

Veritas Access 7.3.0.1 Release Notes

Linux

7.3.0.1

Veritas Access Release Notes

Last updated: 2017-11-09

Document version: 7.3.0.1 Rev 0

Legal Notice

Copyright © 2017 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

Veritas Technologies LLC
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	6
	About this release	6
	Important release information	6
	Changes in this release	7
	Support for creating CIFS shares for a scale-out file system	7
	File sharing for a scale-out file system using FTP	7
	WORM support over NFS	7
	Support for RHEL 7.3	7
	GUI enhancements	7
	Not supported in this release	8
	Technical preview features	8
	IP load balancing	8
	Erasure coding for Object Store buckets	9
	Veritas Access Streamer as a storage type for Enterprise Vault	9
Chapter 2	Fixed issues	15
	Fixed issues since the last release	15
Chapter 3	Software limitations	16
	Limitations on using shared LUNs	17
	Flexible Storage Sharing limitations	17
	If your cluster has DAS disks, you must limit the cluster name to ten characters at installation time	17
	Limitations related to installation and upgrade	17
	If required VIPs are not configured, then services like NFS, CIFS, and S3 do not function properly	17
	Limitations in the Backup mode	17
	Veritas Access IPv6 limitations	18
	FTP limitations	18
	Samba ACL performance-related issues	18
	Veritas Access language support	18
	Veritas Access does not support non-English characters when using the CLISH (3595280)	19

	Limitations on using InfiniBand NICs in the Veritas Access cluster	19
	Limitation on using Veritas Access in a virtual machine environment	19
	NFS-Ganesha limitations	20
	Kernel-based NFS v4 limitations	20
	File system limitation	20
	Any direct NLM operations from CLISH can lead to system instability (IA-1640)	20
	Veritas Access S3 server limitation	21
	LTR limitations	21
	Limitation related to replication authentication	21
Chapter 4	Known issues	22
	Veritas Access known issues	22
	AWS issues	22
	Backup issues	23
	CIFS issues	23
	Deduplication issues	25
	Enterprise Vault Attach known issues	26
	FTP issues	26
	GUI issues	28
	Installation and configuration issues	29
	Networking issues	34
	NFS issues	35
	ObjectAccess issues	39
	OpenDedup issues	42
	OpenStack issues	42
	Replication issues	43
	SmartIO issues	49
	Storage issues	49
	System issues	59
Chapter 5	Getting help	61
	Displaying the online Help	61
	Displaying the man pages	61
	Using the Veritas Access product documentation	61

Overview of Veritas Access

This chapter includes the following topics:

- [About this release](#)
- [Important release information](#)
- [Changes in this release](#)
- [Not supported 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/7.3.0.1/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 7.3.0.1 version of Veritas Access.

Support for creating CIFS shares for a scale-out file system

Veritas Access provides support for creating CIFS shares for a scale-out file system.

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

File sharing for a scale-out file system using FTP

Veritas Access provides support for sharing files using the FTP protocol for a scale-out file system.

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

WORM support over NFS

Veritas Access supports WORM over NFS. Once a file is WORM-enabled, writes are not allowed from the NFS client and the server. Deletion of the file is not allowed until the retention period expires.

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

Support for RHEL 7.3

The Veritas Access 7.3.0.1 release supports RHEL 7 Update 3 only. The RPMs are included in the Veritas Access ISO.

GUI enhancements

The following are the enhancements to the GUI:

- Veritas Access provides support for LDAP for clients to access and to manage the user authentication data, which is stored on the LDAP server.

- Veritas Access provides support for FTP for client access to files in Veritas Access for secure and non-secure FTP.

Not supported in this release

- Veritas Access as iSCSI target is not supported.
- FTP with LDAP as security is not supported.
- Veritas Access 7.3.0.1 release does not support upgrade of any type from an earlier release.

Technical preview features

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

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 preview 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 technical 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:

`dvdl-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 `EvExtendedSettings.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 10.
- 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 `EVEExtendedSettings.xml`.
- 3 Replace `EVEExtendedSettings.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 9.
- 4 Update the registry value:

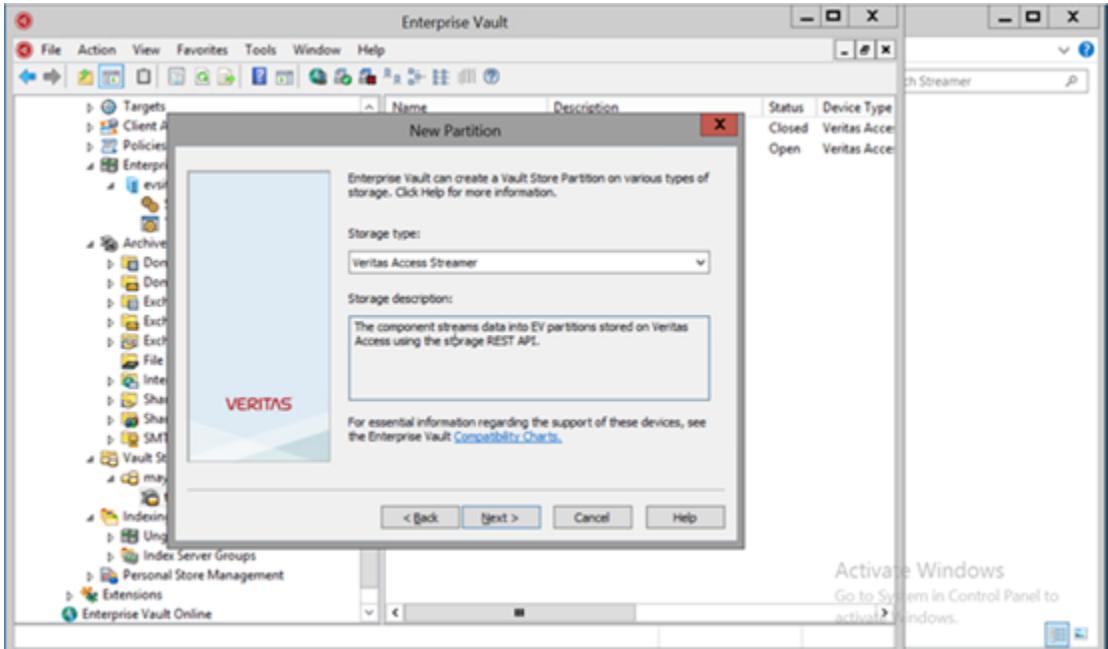
```
[HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\KVS\Enterprise Vault\Admin] "PopulateExtendedSettingTypes"="1".
```
- 5 Close and then relaunch the Enterprise Vault Administration Console (VAC).
- 6 Navigate to **Policies > Exchange**.
- 7 Right-click **Exchange**, and then click **Populate Setting Types**.

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**.

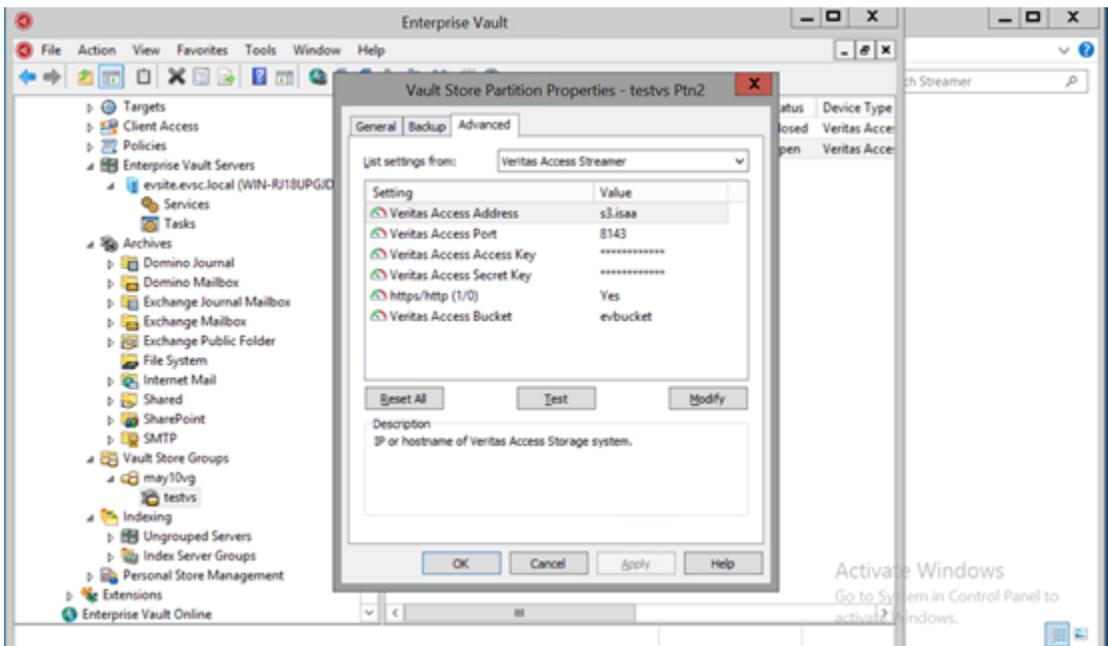


Now, you have to configure the Veritas Access Streamer.

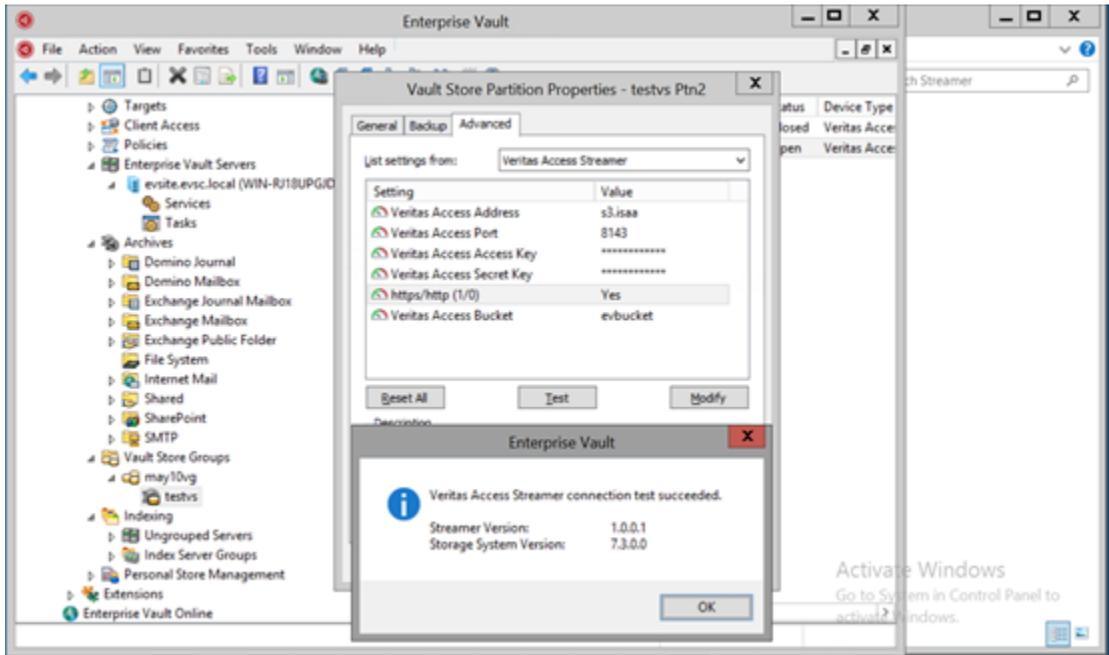
To configure the Veritas Access Streamer

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.



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-3570	If installing using a response file is started from the cluster node, then the installation session gets terminated after the configuring NICs section
IA-3240	Rollback cache grow option missing in CLISH

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 limitations](#)
- [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](#)
- [Limitation related to replication authentication](#)

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]
```

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 limitations

The following limitations apply to FTP.

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

Samba ACL performance-related issues

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

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

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

Veritas Access language support

Veritas Access supports only English.

Veritas Access does not support non-English characters when using the CLISH (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.
- NFS v4 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.

Limitation related to replication authentication

While creating a replication link, you have to provide the "master" user credentials to authenticate a different cluster for replication.

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

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

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.

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
```

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.

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 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.

Modifying or adding a CIFS share with non-default fs_mode/owner/user options results in these settings getting applied to all the existing shares on the same path

When you modify or add a CIFS share with non-default fs_mode/owner/user options, the owner, the group and the fs_mode for all the existing shares on the same path also change. This behavior is observed only if you have more than one CIFS share on the same path. These three options are applied on path and are not share-specific. Hence, if you change these settings for one share, the change is applied to all the shares that share the path.

Workaround:

If different fs_mode/owner/user options are required for different shares, create the shares out of different file systems or paths.

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.

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 deduplication 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

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

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 `anonymous_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:

Execute the `FTP> session showdetails` command without any filter and manually search for a particular user in the output.

If there are more than 1000 FTP sessions, the FTP> session showdetails command takes a very long time to respond or hangs

The `FTP> session showdetails` command displays information about all the active FTP sessions. As the number of sessions increases, the command takes a longer time to respond. If some of the sessions close while the command is in progress, the command becomes unresponsive and appears to hang.

Workaround:

Use the filter settings like *client_ip*, *server_ip*, and *user* that are available as optional arguments in the command. In case of a hang, press `CTRL + C` and try to execute the command again.

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

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.

Installation and configuration issues

The following issues relate to Veritas Access installation and configuration.

Running individual Veritas Access scripts may return inconsistent return codes

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

Excluding PCI's 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 PCI's 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.

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

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

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

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

Workaround:

To work around this issue:

- 1 Recover the SSH connection manually.
- 2 Uninstall Veritas Access.
- 3 Reinstall Veritas Access.

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

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

Workaround:

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

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.

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.

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.

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:

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

The state of phantomgroup becomes *ONLINE*.

During installation, detection of network interfaces may fail in some cases

The Veritas Access installer uses the `lshal` utility to detect the network links. But the `lshal` utility may fail to detect the network cards in some cases because the `hald` daemon does not work properly.

Workaround:

Restart the `hald` daemon and its dependency, `dbus` on the system.

Enter the following commands to re-run the Veritas Access installer:

```
# /etc/init.d/messagebus restart
# /etc/init.d/haldaemon restart
```

After the installation and configuration of Veritas Access 7.3.0.1, one of the nodes fails to restart

After installing and configuring Veritas Access 7.3.0.1, the installer tries to restart all the nodes of cluster. Even though the installer gives the message that it is restarting all the nodes, one of the nodes fails to restart because the ssh connection breaks before the restart command is triggered.

Workaround:

Restart the node manually to add the node to the cluster.

The `nsLCD.service` does not restart after failure

The `nsLCD` service fails to restart because of an RHEL7 bug. Bug #1312297 is already open for it.

For more information, refer to https://bugzilla.redhat.com/show_bug.cgi?id=1312297

Workaround:

There is no workaround for this issue. The bug is still in open state from Red Hat side.

Networking issues

This section describes known issues related to networking.

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.

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.

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

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

Workaround:

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

```
# hastart
```

NFS issues

This section describes NFS issues.

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.

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.

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.

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-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> 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 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.

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

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

Workaround

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

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

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

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

Workaround:

You can verify the given IP addresses by using DNS.

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

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

Workaround:

The client can remount the exports to access the shares.

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

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

Workaround:

There is no workaround for this issue.

Kernel NFS v4 export mount for Netgroup does not work correctly

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

Workaround:

Restart the NFS service.

ObjectAccess issues

This section describes ObjectAccess issues.

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

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

Workaround:

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

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

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

Workaround:

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

If the cluster name does not follow the DNS hostname restrictions, you cannot work with 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>

An erasure coded file system may report the file system as full even if free space available in the file system

While creating an erasure coded file system, the data volumes and the metadata volumes are created separately. Whenever data gets copied, there is an option to move data on either of the volumes. Placement policies decide the volume on which

data will reside. Currently, while creating a file system by using erasure coded layout, the placement policy does not get set properly which may lead to data being stored on the metadata volume. Thus, the metadata volume gets filled faster.

Workaround:

Use the `grow` command frequently to create space in the file system.

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

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

Workaround:

Use path-style access when SSL is enabled.

ObjectAccess server enable operation fails on a single node

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

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"

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

Workaround:

Client needs to verify bucket deletion even if request fails.

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

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

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

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

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.

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.

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.

Replication issues

This section describes known issues related to replication.

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 replication cycle fails with the error:

```
SSL certificate error.
```

Workaround:

Follow these steps to solve this issue:

- 1 Execute the `Replication> 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> job add_link` command to add the link back to the job, and enable or sync the replication job.

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

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

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.

Replication job status shows the entry for a link that was removed

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

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

Workaround:

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

The job uses the schedule on the target after replication failover

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

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.

Replication job modification fails

Replication has a facility to have a multiple recovery point objective (RPO) report on the target side. The `Replication> 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 replication can fail with an ENOSPC error. In this case, any replication job modification command fails.

Workaround:

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

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 replication cluster is not successful.

Workaround:

There is no workaround for this issue.

Synchronous replication shows file system layout as mirrored in case of simple and striped file system

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

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

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

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

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

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

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.

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

When you destroy a file system, the `/etc/mstab` entry should be removed. If the file system `umount` command hangs during the destroy operation, the `/etc/mstab` 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.

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 exists 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
```

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

The `Storage> fs growto` and `Storage> fs growby` commands give an *"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.

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.

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 fencing on a setup with volume manager objects present fails to import the disk group

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 `vxldg -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.

```
# vxldisk -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

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

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> hagrps -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.

Storage> dedup commands are not supported on RHEL 7.x

The `Storage> dedup` commands are not supported in Veritas Access 7.3.0.1 on RHEL 7.x platform.

Workaround:

There is no workaround for this issue.

A scale-out file system may be in offline state after the execution of Storage> fencing off/on command

The `Storage> fencing off/on` command is executed to configure fencing on a Veritas Access cluster. The operation involves restarting of services. After this operation, you may find that some of the scale-out file systems fail to come online on their own. Their state is reported as offline when you run the `Storage fs> list` command.

Workaround:

The cluster can recover from this state. Use the `Support> service autofix` command to fix the file systems that are offline. This command clears all the faults and brings the file systems online.

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.

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 Command Reference 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*