice under **Infrastructure > Hardware** in the Access GUI. The second instance is actually the **Temperature** menu.
- The Ethernet interface names are not automatically reverted to default value, for example, eth4, eth5, after an appliance node is removed from cluster. Instead, the Access cluster Ethernet interface names remain, for example, priveth0, pubeth1.
  To refresh the Ethernet interface status, you must restart the appliance after you remove the appliance node. Then, use the `Network > Show Status` command to check the latest Ethernet interface status.
- When an appliance node leaves the Access cluster, the logon credential remains unchanged. The logon credential does not revert to the state before the node was added to the Access cluster. This behavior is by design.
  To log on the node after the node has been removed from the cluster, use the same credential as what was used when the node was in the cluster.
- In the following scenarios, the `Manage > Cluster > Configure` command is available, but running the command results in a failure.

- While the cluster configuration process is ongoing on one node, you run the `Cluster > Configure` command on another un-configured node of the Access Appliance.

  In this scenario, the `Cluster > Configure` command fails. However, you can safely ignore this failure because the existing cluster configuration process will not be affected.

- One node in a configured Access Appliance is factory reset without resetting the shared storage (the cluster is intact), and then you run the `Cluster > Configure` command on the factory state node in the appliance.

  In this scenario, the `Cluster > Configure` command fails. However, you can safely ignore this failure because it does not affect the node that remains in the Access cluster (single-node cluster).

- If SSL is enabled for S3 and then you run the `Support > Test Software` command, the S3 Service test always shows `[WARNING]`. The following warning message appears incorrectly:

  - [Warning] V-409-830-5451: The S3 service is not ready. Refer to the Access Appliance documentation for how to configure the S3 service.

  To work around this issue, you can temporarily disable SSL for S3, and then run the test.

- After an Access Appliance is configured, and S3 is ready, if you run the `Support > Test Software` command on one node while a software self-test is ongoing on another node, the following occur:

  - The S3 Service test results on both nodes always incorrectly show `[WARNING]`, and a `[WARNING]` message as follows:

    - [Warning] V-409-830-5451: The S3 service is not ready. Refer to the Access Appliance documentation for how to configure the S3 service.

  - At the same time, a '`selftest`' bucket is created on the shared storage. You cannot delete the bucket from the Access GUI. The '`selftest`' bucket size is 5 GB. If you have changed the default bucket size before this issue occurred, this '`selftest`' bucket size defaults to that size.

  This issue does not affect the S3 service working. You cannot delete the '`selftest`' bucket unless you do a full factory reset and re-configure the Access Appliance.

- The file system type for S3 buckets must be set to **Scale-Out FS** from Access GUI or **largefs-simple** from Access shell menu. Otherwise, the bucket pool size can only be resized to 512 TB at maximum.

  - When you create file system from the Access GUI, the default type is **Scale-Out FS**. Make sure that you use the default type.

- When you create file system from the Access shell menu, make sure that you specify the file system type as **largefs-simple**.

Access services do not restart properly after storage shelf restart

If the Veritas Access Appliance loses connectivity to an attached Primary or Expansion storage shelf, the underlying storage connectivity is lost and the VxVM disk group goes into a deported state. This issue occurs whenever a storage shelf intentionally or unintentionally restarts. To correct this issue, you need to restart the Access services.

To restart the Access services after the appliance storage shelves restart

1. Log onto the Access shell menu over the console IP address.
2. Run the following command to import the VxVM disk group and other Access configurations:

   ```
   ltrcluster> storage scanbus
   ```

3. Restart the services that were configured before the storage shelf restart.

   For example, if the S3 server is configured, use the following commands

   ```
   ltrcluster> objectaccess server status
   ObjectAccess Status on ltrcluster_01 : OFFLINE|FAULTED
   ObjectAccess Status on ltrcluster_02 : OFFLINE|FAULTED
   ltrcluster> objectaccess server stop
   ACCESS ObjectAccess ERROR V-493-10-4 ObjectAccess server already stopped.
   ltrcluster> objectaccess server start
   ACCESS ObjectAccess SUCCESS V-493-10-4 ObjectAccess started successfully.
   ltrcluster> objectaccess server status
   ObjectAccess Status on ltrcluster_01 : ONLINE
   ObjectAccess Status on ltrcluster_02 : ONLINE
   ltrcluster>
   ```

Chapter 4

Getting help

This chapter includes the following topics:

Displaying the Online Help

You can access the Online Help through the management console of Veritas Access by clicking the question mark icon.

Displaying the man pages

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

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

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

Using the Veritas Access product documentation

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

https://sort.veritas.com/documents

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

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

The following guides are available for the Access Appliance:

- *Veritas Access Appliance Initial Configuration and Administration Guide*

  Note: This guide replaces the *Veritas Access Appliance Getting Started Guide*

- *Veritas Access Appliance Commands Reference Guide*
- *Veritas Access 3340 Appliance Product Description*
- *Veritas Access 3340 Appliance Hardware Installation Guide*
- *Veritas Appliance AutoSupport 2.0 Reference Guide*
- *Veritas Access Safety and Maintenance Guide*

The following Veritas Access 7.3.1 documents are available on the SORT site:

- *Veritas Access Administrator's Guide*
- *Veritas Access Command Reference Guide*
- *Veritas Access NetBackup Solutions Guide*
- *Veritas Access Release Notes*
- *Veritas Access RESTful API Guide*
- *Veritas Access Third-Party License Agreements*
- *Veritas Access Troubleshooting Guide*
- *Veritas Access Enterprise Vault Solutions Guide*