Veritas Access Release Notes

Linux

7.4

VERITAS

Veritas Access Release Notes

Last updated: 2018-07-24

Document version: 7.4 Rev 1

Legal 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

Chapter 1

Overview of Veritas Access

This chapter includes the following topics:

About this release

Veritas Access is a software-defined, scale-out network-attached storage (NAS) solution for unstructured data that works on commodity hardware. Veritas Access provides resiliency, multi-protocol access, and data movement to and from the public cloud based on policies.

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

Important release information

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

The hardware compatibility list contains information about supported hardware and is updated regularly. You can use any commodity hardware that is certified and mentioned in the hardware compatibility list.

For the latest information on supported hardware, see the compatibility list at:

https://sort.veritas.com/documents/doc_details/isa/7.4/Linux/CompatibilityLists/

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

https://www.veritas.com/support/en_US/article.100042732

Changes in this release

This section shows the major new features and enhancements added in the 7.4 version of Veritas Access.

Integration with the SDS Management Platform

The Software-Defined Storage (SDS) Management Platform when integrated with Veritas Access provides a platform to integrate data from various sources within an enterprise (not just the data center), which presents a single, integrated view in to the whole environment. The SDS Management Platform brings together Veritas Access and Veritas NetBackup using a simple and intuitive platform to address long-term retention and other use cases.

See the *Veritas Access Software-Defined Storage (SDS) Management Platform Solutions Guide* for more information.

Capacity-based licensing model

In this release, Veritas has introduced the TB-per-core licensing model for Veritas Access. The per-core and per-terabyte licensing model of earlier releases is also supported in this release.

The TB-per-core licensing model is based on both capacity per-core and time period. You can now license Veritas Access as per your requirement for raw capacity. This is managed through the software.

See the *Veritas Access Installation Guide* for more details on licensing.

Support for AWS signature version 2 or version 4 authentication

Veritas Access supports either AWS signature version 2 or version 4 authentication.

Changes to the GUI

The following updates are made to the GUI:

- On the **Service Management** page, the **iSCSI Target Service Management** option is added. You can use this option to set the iSCSI server online or offline.
- On the Dashboard page, under **Quick Actions** > **Provision Storage**, the following options are added:
  - CIFS
    - With continuous replication
    - With episodic replication
    - With encryption
    - With replication and encryption
    - Without replication and encryption
  - Storage for Enterprise Vault
    - With replication
    - Without replication
  - NFS
    - With continuous replication
    - With episodic replication
    - With encryption
    - With replication and encryption
    - Without replication and encryption
  - iSCSI Block storage
    - Provision iSCSI LUNs
  - S3 storage for NetBackup
    - With a cloud tier
    - Without a cloud tier

Support for RHEL and OL operating systems

The Veritas Access 7.4 release supports the following operating systems:

- Red Hat Enterprise Linux (RHEL)
  - RHEL 7 Update 3 and 4
- Oracle Linux (only in RHEL compatible mode)

- OL 7 Update 3 and 4

Support for internationalization (I18N)

Veritas Access provides Chinese, Japanese and Korean (CJK) character support for I18N in this release. Veritas Access accepts inputs in CJK characters. The output and other header messages are in English.

The support includes:

- CJK character support is present for the **Storage**, **ObjectAccess**, **NFS**, **CIFS**, **FTP**, and **SmartIO** modules.
- Only the object names are accepted in CJK characters.
- The object that is created using CJK characters is displayed in the same way in the `list`, `show`, and `status` commands.
- You can perform all the operations on the object that are created using CJK characters in the same way as for the remaining I18N ported commands.
- CJK character support is available from the CLISH and GUI.
- Testing has not been done for the other languages that are supported by RHEL.

The support does not include:

- The modules apart from the ones mentioned do not accept inputs in CJK characters.
- Objects created using names containing CJK characters through the **Storage**, **ObjectAccess**, **NFS**, **CIFS**, **FTP**, and **SmartIO** modules are not accepted as inputs for commands in the modules that have not been I18N ported.
  For example, a file system that is created using a CJK language cannot be used to create an object in a module that has not been I18N ported.
- A **scale-out file system** cannot have its name in any other language than English.

Table 1-1 I18N supported features

Module name	Object
Pool	`pool_name`
File system	`pool_name` `fs_name`
Policy	`policy_name` `pattern`

Table 1-1 I18N supported features *(continued)*

Module name	Object
Worm	`fs_name`
Quota	`fs_name`

Rollback	`fs_name`
	`rollback_name`
	`cache_name`

Snapshot	`fs_name`
	`snapshot_name`
	`schedule_name`

FSCK or Defrag	`fs_name`

Compression	`fs_name`
	`pattern`
	`schedule_name`

NFS	`share path` **Note:** GNFS supports share names only in English.

CIFS	`file system path`, `directory path`
	`share_name`
	`fs_name` of the home directory
	Note: `Local user name` and `group management name` is supported only in English.

FTP	`fs_name` in the home directory `fs_name` in the anonymous log in the directory

SmartIO	`fs_name`
	`cache_name`
	`file_name`

Not supported in this release

Support for the following features is not present in this release:

- Veritas Access does not support changes to PAM or other OS- security settings. If you make any changes to PAM or other security settings, the authentication settings may not work.
  For example, you may not be able to change passwords for the master and new users that are added by using CLISH.
- Replication is not supported for IPv6 addresses.

Technical preview features

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

Veritas Access Streamer as a storage type for Enterprise Vault

Choosing Veritas Access Streamer as a storage type for Enterprise Vault is a technical preview feature in this release.

Note: This feature is supported only on test and development environments. It is not supported on production environments.

You are required to run Enterprise Vault 11 and later versions.

You can use the Veritas Access Streamer setup wizard to install Veritas Access Streamer. The Veritas Access Streamer installer can be found at the following location:

`dvd1-redhatlinux/rhel7_x86_64/EV_Streamer/Veritas_Access_Streamer_Setup.msi`

To install Veritas Access Streamer

1. Run the Veritas Access Streamer installer. You are prompted to choose the location where you want to install it. You have to choose the default location.

   Click **Next**.
2. The installer is ready to install Veritas Access Streamer on your system. Click **Next** to start the installation.
3. A window pops up which shows the progress of the installation. Once the installation is complete, click **Close** to exit the installation.
4. Open an administrator command prompt and navigate to `C:\program files(x86)\Enterprise Vault\Veritas Access Streamer`.

5 Run `regsvr32 VeritasAccessStreamer.dll`. You get a pop-up message that the registration is successful.

6 Go to `C:\program files(x86)\Enterprise Vault\Veritas Access Streamer\xml` to get the `EvExtendedSetttings.xml` file and configure Veritas Access Streamer as a storage type for Enterprise Vault to make the Veritas Access Streamer device known to the Enterprise Vault Administration Console.

See "To configure Veritas Access Streamer as a storage type for Enterprise Vault" on page 13.

7 Create a new partition and verify that Veritas Access Streamer is listed as one of the storage options.

You can perform the following steps on the Enterprise Vault server.

To configure Veritas Access Streamer as a storage type for Enterprise Vault

1 Open Windows Explorer, and then navigate to **<Program Files (x86)>\Enterprise Vault\InitialConfigurationData\en\Policies**.

2 Make a copy of `EVExtendedSettings.xml`.

3 Replace `EVExtendedSettings.xml` with the version provided by Veritas. Use the xml file created in `C:\Program Files(x86)\Enterprise Vault\Veritas Access Stream\xml`. The xml file is available after the `Veritas Access Streamer setup.msi` is installed.

See "To install Veritas Access Streamer" on page 12.

4 Update the registry value:
`[HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\KVS\Enterprise Vault\Admin] "PopulateExtendedSettingTypes"="1"`.

5 Close and then relaunch the Enterprise Vault Administration Console (VAC).

6 Navigate to **Policies > Exchange**.

7 Right-click **Exchange**, and then click **Populate Setting Types**.

A message is displayed that indicates that the **SettingsType** table in the Directory database has been successfully populated.

8 Restart the Storage service on all Enterprise Vault storage servers using Veritas Access Streamer as storage.

Once the services start, Veritas Access Streamer is displayed as a storage type when configuring a partition.

9 Select **Veritas Access Streamer** and click on **Next**.

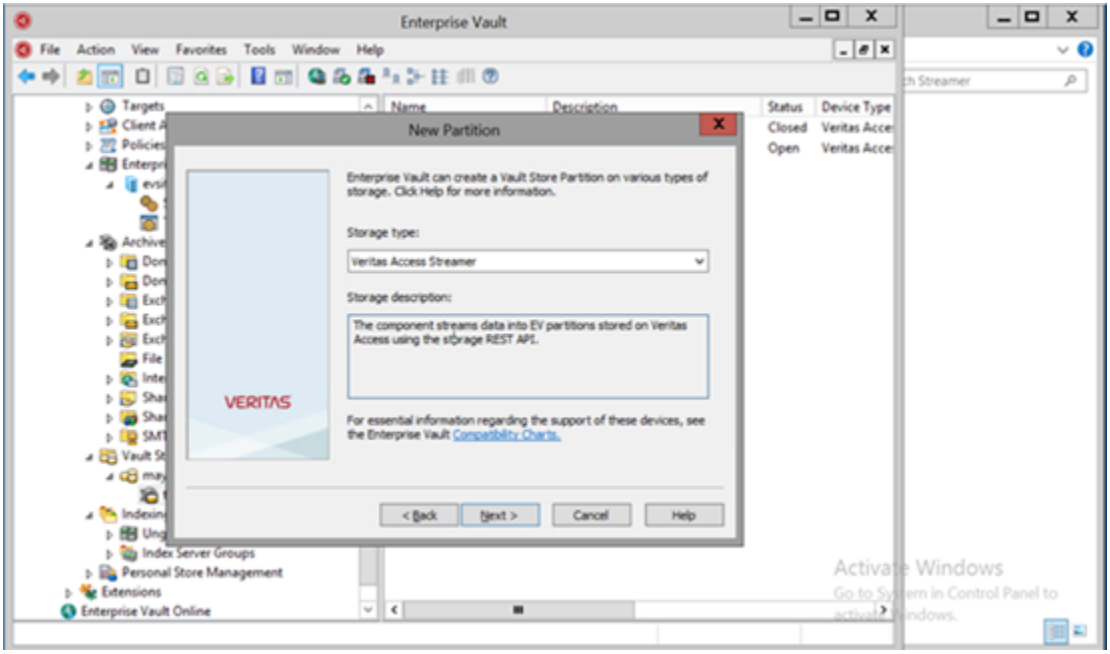

Now, you have to configure the Veritas Access Streamer.

To configure the Veritas Access Streamer

1 Configure the properties of the Veritas Access Streamer.

Name	Definition	Example value
Veritas Access Address	The hostname of the Veritas Access storage system.	s3.isaa
Veritas Access Port	The port where the server on Veritas Access Storage system listens for http requests.	8143
Veritas Access Access Key	The access key of the user to access the bucket.	********
Veritas Access Secret Key	The secret key of the user to access the bucket.	********
http/https	Denotes whether SSL should be used to connect to server.	Yes
Veritas Access Bucket	The bucket where the partition data will be stored.	evbucket

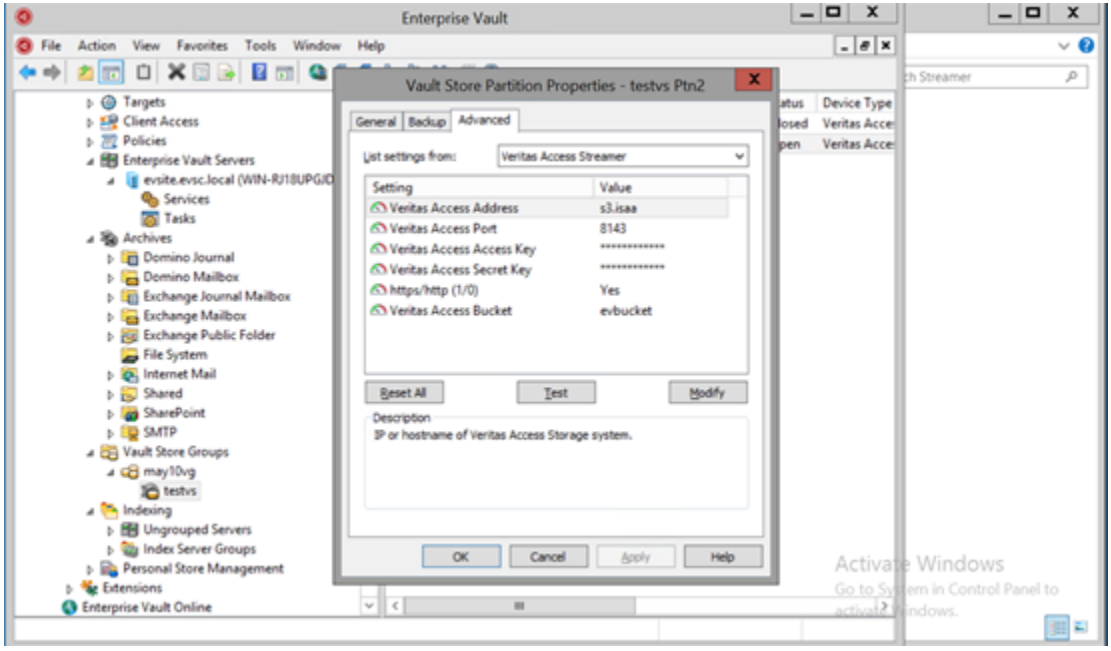

2 Go to **Advanced** on the partition settings and click **test**. You will get a pop-up message that tells you that your connection test is successful.

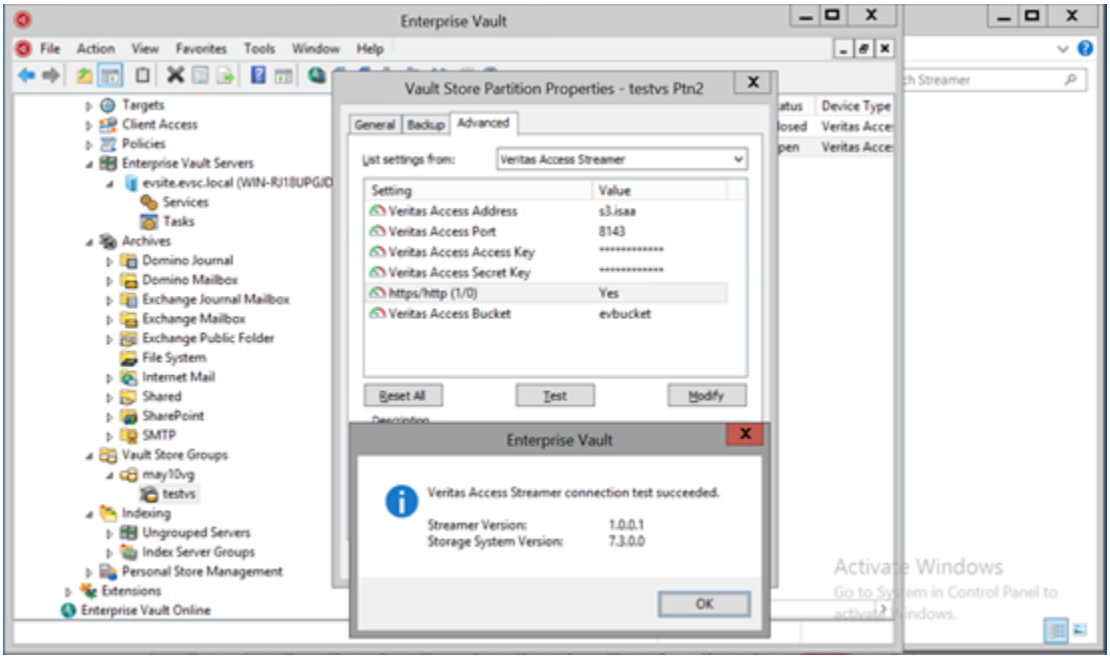

Support for erasure coding in a scale-out file system for an LTR use case over S3 protocol

Support for erasure coding in a scale-out file system for an LTR use case over S3 protocol is a technical preview feature in this release.

A scale-out file system can be created with an erasure-coded layout for an LTR use case.

Erasure coding (EC) is configured without the EC log for better performance.

The scale-out file system erasure-coded buckets are created through LTR policies and Veritas Access.

See the *Veritas Access Online Help* for more information on LTR policies.

Use the following command to set the parameters in Veritas Access (ObjectAccess) for creation of scale-out erasure-coded buckets through Veritas Access.

```
ObjectAccess> set fs_type largefs ecoded 4 2 16k stripe_aligned=yes
stripe_tag=disk rotating_parity=yes
```

Note: In this release, erasure coding for a scale-out file system is not supported for other use cases apart from the LTR use case over S3 protocol.

See the *Veritas Access Command Reference Guide* for information on parameters for a scale-out erasure-coded file system.

Veritas Access simple storage service (S3) APIs

Table 1-2 gives a list of the Veritas Access simple storage service (S3) APIs.

Table 1-2 Veritas Access simple storage service (S3) APIs

API	Description
`abort-multipart-upload`	Abort a multipart upload.
`complete-multipart-upload`	Complete a multipart upload by assembling previously uploaded parts.
`create-multipart-upload`	Start a multipart upload.
`delete-bucket`	Delete the bucket.
`delete-object`	Delete the specified object in the bucket.
`get-bucket-acl`	Get the ACLs for a bucket.
`get-bucket-(list objects) Version 1`	List of all the objects in a bucket.
`get-bucket-(list objects) Version 2`	List of all the objects in a bucket.
`get-bucket-location`	Get the bucket's region of the object.
`get-object`	Retrieve objects from a Veritas Access S3 bucket.
`get-service`	List of all buckets which are owned by the authenticated sender.
`head-bucket`	Determine if a bucket exists or not.
`head-object`	Retrieve metadata from an object without returning the object itself.
`initiate-multipart-upload`	Initiate a multipart upload and returns an upload ID.

Table 1-2 Veritas Access simple storage service (S3) APIs *(continued)*

API	Description
`list-multipart-uploads`	List the in-progress multipart uploads.
`list-parts`	List the parts that have been uploaded for a specific multipart upload.
`put-bucket`	Create a new bucket.
`put-bucket-acl`	Set permission on the existing bucket by using an ACL.
`put-object-copy`	Create a copy of an object that is already stored in the Veritas Access S3 server.
`put-object`	Add an object to a bucket.
`upload-part`	Upload a part in a multipart upload.
`upload-part-copy`	Upload a part by copying data from an existing object.

See the *Veritas Access Restful API Guide* for more information on the Veritas Access simple storage service (S3) APIs.

Chapter 2

Fixed issues

This chapter includes the following topics:

Fixed issues in this release

This section includes the issues fixed since the last release.

Table 2-1 Fixed issues since the last release

Fixed issues	Description
IA-9087	Removing or modifying the virtual IP associated to an OpenDedup volume leads to the OpenDedup volume going into an inconsistent state
IA-9659	OpenDedup port is blocked if the firewall is disabled and then enabled again
IA-9730	The OpenDedup volume is not mounted automatically by the /etc/fstab on the media server after a restart operation
	Removing a network device whose IP address is used by OpenDedup volume(s), affects the backup jobs which use that volume
IA-1943	ObjectAccess server goes in to faulted state while doing multi-part upload of a 10-GB file with a chunk size of 5 MB
IA-5737	ObjectAccess operations do not work correctly in virtual hosted-style addressing when SSL is enabled
IA-7434	Temporary objects may be present in the bucket in case of multi-part upload
IA-8747, IA-8649	If there are more than 1000 FTP sessions, the FTP> session showdetails command takes a very long time to respond or hangs
IA-9697	LIO does not support target name in uppercase

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

Fixed issues	Description
IA-7685	Error while setting full access permission to Enterprise Vault user for archival directory

Chapter 3

Software limitations

This chapter includes the following topics:

Limitations on using shared LUNs

The following limitations relate to shared LUNs in Veritas Access.

Veritas Access does not support thin LUNs

Veritas Access does not support thin LUNs. Some CLISH commands may fail if thin LUNs are used.

Flexible Storage Sharing limitations

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

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

When formatting the DAS disks, the disks are given unique names. The names include the embedded cluster name. There is a limit of 25 characters for a DAS disk name. When choosing the cluster name for a cluster that has DAS disks, you must limit the cluster name to ten characters.

Limitations related to installation and upgrade

The following limitations are related to installation and upgrade.

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]
```

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

The following Veritas Access modules are not supported for IPv6:

- NIS

FTP limitations

The following limitation applies to FTP.

- Multiprotocol access of FTP with other protocols such as NFS, CIFS is not supported.

Intel Spectre Meltdown limitation

The following limitation applies to Intel Spectre Meltdown.

Veritas Access recommends upgrading the kernel to one of the following versions which will resolve the Intel Spectre Meltdown issue that is caused by the kernel.

- RHEL 7.4: 3.10.0-693.21.1.el7.x86_64
- RHEL 7.3: 3.10.0-514.36.5.el7.x86_64

Retpoline is a spectre v2 mitigation technique. The kernel and product code privileged part should be compiled with retpoline support. Veritas Access code is not compiled with retpoline support.

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.

Limitations on using InfiniBand NICs in the Veritas Access cluster

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

  Note: The case is observed on Mellanox card.

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

  Note: The case is observed on Mellanox card.

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

Limitations related to commands in a non-SSH environment

Some CLISH commands work only when passwordless SSH is configured for the root user. If the `/opt/VRTSnas/con/communication.conf` file exists then, `CommunicationType` key is set to `SSH`.

For example:

```
# cat /opt/VRTSnas/conf/communication.conf
{
    "WorkingVersion": "1",
    "Version": "1",
    "CommunicationType": "SSH"
}
```

The following commands work only when passwordless SSH communication is enabled for the root user:

- `Backup> install`
- `Cluster> addnode`
- `Cluster> delnode`
- `Cluster> reboot`
- `Cluster> shutdown`
- `FTP> logupload`
- `License> add`
- All `Replication>` commands
- `Report> exportevents`
- `Report> snmp exportmib`
- `Storage> fencing on` (for majority-based fencing)
- `Storage> fencing off` (for majority-based fencing)
- `System> config import`
- `System> config import remote`
- `System> config export`
- `System> config export remote`
- All `Support>` commands
- `Upgrade> add`
- `Upgrade> install`

Limitation on using Veritas Access in a virtual machine environment

Veritas Access is not supported on KVM-based virtual machines.

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.

File system limitation

The following limitations relate 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.
- When a file system is created, an additional file system is also created for the purpose of keeping the lock and configuration information. The additional file system is not directly accessible to the user. It is meant for internal use only.
  It is recommended to use disks from as many nodes as possible when creating the first storage pool in Veritas Access. In case of a shared nothing environment

where the disks are local to the cluster nodes, the additional file system mirrors are created across all those nodes. This ensures that the Veritas Access configuration is available even if one of the nodes on which the additional file system was created is available.

In case of SAN environments the additional file system is mirrored across two disks.

- On-premises tiering is not supported for a scale-out file system.
- On-premises tiering in a cluster file system only supports one primary and one secondary.

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 (LTR) limitations

The following limitations are related to LTR.

- Veritas Access does not support the HTTPS application protocol for an S3 bucket from the GUI in Veritas NetBackup long-term retention (LTR) use cases.
- In case a cluster node serving the OpenDedup volume crashes, the ongoing NetBackup jobs on that particular OpenDedup volume may fail. But the same NetBackup job will be successful in the next retry which is triggered automatically by NetBackup. The NetBackup job may again restart when the crashed node comes up and the IP failbacks.

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 "master" 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.

Chapter 4

Known issues

This chapter includes the following topics:

Veritas Access known issues

The following known issues relate to the Veritas Access commands.

Admin issues

This section describes known issues related to the admin module.

The user password gets displayed in the logs for the Admin> user add username system-admin|storage-admin|master command

If you execute the `Admin> user add` *`username`* `system-admin|storage-admin|master` command and enter the password with the command (which is an optional parameter), the user password gets displayed in the logs. This happens because every command that is executed on CLISH is logged on the `admin.log` and `command.log`. Since the password is also a part of the CLISH command, the password also gets logged.

Workaround:

There is no workaround for this issue.

Veritas recommends that when you create new users, you provide the password on the CLI only when prompted.

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.

Windows client displays incorrect CIFS home directory share size

After you map a CIFS share on a Windows client, the CIFS home directory share size that is displayed on the Windows client is incorrect.

Workaround:

There is no workaround. You can run the `Storage> fs list` command from CLISH to get the correct size of the CIFS home directory share.

Deduplication issues

This section describes known issues related to deduplication.

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

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

Workaround:

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

To enable the duduplication job:

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

Enterprise Vault Attach known issues

The following known issues relate to Enterprise Vault Attach.

The Enterprise Vault archival policy does not create the ev_archival folder in the share 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. If the archival policy does not create the `ev_archival` folder, you have to explicitly create this folder and check the ownership of the folder before you assign this folder as partition to Enterprise Vault.

Workaround:

Create an empty folder and configure it as a vault store partition inside the share created by the archival policy.

FTP issues

The following issues relate to the Veritas Access FTP commands.

If a file system is used as homedir or anonymous_login_dir for FTP, this file system cannot be destroyed

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

Workaround:

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

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

The FTP> server start command reports the FTP server to be online even when it is not online

The `FTP> server start` command sometimes reports that it successfully started the FTP server but due to an internal issue, the online operation actually fails.

Workaround:

Use the `FTP> server status` command to verify the status of the FTP service. If the FTP service is offline, run the `FTP> server start` command again or run the `Support> service autofix` command to fix the faults, if any.

The FTP> session showdetails user=<AD username> command does not work

The `FTP> session showdetails` command takes the *AD username* as an additional filter parameter. You can specify the user name to filter out sessions belonging to that particular user. This command does not work if *AD username* is in the format of DOMAINNAME/USERNAME. This is due to an internal parsing issue.

Workaround:

Add an escape character (\) in between the domain name and the AD user name.

For example, if the user name is *domain**username*, use *domain**username* in the `FTP> session showdetails user=<AD username>` command.

If the security in CIFS is not set to Active Directory (AD), you cannot log on to FTP through the AD user

The AD configuration in Veritas Access is common across protocols. It is configured through the `CIFS> set domain/domaincontroller/domainuser` commands. Hence, if the user wants to use AD as security in FTP, the AD configuration has to be done through CIFS. If the security in CIFS is not set to AD, you cannot log on to FTP through the AD user. Any changes to the AD configuration that are done through a CIFS session have implications on FTP also.

Workaround:

There is no workaround for this issue as the AD configuration is common across all protocols in Veritas Access. Make sure that the AD is configured correctly using the `CIFS> set` commands. To use it in protocols other than CIFS, set the security in CIFS to AD.

If security is set to local, FTP does not work in case of a fresh operating system and Veritas Access installation

When security is set to local, the `+/home/ftpuse+r` directory is not present on the console node. Hence, you cannot login using FTP.

Workaround:

Use a virtual IP of a different node (other than the master node) to log in.

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.

After you restart a node that uses RDMA LLT, LLT does not work, or the gabconifg –a command shows the jeopardy state

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

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

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

Workaround:

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

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

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

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

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

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

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

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

Running individual Veritas Access scripts may return inconsistent return codes

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.

Excluding PCIs from the configuration fails when you configure Veritas Access using a response file

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

Workaround:

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

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

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 the same driver node is used for two installations at the same time, then the second installation shows the status of progress of the first installation

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

Workaround:

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

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

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

Workaround:

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

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

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

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

Workaround:

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

If duplicate PCI IDs are added for the PCI exclusion, the Cluster> add node name command fails

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

Workaround:

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

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

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

Workaround:

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

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

After finishing system verification checks, the installer displays a warning message about missing third-party RPMs

After finishing system verification checks, the installer displays a warning message about missing required third-party RPMs or that the RPMs need to be upgraded. The warning message indicates that the verification checks completed successfully.

The missing third-party required RPMs are installed or upgraded from the Veritas Access ISO image during the installation process.

Workaround:

You can safely ignore this warning message.

Installer appears to hang when you use the installaccess command to install and configure the product from a node of the cluster

If you try to install and configure the product from a node of the cluster by using the `installaccess` command, the installer appears to hang after the 'Redefining network configurations' session. The installer does not hang, it just takes a long time to execute.

Workaround:

Wait for the installer to complete the configuration. Once the network configurations are redefined, the installer takes around 20 minutes to complete the remaining tasks. You can also avoid this issue by installing and configuring the product from the third node using the `access72` command.

After phase 1 of rolling upgrade is complete on the first node, a panic occurs on the second node

When you perform a rolling upgrade on the Veritas Access cluster, after phase 1 is complete on the first node, a panic occurs on the second node. The panic is triggered because the second node still remains in the old product version while the first node has the new product version.

Workaround:

Wait for the second node to come out from the panic. This takes about 10 minutes. Then, you can continue with the rolling upgrade procedure on the cluster.

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 fails to install if LDAP or the autofs home directories are preconfigured on the system

The Veritas Access installation (7.x) may fail if the following conditions exist:

- LDAP is configured on the system
- The autofs home directories are configured on the system

This can create problems during the installation of the user home directories that are required for the installation of Veritas Access.

Workaround:

Do not configure LDAP or the autofs home directories on systems before installing Veritas Access.

When performing a rolling upgrade from Veritas Access 7.3.0.1 to 7.4 on RHEL 7.3, CIFS services get into a faulted stated after the nodes are upgraded to Veritas Access 7.4

After the rolling upgrade is completed on the cluster nodes, soft links are not created for the CIFS configuration files. Hence, the CIFS service do not come online. This causes the CIFS I/O path to be unavailable for a long time.

Workaround:

Create the soft links for the CIFS configuration files on every node of cluster manually as soon as the upgrade is complete on that node. This brings the CIFS services online on that node. Run the following commands on each node of the cluster after the rolling upgrade is complete on that node:

```
# ln -sf /opt/VRTSnas/scripts/cifs/SambaServer_online
/opt/VRTSvcs/bin/SambaServer/online
 # ln -sf /opt/VRTSnas/scripts/cifs/cifs_va_options /etc/sysconfig/samba
```

After the Veritas Access installation is complete, the installer does not clean the SSH keys of the driver node on the Veritas Access nodes from where the installation is triggered.

When Veritas Access is installed from the driver node, the installer does not delete SSH keys, which are saved in */root/.ssh/authorized_keys*. Therefore, you can SSH to the Veritas Access nodes without a password even after the installation is complete.

Workaround:

Check the SSH key of the driver node and delete the key from all the Veritas Access nodes of the cluster.

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.

Some Phantomgroups do not come online after a rolling upgrade

After rolling upgrade some Phantomgroups do not come online because some of the resources of the Phantomgroups are in offline state. The `phantomproc_<interface_name>` resource of the `Phantomgroup_pubeth<number>` does not come online. Hence, the `Phantomgroup_pubeth<number>` does not come online.

Workaround:

Use the following command to bring the resource online.

```
[root@varnic_01 ~]# hares -online phantomproc_ens161 -sys varnic_01
[root@varnic_01 ~]# hares -state phantomproc_ens161
#Resource          Attribute             System     Value
phantomproc_ens161 State                 varnic_01  ONLINE
phantomproc_ens161 State                 varnic_02  ONLINE
[root@varnic_01 ~]#
```

If you do not know the complete resource name, then perform the following steps to find the resource name and then bring it online.

To find the resource name and then bring it online

1. Find the name of the Phantomgroup which is offline using the `hastatus -sum` command.
2. Find the resources for the Phantomgroup using the following command.

```
[root@varnic_01 ~]# hagrp -resources Phantomgroup_pubeth0
phantomproc_ens161
phantomNIC_ens161
```

3. Check the state of the resources using the following command.

```
[root@varnic_01 ~]# hares -state phantomproc_ens161
#Resource          Attribute             System     Value
phantomproc_ens161 State                 varnic_01  OFFLINE
phantomproc_ens161 State                 varnic_02  ONLINE
```

4. Bring the resource online using the following command.

```
[root@varnic_01 ~]# hares -online phantomproc_ens161 -sys varnic_01
[root@varnic_01 ~]# hares -state phantomproc_ens161
#Resource          Attribute             System     Value
phantomproc_ens161 State                 varnic_01  ONLINE
phantomproc_ens161 State                 varnic_02  ONLINE
[root@varnic_01 ~]#
```

5. Verify the state of the Phantomgroup. The status should be online now.

```
[root@varnic_01 ~]# hagrp -state Phantomgroup_pubeth0
#Group               Attribute             System     Value
Phantomgroup_pubeth0 State                 varnic_01  |ONLINE|
Phantomgroup_pubeth0 State                 varnic_02  |ONLINE|
[root@varnic_01 ~]#
```

Protocol versions are different on nodes after rolling upgrade is performed

While performing a rolling upgrade on cluster nodes, it may happen that Phase 1 is not successful on any one of the cluster nodes and Phase 2 is performed on all the nodes. In such a scenario, the protocol version on the failed node is lower. Hence, the nodes are not able to form one cluster as protocol versions are different on the nodes.

Workaround:

Perform Phase 1 of the rolling upgrade on all the nodes of the cluster. If Phase 1 is successful on all the nodes, then perform Phase 2 on the nodes. This way, the protocol versions will be upgraded in all the nodes and the nodes will be able to form a cluster after the rolling upgrade.

Installing Veritas Access with preconfigured VLAN and preconfigured bond fails

If you try to install Veritas Access with preconfigured VLAN and preconfigured bond, then the installation fails. This is because during installation, you can either preconfigure VLAN or you can preconfigure the bond as public device but not both at the same time.

Workaround:

After installation, you can create a bond over a particular network interface by using the `Network> bond create` command. You can create VLAN using the `Network> vlan create` command.

Internationalization (I18N) issues

This section describes known issues related to I18N.

The CLISH prompt disappears when characters in a foreign language are present in a command

English and non-English language characters have different character encoding. Hence, the CLISH prompt disappears when there are foreign characters in the command and you try to modify the command using the up and down arrow keys.

Workaround:

You can use any one of the following methods:

- Log out and log into CLISH again.
- Press `Ctrl + C`.
- Set the locale to the intended non-English language. Start CLISH. The supported languages are Chinese, Japanese, and Korean.

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.

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

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

Workaround:

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

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

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

Workaround:

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

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

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

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

Workaround:

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

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
```

Unable to import the network module after an operating system upgrade

The Veritas Access 7.4 release supports NIC name retention feature. Hence, you cannot import the network module if you perform an operating system upgrade.

Workaround:

Before you install Veritas Access 7.4, rename the public NICs as public0, public1 ... and so on. Rename the private NICs as priveth0 and priveth1.

LDAP with ‘SSL on’ does not work if you upgrade Veritas Access

If you perform an upgrade of Veritas Access from 7.3.x to 7.4, the following command does not work because there is a bug in upgrade path which does not ask for correct LDAP certificate from the user.

```
# network> ldap set ssl on
```

Hence, LDAP with ‘SSL on’ does not work after an upgrade.

Workaround:

After the upgrade is complete, again set the 'SSL on' option using the following command:

```
# network> ldap set ssl on
```

Network loadbalancer does not get configured with IPv6

If you configured loadbalancer using CLISH with an IPv6 virtual IP, the loadbalancer configuration appears to be successful but does not balance the load in the background. This is because the loadbalancer is not supported with IPv6.

Workaround:

There is no workaround.

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.

NFS-Ganesha share for IPv6 subnet does not work and NFS share becomes faulted

When NFS-Ganesha (GNFS) server is enabled, then NFS export to IPv6 subnet is not supported. Hence, the NFS share becomes faulted. This is a limitation at the GNFS level.

Workaround:

There is no workaround. You have to create separate NFS exports for each client.

ObjectAccess issues

This section describes ObjectAccess issues.

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 you have upgraded to Veritas Access 7.4 from an earlier release, access to S3 server fails if the cluster name has uppercase letters

If the cluster name has uppercase letters, access to the S3 server fails. This is due to a limitation of the underlying library that is used to accept S3 requests.

Workaround:

Use all lowercase letters to access the S3 server.

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

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.

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.

The pool name is not displayed in the output of the Objectaccess> bucket show command for the existing buckets after OpenDedup is upgraded from Veritas Access 7.3.0.1 to 7.4

The ObjectAccess configuration is not available for import when you upgrade from Veritas Access 7.3.0.1 to 7.4. Hence, after the upgrade, the pool field is empty for the existing buckets. The pool name is not displayed in the output of the `objectaccess> bucket show` command. This does not break any functionality.

Workaround:

There is no workaround for this issue.

Portald crashes for the get_object API while sending response for bigger payloads if SSL is enabled

If SSL is enabled, `portald` for the `get_object` API crashes while sending response for bigger payloads. This problem is also observed when conditional parameters are used while SSL is enabled.

Workaround:

Disable SSL when using conditional parameters.

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.

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.

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.

The OpenDedup> volume create command does not revert the changes if the command fails during execution

If the `Opendedup> volume create` command fails during execution, the changes that have been made are not reverted.

Workaround:

Execute the `Opendedup> volume list` command to check if the volume has been created but is in offline state. If the volume is in offline state, delete the volume using the `Opendedup> volume delete <volume-name>` command and try to re-create the volume.

Some of the OpenDedup volume stats reset to zero after upgrade

If the OpenDedup volume configuration is exported when volumes are online, the exported data does not reflect the state of the volumes properly. This results in some of the stats being reset to zero while importing the OpenDedup volume configuration.

Workaround:

Before you export the OpenDedup volume configuration, offline all the volumes and then export the configuration.

OpenDedup volume mount operation fails with an error

When you try to mount an OpenDedup volume, it fails with the following error:

```
Still running according to PID file /var/run/S3fs********.pid,
PID is ****. Service exit with a return value of 122
```

Stale `jsvc` processes are observed after unmounting the OpenDedup volume. Hence, volume re-mount fails.

Workaround:

Kill the stale `jsvc` processes manually and try to mount the OpenDedup volume again.

Restore of data from AWS glacier fails

When data is uploaded on the AWS Glacier cloud, and you try to restore the data from the cloud using NetBackup restore, the operation fails with the following error:

```
socket read failed errno = 62 - timer expired
```

Workaround:

- In the `/etc/sdfs/*-volume-cfg.xml` file, change the value of `glacier-archive-days` from 0 to 30.
- Check the value of the `Client Read Timeout` for the media server that is used for the restore operation.

In the NetBackup Administration Console, you can find the `Client Read Timeout` at **Host properties** > **Media Servers** > **Timeouts** > **Client Read Timeout**.
On the UNIX media server, see the value of `CLIENT_READ_TIMEOUT` in the `/usr/openv/netbackup/bp.conf` file. Set the `CLIENT_READ_TIMEOUT` value to a big number such as 3600, 7200 or 10800.
For more details, see https://www.veritas.com/support/en_US/article.100006172

OpenDedup volumes are not online after an OpenDedup upgrade if there is a change in the cluster name

The configuration file of the OpenDedup volumes contains S3 endpoint based on the previous cluster name. Hence, after an upgrade, the volumes do not come online on a cluster with a different name due to change in S3 endpoint.

Workaround:

Contact Veritas Technical Support to update the S3 endpoint in the configuration file.

If the Veritas Access master node is restarted when a restore job is in progress and OpenDedup resides on the media server, the restored files may be in inconsistent state

If OpenDedup resides on the media server, and you restart the Veritas Access master node when a restore job is in progress, the restored files may not be consistent with the source file which has been backed-up.

Workaround:

Perform the following:

- Ensure that all the backup and restore jobs that are currently running are stopped.
- Unmount the OpenDedup volume. If any stale `jsvc` processes are running, kill them manually.
- To delete the cache, delete everything from the `/opt/sdfs/volumes/<volume-name>/chunkstore/chunks/` file and remount the OpenDedup volume.
- Restart the restore jobs.

The OpenDedup> volume list command may not show the node IP for a volume

In some cases, the `Opendedup> volume list` command does not display the node IP for the OpenDedup volume.

Workaround:

You can get the associated node IP for the OpenDedup volume using the following command.

```
# hares -display $(grep <OpenDedup_volume_name>
/opt/VRTSnas/conf/odd_vipgrp_map.conf | awk '{ print $2 }' |
tr -d 'group') -attribute Address | tail -1 | awk '{ print $NF }'
```

Where <*OpenDedup_volume_name*> is the name of the OpenDedup volume whose node IP does not get displayed.

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.

Cinder volume creation fails after a failure occurs on the target side

Sometimes a cinder volume creation operation fails. This is an intermediate issue. The volume creation fails as the `_vrts_get_targets_store` function returns a blank target list output. If you check in the Cinder logs, you see the following error message:

```
ERROR oslo_messaging.rpc.server return target_list
['output']['output']['targets']
ERROR oslo_messaging.rpc.server TypeError:
'bool' object has no attribute '_getitem'_
```

Workaround:

Restart the Cinder volume service using the `service openstack-cinder-volume restart` command.

Cinder volume may fail to attach to the instance

The Cinder volume fails to attach to the instance as the `_/lib/udev/scsi_id --page 0x83 --whitelisted /dev/disk/by-path/ip-10.182.97.58:3260-iscsi-iqn.2018-02.com.veritas:target02-lun-4_` command returns an error.

Workaround:

Check the Cinder logs. If the volume attach has failed and you get an error that *_ scsi_id: cannot open /dev/disk/by-path/ip-10.182.97.58:3260-iscsi-iqn.2018-02.com.veritas:target02-lun-4: No such device or address_* , then delete that volume. Create a new volume and try to attach the volume.

Replication issues

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

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
```

SDS known issues

The following issue relates to the SDS module.

After the SDS log is rotated, the log messages from either Veritas Access or the SDS plugin go to the rotated file instead of the new file

The SDS log (located at `/var/log/sds.log`) is rotated every day. SDS contains two services, the SDS plugin and the Veritas Access plugin, which share the log file. Because of a bug in Python, a race condition occurs and one of the two services

starts logging to a new file while the other service continues logging in the rotated file.

Workaround:

Check the last updated timestamps of the current log file and the rotated log file and the, co-relate the logs from different services for debugging.

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.

File system creation fails with SSD pool

The file system creation with `layout=mirror` operation fails when the pool has SSDs from two or more nodes.

Workaround:

Create the file system using available SAN/DAS disks.

For the disks present in the pool of type SSD, run the following command from the bash shell as `Support` user to export the disks on all the nodes from where the disks are physically present.

```
Support> vxdisk export disk name
```

After all the disks in the pool are exported from the respective cluster nodes, proceed with the file system creation from the Veritas Access CLISH.

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
```

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.

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>
   ```

The Storage> fs create command does not display the output correctly if one of the nodes of the cluster is in unknown state

If one of the nodes of the cluster is in unknown state, then the `Storage> fs create` command behaves differently. Though the file system is created successfully, the output does not get displayed correctly.

Workaround:

If you want to create the file system using GUI, then bring the node online. Else, If you want to create the file system even if one node is in unknown state, then create the file system from CLISH. You can verify that the file system has been created using the `Storage> fs list` command.

Storage> fs growby and growto commands fail if the size of the file system or bucket is full

The `Storage> fs growby` and `Storage> fs growto` commands fail if there is no free space in the file system or the bucket.

Workaround:

There is no workaround. You can delete files manually to create free space.

While provisioning an S3 bucket for NetBackup, the bucket creation fails if the device protection is selected as erasure-coded and the failure domain is selected as disk

When you try to provision an S3 bucket for NetBackup using the Veritas Access GUI, the bucket creation fails if the device protection is selected as erasure-coded and the failure domain is selected as disk. This happens because the selection of the failure domain as disk is not passed to the `vxassist` command. If the failure domain is not specified, by default, the command take the failure domain as node. Hence, bucket creation fails.

Workaround:

There is no workaround. You cannot select the failure domain as disk for an erasure-coded layout while provisioning storage from the Veritas Access GUI as it is not supported in this release.

The operating system names of fencing disks are not consistent across the Veritas Access cluster which may lead to issues

The disks that are used for fencing across the cluster may not have the same operating system names. For example, a disk that is called *sda* on one node may be called as *sdf* on another node. This means that the *sda* disk on both the nodes are not the same. This can lead to writes on unintended disks when setting up disk-based SCSI3 fencing.

Workaround:

Ensure that the same operating system disk names are used for all the disks which are used for fencing across the cluster.

The disk group import operation fails and all the services go into failed state when fencing is enabled

If the disks are not SCSI-3 compliant, the SCSI-3 persistent reservation inquiries have to be turned off from the Volume Manager side. Else, all the services go into faulted state when you try to enable fencing,

Workaround:

You can enable fencing with non-SCSI3 disks by following any one of the following methods.

To enable fencing with non-SCSI3 disks using the cluster> reboot all command

1. Install Veritas Access without enabling fencing.
2. Execute the `vxdctl scsi3_pr off` on all the nodes.
3. From CLISH, execute the `Cluster> reboot all`.
4. After the system restart, execute the `Storage> fencing on majority` from CLISH.
5. Create the pool and file system.

To enable fencing with non-SCSI3 disks without restart

1. Stop the cluster services.

   ```
   # hastop -all
   ```

2. After all the services go down, turn off the SCSI3 persistent reservations on all the nodes in the cluster.

   ```
   # vxdctl scsi3pr off
   ```

3. Get the process ID of `vxconfigd` and kill the `vxconfigd` process on all the nodes of the cluster.
4. Restart `vxconfigd` on all the nodes of the cluster.

   ```
   # /sbin/vxconfigd -k -x syslog
   ```

5. Start all the nodes of the cluster.

   ```
   # vxclustadm -m vcs startnode
   ```

   Wait for the disk group to get imported.

6. Start HA service on all the nodes of the cluster.

   ```
   #hastart
   ```

Now, you can enable fencing.

While creating an erasure-coded file system, a misleading message leads to issues in the execution of the storage> fs create command

When you create an erasure-coded file system using the `Storage> fs create` command, the following content appears as part of the help message:

```
eclogdisk : comma separated list of disks from at least ndata+nparity
failure domains to be used for EC log allocation. To use default disks,
pass eclogdisk=default []
```

This message is misleading and if you specify `eclogdisk=default` as an option, the command does not get executed successfully.

Workaround:

Ignore the help description for `eclogdisk`. Pass the value for `eclogdisk` as *"default"* instead of `eclogdisk=default` when you create a file system.

The Veritas Access cluster node can get explicitly ejected or aborted from the cluster during recovery when another node joins the cluster after a restart

This issue is observed if any file system resources are in faulted state on all the cluster nodes.

Workaround:

Trigger a cluster join for the aborted node and manually online the file system(s).

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 the Veritas Access target commands.

Storage provisioning commands hang on the Veritas Access initiator when LUNs from the Veritas Access target are being used

When file system holding iSCSI LUNs come in online state from faulted state on the Veritas Access target, the Veritas Access initiator is not able to recognize the LUNs.

Workaround:

Restart the target service on the Veritas Access target.

After the Veritas Access cluster recovers from a storage disconnect, the iSCSI LUNs exported from Veritas Access as iSCSI target may show wrong content on the initiator side

There is a synchronization issue between the Veritas Access as iSCSI target and any other iSCSI initiator. Hence, when the Veritas Access cluster recovers from a storage disconnect, the iSCSI LUNs that are exported from the Veritas Access cluster appear to have wrong content. This is because the initiator does not correctly reconnect to the disks.

Workaround:

Perform the following steps:

1. Detach the volume from OpenStack instances.
2. Restart the iSCSI target services using the `target> iscsi service stop` and `target> iscsi service start` commands.
3. Attach the volumes back the respective instances.
4. Mount the volume to the respective instances. You can see the correct content on the initiator side.

Chapter 5

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 documents are available on the SORT site:

- *Veritas Access Administrator's Guide*
- *Veritas Access Cloud Storage Tiering Solutions Guide*
- *Veritas Access Command Reference Guide*
- *Veritas Access Enterprise Vault Solutions Guide*
- *Veritas Access Getting Started Guide*
- *Veritas Access Installation Guide*
- *Veritas Access NetBackup Solutions Guide*
- *Veritas Access Quick Start Guide*
- *Veritas Access Release Notes*
- *Veritas Access RESTful API Guide*
- *Veritas Access Software-Defined Storage (SDS) Management Platform Solutions Guide*
- *Veritas Access SDS Management Platform Quick Start Guide*
- *Veritas Access Third-Party License Agreements*
- *Veritas Access Troubleshooting Guide*