Veritas Access 7.3 Release Notes

Linux

7.3

April 2019

VERITAS™

Veritas Access Release Notes

Last updated: 2019-04-04

Document version: 7.3 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 7

About this release 7
Important release information 7
Changes in this release 8
Changes to the GUI 8
Additional cloud providers 8
Scale-out file system enhancements 9
Installer enhancements 9
Multi-protocol support for NFS with S3 9
Support for kernel-based NFS version 4 9
Managing application I/O workloads using maximum IOPS settings 9
Veritas Access sync replication 10
WORM storage for Enterprise Vault archiving 10
Creation of Partition Secure Notification (PSN) file for Enterprise Vault archiving 10
Changing firewall settings 10
Setting retention in files 10
Not supported in this release 11
Technical preview features 11
Veritas Access as an iSCSI Target 11
IP load balancing 12
Erasure coding for Object Store buckets 13
Veritas Access Streamer as a storage type for Enterprise Vault 13

Chapter 2 Fixed issues 19

Fixed issues since the last release 19

Chapter 3 Software limitations 20

Limitations on using shared LUNs 20
Flexible Storage Sharing limitations 21
If your cluster has DAS disks, you must limit the cluster name to ten characters at installation time 21

Limitations related to installation and upgrade 21
If required VIPs are not configured, then services like NFS, CIFS, and S3 do not function properly 21
Rolling upgrade is not supported from CLISH 21
Limitations in the Backup mode 21
Veritas Access IPv6 limitations 22
FTP create_homedirs limitation 22
Samba ACL performance-related issues 22
Veritas Access language support 22
Veritas Access does not support non-English characters when using the CLISH (3595280) 23
Limitations on using InfiniBand NICs in the Veritas Access cluster 23
Limitation on using Veritas Access in a virtual machine environment 23
NFS-Ganesha limitations 24
Kernel-based NFS v4 limitations 24
File system limitation 24
Any direct NLM operations from CLISH can lead to system instability (IA-1640) 24
Veritas Access S3 server limitation 25
LTR limitations 25

Chapter 4 Known issues 26

Veritas Access known issues 26
AWS issues 26
Backup issues 27
CIFS issues 27
Deduplication issues 29
Enterprise Vault Attach known issues 29
FTP issues 30
GUI issues 30
Installation and configuration issues 32
Networking issues 39
NFS issues 40
ObjectAccess issues 44
OpenDedup issues 47
OpenStack issues 47
Replication issues 48
SmartIO issues 52
Storage issues 53

Chapter 5 Getting help 62

Displaying the online Help 62
Displaying the man pages 62
Using the Veritas Access product documentation 62

Chapter 1

Overview of Veritas Access

This chapter includes the following topics:

- About this release
- Important release information
- Changes in this release
- Technical preview features

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/7.3/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.000127371

Changes in this release

This section shows the major new features and enhancements added in the 7.3 version of Veritas Access.

Changes to the GUI

The following updates were made to the GUI:

- Support for growing disks, formatting disks, marking disks as spare, and removing disks from the cluster.
- Ability to set the maximum IOPS limit.
- Support for volume-level and file system replication.
- Storage Policies - Synchronous Data Protection policy for protection against device, node, and site failures, including replication support.
- Archival Policies - Archival of data using Enterprise Vault (EV) for CIFS shares.
- Support for additional cloud providers, Alibaba, Azure, Amazon GovCloud(US), IBM Cloud Object Storage, and Google cloud for scale-out file systems
- NTP server management.
- Licencing updates.
- Configuration of the Key Management Service (KMS)
- Time management support.
- Default parameters and group-related parameters for S3 buckets.
- SmartIO cache creation and deletion.

Additional cloud providers

You can use the following cloud providers for moving data from on-premises storage to cloud storage for a scale-out file system:

- Amazon S3
- Amazon Glacier

- Amazon GovCloud (US)
- Azure
- Google cloud
- Alibaba
- IBM Cloud Object Storage

Scale-out file system enhancements

The following scale-out file system enhancements were added in this release:

- Support for growing a file system
- Support for multiple cloud tiers (up to eight cloud tiers)
  For example, you can configure Azure and AWS Glacier as two cloud tiers and move data between these clouds.

Installer enhancements

You can reduce the number of required IP addresses at installation time by not configuring virtual IP addresses. Virtual IP addresses can be added post-installation using the `Network> ip addr add` command in the CLISH.

Multi-protocol support for NFS with S3

Beginning with the Veritas Access 7.3 release, you can add multi-protocol support for NFS with S3. If an NFS share is present (and objects may be present in the exported path), the storage administrator can map that path as an S3 bucket (S3 over NFS). In addition, a normal file system path can also be mapped as an S3 bucket. The buckets created by S3 APIs cannot be exported as an NFS share (NFS over S3).

Support for kernel-based NFS version 4

Veritas Access supports both NFS v3 and NFS v4. Both NFS v3 and NFS v4 support Kerberos authentication.

Managing application I/O workloads using maximum IOPS settings

When multiple applications use a common storage subsystem, it is important to balance the application I/O requests in a way that allows all the applications to co-exist in a shared environment so that a particular application does not monopolize the storage bandwidth. You can address this need by setting a maximum threshold

on the I/O operations per second (MAXIOPS) for the file system. The MAXIOPS limit determines the maximum number of I/Os processed per second collectively by the storage underlying the file system.

Veritas Access sync replication

The Veritas Access sync replication solution provides high performance, robustness, ease of use, and synchronous replication capability which is designed to contribute to an effective disaster recovery plan. Veritas Access provides both command-line interface (CLISH) and the graphical user interface (GUI) for online management of the synchronous replication. It also maintains write-order fidelity and performs replication of volumes in synchronous mode, ensuring data integrity and consistency.

WORM storage for Enterprise Vault archiving

Veritas Access can be configured as WORM primary storage for archival by Enterprise Vault. Veritas Access 7.3 is certified as a CIFS primary WORM storage for Enterprise Vault 12.1.

See the *Veritas Access Enterprise Vault Solutions Guide* for more details on this feature.

Creation of Partition Secure Notification (PSN) file for Enterprise Vault archiving

A Partition Secure Notification (PSN) file is created at a source partition after the successful backup of the partition at the remote site. See the *Veritas Access Enterprise Vault Solutions Guide* for more details on this feature.

Changing firewall settings

The `Network> firewall` command can be used to view or change the firewall settings.

See the `Network> firewall` man page for detailed examples.

Setting retention in files

The retention feature provides a way to ensure that the files are not deleted or modified until retention is applied on the files. You can set, clear, and show the retention on files from CLISH.

See the `Storage> fs` man page for detailed examples.

Not supported in this release

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

- Metadata-only replication
- Multi-source multi-target replication
- On the wire encryption
- Wild-card support for replication
- Upgrade of replication jobs is not supported from previous versions of Veritas Access. Contact Veritas Technical Support if upgrade is a mandatory requirement.

Technical preview features

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

Veritas Access as an iSCSI Target

Veritas Access as an iSCSI target is added as a preview feature in the 7.2.1 release.

The following functionality is available in this feature:

- Veritas Access can be configured as an iSCSI target to serve block storage.
- The iSCSI target service is hosted in active-passive mode in the Veritas Access cluster.
- Once configured, the cluster is available to any standard iSCSI initiator over a portal IP.
- You can perform the following functions on an iSCSI target:
  - Starting and stopping of the iSCSI target service
  - Addition and deletion of targets
  - Addition and deletion of LUNs
  - Map and un-map initiators
  - Addition and deletion of users
- See the `target` manual pages for more information.

The following limitations are present in this feature:

- Fault injection scenarios have not been covered during testing. Hence, iSCSI functionality may not behave as per expectation.

  - In case of node reboot and cable pull scenarios, the feature may not behave as expected
  - LUN add or destroy operation when interleaved with the target service restart may put the cluster in an inconsistent state.
- Strong integration with the rest of the Veritas Access code is incomplete. For example, integration with the network bonding and VLAN feature is incomplete.
- Performance testing has not been done.

IP load balancing

IP load balancing is added to Veritas Access as a technical preview feature.

The purpose of this feature is to reduce the number of virtual IPs required for Veritas Access. With IP load balancing, a single virtual IP is used to act as a load balance IP which distributes the incoming request to the nodes.

The following functionality is available in this feature:

- One of the existing Veritas Access virtual IP is configured as the load balancer IP.
- All clients can connect to the Veritas Access cluster using this single virtual IP.
- Veritas Access makes use of load balancer algorithms internally to allocate the next available Veritas Access node to serve the client.
  Currently, the Veritas Access cluster makes use of the round-robin algorithm in the implementation of the load balancer.
- Enter the following command to configure the load balancer.

  ```
  Network> loadbalancer configure <VIP>
  ```

- Enter the following command to destroy the load balancer configuration.

  ```
  Network> loadbalancer remove
  ```

Note: Only NFSv3 is supported for this technical preview.

The following limitation is present in this feature:

- If a new virtual IP is added or a virtual IP fails over after the load balancer is set up, the load balancer configuration gets deleted. The load balancer needs to be set up again.

Erasure coding for Object Store buckets

Erasure coded (ecoded) file system is added as a technical review feature in Veritas Access. Erasure coding offers a more robust solution in redundancy and fault tolerance for critical storage archives. It is supported in DAS, SAN, FSS, and standalone environments. ObjectAccess buckets can be created to use erasure-coded volumes.

Veritas Access Streamer as a storage type for Enterprise Vault

Choosing Veritas Access Streamer as a storage type for Enterprise Vault is added as a preview feature in the 7.3 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-rhel6_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 14.

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 13.

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**

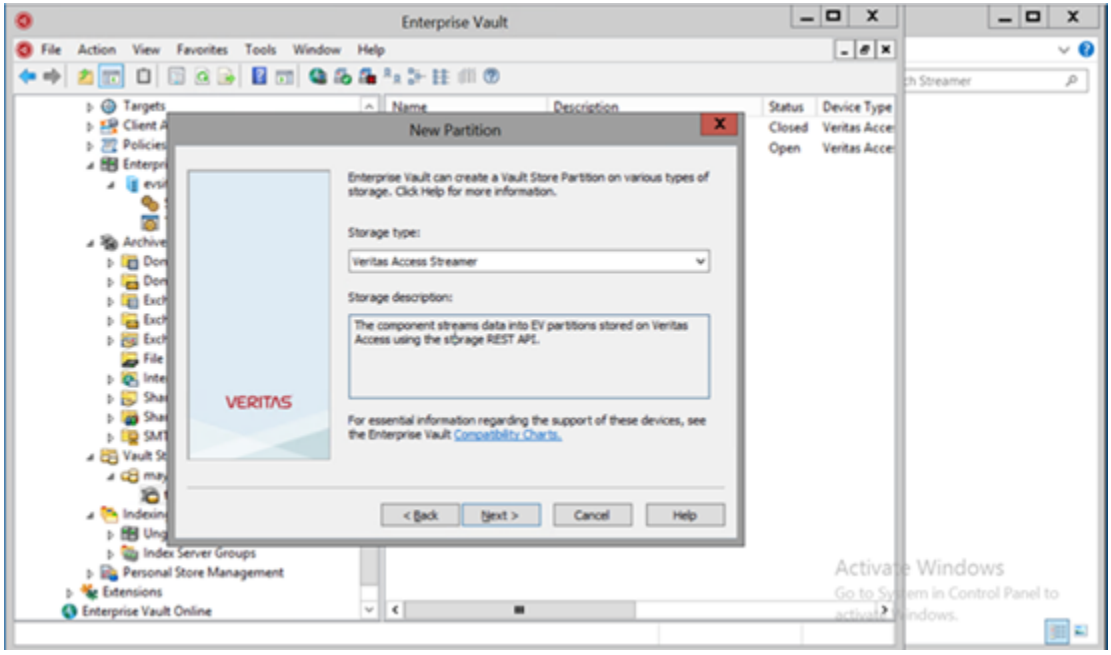

Now, you have to configure the Veritas Access Streamer

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

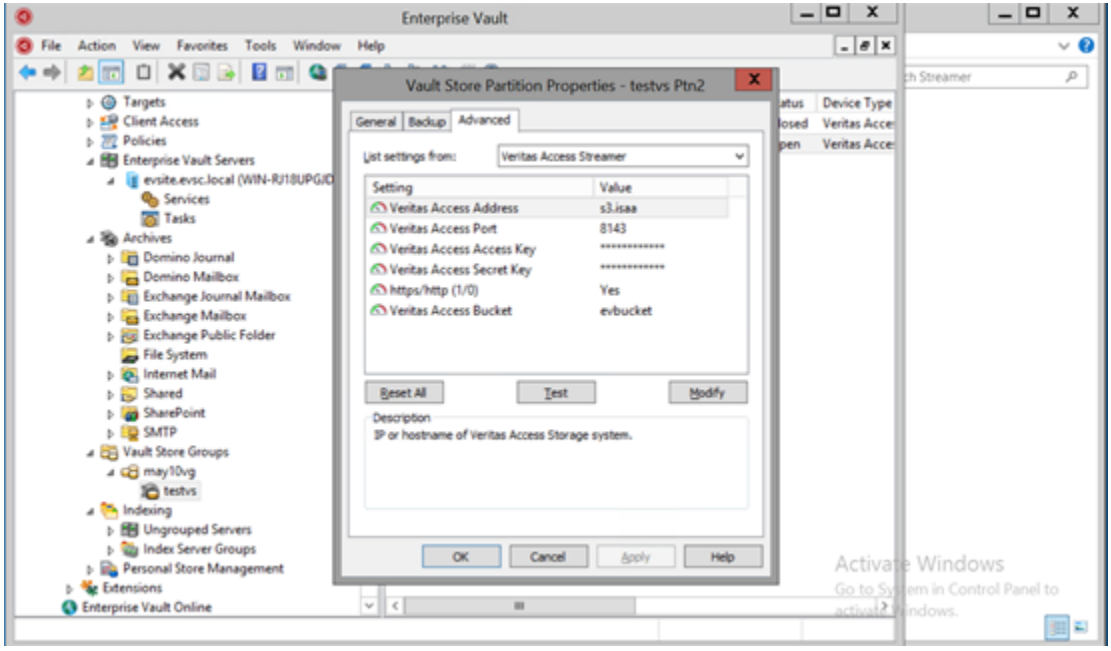

2 Go to **Advanced** on the partition settings and click **test**. You will get a pop-up message that tells you that your connection test is successful.

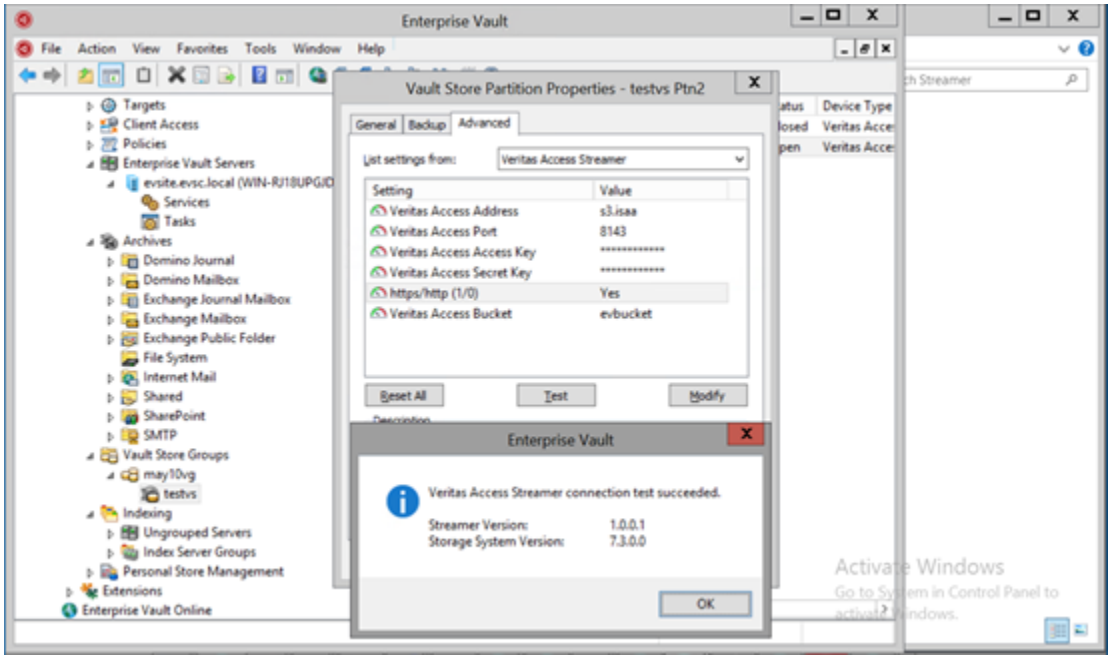

Chapter 2

Fixed issues

This chapter includes the following topics:

- Fixed issues since the last release

Fixed issues since the last release

This section includes the issues fixed since the last release.

Table 2-1 Fixed issues since the last release

Fixed issues	Description
IA-5621	After phase 1 of rolling upgrade is complete on the first node, a panic occurs on the second node
A-5761	Argparse module does not get installed during OS installation in RHEL 6.6
IA-4061	Storage> fs-growto and Storage> fs-growby commands give error with isolated disks
IA-3432	Pattern given as filter criteria to Storage> fs policy add sometimes erroneously transfers files that do not fit the criteria
IA-3240	Rollback cache grow option missing in CLISH
IA-5734	An erasure coded file system may report the file system as full even if free space available in the file system

Chapter 3

Software limitations

This chapter includes the following topics:

- Limitations on using shared LUNs
- Flexible Storage Sharing limitations
- Limitations related to installation and upgrade
- Limitations in the Backup mode
- Veritas Access IPv6 limitations
- FTP create_homedirs limitation
- Samba ACL performance-related issues
- Veritas Access language support
- Limitations on using InfiniBand NICs in the Veritas Access cluster
- Limitation on using Veritas Access in a virtual machine environment
- NFS-Ganesha limitations
- Kernel-based NFS v4 limitations
- File system limitation
- Veritas Access S3 server limitation
- LTR limitations

Limitations on using shared LUNs

The following issues 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

The following IPv6 functionality is not supported for CIFS:

- CIFS does not support IPv4/IPv6 mixed mode for the domain controller. The IPv4 DNS entry needs to be removed from the DNS server.
- CIFS does not accept IPv6 addresses for the domain controller in the Veritas Access CLI. Only hostnames are allowed for the domain controller entry.

FTP create_homedirs limitation

Due to a limitation, you must manually create the user's logon directory even if the `create_homedirs` option is set to `yes`.

Samba ACL performance-related issues

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

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

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

Veritas Access language support

Veritas Access supports only English.

Veritas Access does not support non-English characters when using the CLISH (3595280)

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

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

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.

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.
- NFSv4 share reservations are not supported.
- NFS v4 delegation is not supported.

File system limitation

The following issue relates to the Veritas Access file system.

Any direct NLM operations from CLISH can lead to system instability (IA-1640)

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

Veritas Access S3 server limitation

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

The object has to be downloaded in parts.

LTR limitations

Veritas Access does not support the https application protocol for S3 bucket from the GUI in NBU LTR use cases.

Chapter 4

Known issues

This chapter includes the following topics:

- Veritas Access known issues

Veritas Access known issues

The following known issues relate to the Veritas Access commands.

AWS issues

The following known issues relate to the AWS support.

The CLISH storage commands appear to hang when EBS disks are forcibly detached from the AWS console (IA-5042)

Forceful removal of EBS disks from the AWS console may lead to unexpected behavior like command execution slowing down and may require restarting the cluster nodes. So the CLISH storage commands appear to hang.

Workaround:

Contact Veritas Technical Support for guidance when detaching EBS volumes from the AWS console.

CIFS server start command fails on one of the nodes if the clustering mode is set to CTDB

The CIFS server in the CTDB clustering mode depends on the CTDB daemon to be started. The CTDB daemon gets stuck during the recovery process on one of the nodes. Usually, the first node on which VCS tries to start CTDB has both CIFS and CTDB in the ONLINE state while the second node remains in OFFLINE state.

Workaround:

There is no workaround. You can access the CIFS shares in normal clustering mode.

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 (3606322)

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 (3853674)

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 (3824576, 3836861)

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 (IA-4771)

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.

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 (3472414)

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:

Error while setting full access permission to Enterprise Vault user for archival directory (IA-7685)

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. In the 7.3 release, after the creation of the `ev_archival` folder by the archival policy, you have to explicitly change the ownership of the folder before you give full access permission to this folder from Windows.

Workaround:

Perform the following steps to change ownership:

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

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

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

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

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 (IA-1876)

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
```

GUI issues

The following issues relate to the GUI.

When both volume-level and file system replication links are set up in Veritas Access 7.3, provisioning of storage using High Availability and Data Protection policies does not work (IA-7646)

Performing the following steps leads to this scenario:

- Setting up both volume-level and file system replication links.
- Activating High Availability and Data Protection policies.
- Provisioning storage using either of these policies using the Provision Storage wizard.

- Setting up the replication job task fails.

This happens because even though you selected the file system replication link, during storage provisioning, the GUI selects the volume-level replication link for setting up the replication job, which causes the task to fail.

Workaround:

Do not create both volume-level replication link and file system 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 file system 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 (IA-7127)

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
```

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 (IA-1796)

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 (3796864)

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 (3794964)

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 (3686704)

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 (3659716)

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 (IA-3446)

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 (IA-3436)

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 (IA-900)

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 (IA-1850)

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 (IA-3570)

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 (IA-3611)

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 (IA-5300)

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.

Phantomgroup for the VLAN device does not come online if you create another VLAN device from CLISH after cluster configuration is done (IA-6671)

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*.

Rolling upgrade fails to bring the cfsmount resources online when you perform an upgrade from 7.2.1 version (IA-7388)

Rolling upgrade from Veritas Access 7.2.1 to 7.3 is not supported. Cfsmount resources do not become online on nodes on which rolling upgrade phase 1 has been performed when you perform a rolling upgrade from the 7.2.1 version.

Workaround:

There is no workaround for this issue.

Rolling upgrade is not allowed if the scale-out file system is online and being used by the NFS or S3 server

Rolling upgrade is not supported in scale-out file systems if the file systems are online and are being used by the NFS or S3 server.

Workaround:

If you want to perform a rolling upgrade from 7.2.1.1 to 7.3, you have to perform the following set of operations to bring some of the services down before you start rolling upgrade and offline the scale-out file systems. This leads to downtime for the applications using the scale-out file systems.

Log on to CLISH and run the following commands

1. If scale-out file system is NFS shared using the Ganesha server, then unmount the NFS share from the client and stop the NFS service on the Veritas Access management cluster node.

   ```
   CLISH> nfs server stop
   ```

2. If a bucket is created over the scale-out file system, then stop the ObjectAccess service on the Veritas Access management cluster node.

   ```
   CLISH> objectaccess server stop
   ```

3. List the file systems.

   ```
   CLISH> storage fs list
   ```

   Make a note of all the scale-out file systems with layout as 'largefs-simple', 'largefs-mirrored', 'largefs-mirrored-stripe', 'largefs-striped', and 'largefs-striped-mirror'.

4. Offline all the scale-out file systems of type 'largefs'.

   ```
   CLISH> storage fs offline <scaleout filesystem name>
   ```

5. Confirm if the scale-out file systems are offline.

   ```
   CLISH> storage fs list
   ```

Log on to any Veritas Access cluster node using the root/support user credentials and run the following command

- Offline the VCS group for scale-out file system.

```
# hagrp -offline vrts_tfs_infra -any -clus <Veritas Access cluster name>.
```

After rolling upgrade procedure is complete, you have to run the following commands to bring all the services online on the Veritas Access cluster nodes.

Log on to any Veritas Access cluster node using the root/support user credentials and run the following commands

1. Enable NFS v4 support.

```
# /opt/VRTSnas/scripts/nfs/nfs_agent.sh install
```

2. Run the following command on all the nodes in the cluster.

```
# echo "no" > /opt/VRTSnas/conf/nfs_force_stop
```

Log on to CLISH and run the following commands

1. Online all the scale-out file systems.

```
CLISH> storage fs online <scaleout filesystem name>
```

2. Confirm if the scale-out file systems are online.

```
CLISH> storage fs list
```

3. Start the NFS server and remount the NFS shares on the client.

```
CLISH> nfs server start
```

4. Verify that the NFS server has started.

```
CLISH> nfs server status
```

5. Start the S3 ObjectAccess server.

```
CLISH> objectaccess server start
```

6. Verify that the S3 ObjectAccess server has started.

```
CLISH> objectaccess server status
```

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.

Networking issues

This section describes known issues related to networking.

CVM service group goes into faulted state unexpectedly (3793413)

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 (3616561)

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 (3559219)

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 (3596284)

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.

NFS issues

This section describes NFS issues.

Slow performance with Solaris 10 clients with NFS-Ganesha version 4 (IA-1302)

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 (IA-1304)

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 (IA-1002)

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 (IA-1838)

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(IA-849)

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 (IA-1844)

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 (3816074, 3819836)

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 (3847271)

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 (3828442)

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 (IA-2758)

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 (IA-1938)

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 (IA-6629)

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 (IA-5083)

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 (IA-6672)

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

Workaround:

Restart the NFS service.

ObjectAccess issues

This section describes ObjectAccess issues.

ObjectAccess server goes in to faulted state while doing multi-part upload of a 10-GB file with a chunk size of 5 MB (IA-1943)

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

Workaround:

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

When trying to connect to the S3 server over SSLS3, the client application may give a warning like "SSL3_GET_SERVER_CERTIFICATE:certificate verify failed" (IA-5378)

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 7.3 from an earlier release, access to S3 server fails if the cluster name has upper case letters (IA-5628)

If the cluster name has upper case letters, access to the S3 server fails. This is due to a limitation of the underlying library which 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 ObjectAccess service in Veritas Access (IA-5631)

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

Workaround:

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

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

ObjectAccess operations do not work correctly in virtual hosted-style addressing when SSL is enabled (IA-5737)

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

Workaround:

Use path-style access when SSL is enabled.

ObjectAccess server enable operation fails on a single node (IA-5704)

The ObjectAccess server enable operation assumes at least a two-node cluster setup. Hence, the `server enable` command fails.

Workaround:

There is no workaround for this issue.

ObjectAccess (S3) service goes OFFLINE when the node is restarted (IA-6282)

When a node in a cluster comes up after a system restart, the service groups are started automatically as per the AutoStartList attribute. But the designed flow to online service group is interrupted because the `hagrp -online` command also attempts to start the `ReconfigGroup` and `vrts_vea_cfs_int_cfsmount1` service groups. This causes the S3 service to remain in [OFFLINE] state.

Workaround:

Start S3 service from CLISH using the `objectaccess server start` command. You can start the service if the system has an active license.

Bucket creation may fail with "Timed out Error" (IA-7432)

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.

Temporary objects may be present in the bucket in case of multi-part upload (IA-7434)

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

Workaround:

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

Bucket CreationDate is incorrect if the bucket is created by mapping the filesystem path (IA-7227)

If S3 bucket is created by mapping the filesystem path, then subsequent operations on that bucket updates the CreationDate of the bucket.

Workaround:

If the bucket is created by mapping the filesystem path, do not rely on the value of the CreationDate of the bucket.

Group configuration does not work in ObjectAccess if the group name contains a space (IA-7407)

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.

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

An erasure coded file system may show mirrored layout in the Storage> fs list command (IA-7266)

While creating an erasure coded file system, the data volumes and the metadata volumes are created separately. The layout of the metadata volume is mirrored. Sometimes, a mirrored volume creates a Data Change Object (DCO) as well. In such cases, the `Storage> fs list` command shows the layout as mirrored.

Workaround:

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

Accessing a bucket or object in the S3 server fails with S3 internal errors

If the nodes of a cluster fail abruptly or if there is a crash, the named metadata attributes stored in the extended attribute of the file system may not get flushed to the disk correctly. Hence, the correct value of the named attribute data is not recovered.

This causes the operations which are dependent on the named attribute to fail. The S3 server reports ACL corruption errors in the log. The object data or directory information remains valid but named attribute may be in invalid state. This results in failure in accessing the bucket or objects in the S3 server.

Workaround:

The fix for this issue will be available in the next patch release. Veritas recommends that you take preventive measure to avoid system crash when S3 write operations are in progress.

OpenDedup issues

This section describes known issues related to OpenDedup.

If OpenDedup is installed on Veritas Access 7.3, then failover is not supported.

Veritas Access 7.3 does not support failover if OpenDedup is installed.

Workaround:

There is no workaround for this issue.

OpenDedup is not highly available

OpenDedup is not highly available. Hence, if the node that hosts OpenDedup goes down, ongoing and future backups are affected.

Workaround:

There is no workaround available

OpenStack issues

The following issues are related to OpenStack.

Cinder and Manila shares cannot be distinguished from the CLISH (3763836)

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.

Replication issues

This section describes known issues related to replication.

Running replication and dedup over the same source, the replication file system fails in certain scenarios (3804751)

The replication job may fail when the following situations occur on the same source 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 replication keys and jobs (3822515)

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

Workaround:

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

1 Make sure the new target binds the replication IP, because the replication IP is not changed on the new source.

2 Run the `Replication> config import_keys` command on the source and the target.

3 Run the `Replication> 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 replication failover (3668957)

This issue occurs if the schedules on the source cluster and the target cluster have the same name but different intervals. After 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.

Replication fails with error “connection reset by peer” if the target node fails over (IA-3290)

Replication creates a connection between the source and the target to replicate data. 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, replication fails with an error message. If there is a scheduled replication job, the next iteration continues this failed replication session, possibly with a new node from the target.

Workaround:

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

Synchronous replication shows file system layout as mirrored in case of simple and striped file system (IA-7308)

When you configure simple or striped file system under synchronous replication, it displays its layout as mirrored if you execute the `Storage> fs list` command. Even if you unconfigure the sync replication from that file system, it continues to display the layout as a mirrored layout.

Workaround::

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

Synchronous 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> sync status test_fs
Name                    value
====================  ============================================
Replicated Data Set    rvg_test_fs

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 for synchronous 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> sync status test_fs
Name                    value
====================  ======================
Replicated Data Set    rvg_test_fs

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
```

If you restart any node in the primary or secondary cluster, replication may go into PAUSED state (IA-7567)

When you restart any node in the primary or secondary cluster, the IPTABLE rules communication 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 in the primary as well as secondary site.

```
# iptables -F
```

Sync replication failback does not work (IA-7524)

If you try to make the original source cluster as the new target cluster when the source cluster becomes available, the failback command on the original source cluster does not work. Hence, failback of the sync replication is not successful.

Workaround:

There is no workaround for this issue.

Replication jobs created in Veritas Access 7.2.1.1 or earlier versions are not recognized after upgrade to 7.3 version (IA-7597)

If you try to access or modify the 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 replication (IA- 7295)

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

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

Sync replication fails when the 'had' daemon is restarted on the target manually (IA-7357)

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

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 (IA-3423)

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 (IA-1542)

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 (IA-1733)

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 (IA-1639)

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 (IA-1851)

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 (3801216)

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 (3650635)

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 (3452098)

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 (IA-3545)

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 (3588248)

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 (IA-3269)

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 (IA-3427)

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 (IA-3091)

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.

Rollback service group goes in faulted state when respective cache object is full and there is no way to clear the state (IA-3251)

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 (IA-3239)

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

Workaround:

There is no workaround.

Storage device fails with SIGBUS signal causing the abnormal termination of the scale-out file system daemon (IA-2915)

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 (IA-3241)

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

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 (IA-3398)

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 (IA-5434)

`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.

Unable to create space-optimized rollback when tiering is present (IA-5690)

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 fencing on a setup with volume manager objects present fails to import the disk group (IA-7219)

If you enable fencing on a setup 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 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 fencing.

For the rollback cache growto and growby operations, the cache size values cannot be specified in terms of g/G, m/M or k/K (IA-7473)

You cannot specify the cache size values in terms of g/G, m/M, or k/K (like 10G, 10M, or 10K) for the `rollback cache growto` and `rollback cache growby` operations.

Workaround:

Enter the cache size in terms of 512-bytes units.

To calculate the cache size, convert the cache size that you want in KB and multiply by 2.

Example:

To grow the cache size to 10G, do the following:

10G = 10485760 KB

Cache size = 10485760 KB * 2 = 20971520 KB

File system creation fails when the pool contains only one disk (IA-7515)

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 (IA-7174)

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 (IA-7604)

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 (IA-7628)

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.

Chapter 5

Getting help

This chapter includes the following topics:

- Displaying the online Help
- Displaying the man pages
- Using the Veritas Access product documentation

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 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 Third-Party License Agreements*
- *Veritas Access Troubleshooting Guide*
- *Veritas Access Enterprise Vault Solutions Guide*