Veritas Access Appliance Release Notes

8.0 Linux

May 2023

VERITAS™

Access Appliance Release Notes

Last updated: 2023-05-21

Legal Notice

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

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 latest documentation is available on the Veritas website:

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

Documentation feedback

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

APPL.docs@veritas.com

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

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 Access Appliance

This chapter includes the following topics:

About this release

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

Changes in this release

This section shows the major new features and enhancements added in the 8.0 version of Access Appliance.

Enhancements to cluster configuration workflow

The following enhancements are made to the cluster configuration workflow:

- **Specify flexible number of physical IP addresses:**
  You have the flexibility about the number of public data IPs that are required during configuration. You can specify any number between 0 to 4 (considering a 2-node cluster with 2 public data NICs). This feature is supported during initial configuration, while adding a node, and during upgrade scenarios.

- **Configure at least one NTP server for the cluster:**
  At least one NTP server must be configured for the cluster. Unlike in earlier releases, configuring NTP is not optional.
- **Configure a virtual IP address for a public network interface that does not have a physical IP address assigned to it:**
  A physical IP address need not be assigned to a public network interface before assigning a virtual IP address to it. Earlier you were required to assign a physical IP address to a public network interface if you wanted to assign a virtual IP address to it. This feature is supported during initial configuration and while assigning a virtual IP address from the command-line interface after the configuration is complete.
- **Set lockdown mode and retention period for increased security:**
  To protect cluster data from external threats and unauthorized access, during the configuration process, you can configure the lockdown mode for the cluster, normal being the default mode.
- **Create a default storage pool after the configuration:**
  A default storage pool is created automatically at the time of configuration. You no longer need to create a storage pool post configuration.

For more information, see the *Veritas Access Appliance Initial Configuration Guide*.

Support for immutability

Immutability ensures that the backup image is read-only and cannot be modified, corrupted, or encrypted after backup. You can use the lockdown modes to protect your cluster data from internal and external threats by securing all the external endpoints from unauthorized access.

For more information, see the *Access Appliance Administrator's Guide*.

Accessing the WORM storage server instances for management tasks

You can use the Access Appliance WORM storage server shell to configure immutable and indelible data from the appliance and perform some management tasks on the WORM storage server instance.

For more information, see the *Access Appliance Solutions Guide for NetBackup*.

Configuring MSDP-C with Access Appliance

You can now configure Access S3 with MSDP-C and duplicate images to the WORM S3 bucket.

For more information, see the *Access Appliance Solutions Guide for NetBackup*.

Configuring user authentication using digital certificates or smart cards

You can enable smart card authentication from the Access Appliance UI.

For more information, see the *Veritas Access Appliance Administrator's Guide*.

Managing password policies

You can customize the password policies by setting rules for password complexity, password age, and password lockout.

For more information, see the *Veritas Access Appliance Administrator's Guide*.

Setting login banners

You can set a text banner that is displayed when you log in to the appliance. You can use the login banner to communicate various kinds of messages to users.

For more information, see the *Veritas Access Appliance Administrator's Guide*.

New command structure for Access Appliance Shell commands

The 8.0 release introduces new structure for the Access Appliance Shell commands. All the Access Appliance commands that you could previously run from the Access Appliance Shell Menu are now replaced by the Veritas Appliance Shell commands (vxos commands) as part of an effort to create greater consistency across the Veritas appliance products.

For more information, see the *Veritas Access Appliance Command Reference Guide*.

Note: All the legacy technical notes and articles published on the Veritas Support site might not reflect this new command structure. Refer to the *Veritas Access Appliance Command Reference Guide* for the commands.

Support for new hardware model

A new Access 3350 Appliance model is introduced in the Access Appliance 8.0 release.

For more information about the model, see the *Veritas Access 3350 Appliance Product Description* and the *Veritas Access 3350 Appliance Hardware Installation Guide*.

Preupgrade check to determine if the appliance is ready for an upgrade

You can run an independent check before you start the upgrade to identify potential issues that can cause an upgrade failure.

For more information, see the *Veritas Access Appliance Upgrade Guide*.

Terminology changes

Starting with version 8.0, Veritas has begun to replace certain outdated terms. Primary server replaces master server in the UI and documentation.

Note: As Veritas continues to update its terminology, the deprecated term and the new term may be used interchangeably.

Veritas plans to update the following terms in future versions. The UI and documentation will be updated to reflect these changes.

- master
- slave
- whitelist or white list
- blacklist or black list
- whitehat
- blackhat

Supported NetBackup client versions

The following versions of NetBackup client are supported with Access Appliance 8.0.

- 8.1.2.1
- 8.2
- 9.0
- 9.0.0.1

- 9.1
- 10.0

The supported NetBackup clients can be installed as add-on packages.

Access Appliance simple storage service (S3) APIs

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

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

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

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

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

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

Chapter 2

Fixed issues

This chapter includes the following topics:

Fixed issues in this release

This section includes the issues fixed since the last release.

Table 2-1 Fixed issues since the last release

Fixed issue	Description
IA-7127	When an earlier version of the Access Appliance cluster is upgraded, the GUI shows stale and incomplete data
IA-32511	Deprecated policies are displayed after upgrading to version 8.0.
IA-32816	File system status is not updated in the UI.

Chapter 3

Software limitations

This chapter includes the following topics:

Limitations on using shared LUNs

The following limitations relate to shared LUNs in Access Appliance.

Access Appliance does not support thin LUNs

Access Appliance does not support thin LUNs. If thin LUNs are used, some commands may fail when run from the Access Appliance command-line interface.

Limitations related to installation and upgrade

The following limitations are related to installation and upgrade.

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

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

Add the required number of virtual IPs per service using the following command:

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

Underscore character is not supported for host names

Starting with 7.4.3, underscore (_) is not allowed in a host name. You can however upgrade to 7.4.3 from an earlier version where the host name included an underscore.

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.
```

Access Appliance IPv6 limitations

The following Access Appliance modules are not supported for IPv6:

- NIS

FTP limitations

The following limitation applies to FTP.

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

Limitations related to commands in a non-SSH environment

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

For example:

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

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

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

- `Upgrade> add`
- `Upgrade> install`

Limitations related to Veritas Data Deduplication

The following limitation is related to Veritas Data Deduplication.

- If you want to reconfigure Veritas Deduplication using previously used file systems, you have to use the same credentials that you used during the initial configuration.

Kernel-based NFS v4 limitations

The following limitations apply for kernel-based NFS v4:

- NFS v4 ACLs are not supported by Access Appliance.
- NFS v4 share reservations are not supported.
- NFS v4 delegation is not supported.

File system limitation

The following limitations relate to the Access Appliance file system.

- Any direct NLM operations from the Access Appliance command-line interface can lead to system instability
  Do not perform any file system related operations using the Access Appliance command-line interface on the Network Lock Manager (NLM), as it is used for internal purposes. If NLM is used, then Access Appliance cannot guarantee the stability of the cluster.
- When a file system is created, an additional file system is also created for the purpose of keeping the lock and configuration information. The additional file system is not directly accessible to the user. It is meant for internal use only.
  It is recommended to use disks from as many nodes as possible when creating the first storage pool in Access Appliance. In case of a shared nothing environment where the disks are local to the cluster nodes, the additional file system mirrors are created across all those nodes. This ensures that the Access Appliance configuration is available even if one of the nodes on which the additional file system was created is available.
  In case of SAN environments, the additional file system is mirrored across two disks.

- On-premises tiering in a cluster file system only supports one primary and one secondary.

Access Appliance S3 server limitation

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

The object has to be downloaded in parts.

Long-term data retention (LTR) limitations

The following limitations are related to LTR:

- Access Appliance does not support the HTTPS application protocol for an S3 bucket from the GUI in Veritas NetBackup long-term retention (LTR) use cases.

Limitations related to upgrade

The upgrade operation fails if you have configured host-based NetBackup client with Veritas Data Deduplication. Contact Veritas Support to unconfigure the host-based NetBackup client, perform an upgrade and configure container-based NetBackup client.

Limitation related to replication

The following issues relate to replication in Access Appliance.

Limitation related to episodic replication authentication

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

Limitation related to continuous replication

- Continuous replication does not support changing the mode of replication (synchronous or asynchronous) after replication is configured.
- The Access Appliance file system operations such as grow, shrink, resize, addition or removal of column, mirror, or tier (except cloud tier) are not supported for a file system which is configured under continuous replication.

Chapter 4

Known issues

This chapter includes the following topics:

Access Appliance known issues

The following known issues relate to the Access Appliance commands.

Admin issues

This section describes known issues related to the admin module.

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

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

(IA-12819)

Workaround:

There is no workaround for this issue.

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

CIFS issues

This section describes known issues related to CIFS.

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

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

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

(3853674)

Workaround:

To solve this issue, follow these steps:

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

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

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

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

(3824576, 3836861)

Workaround:

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

```
CIFS> share modify
```

Default CIFS share has owner other than root

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

(IA-4771)

Workaround:

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

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 command does not accept the special character '*'.

(IA-8108)

Workaround:

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

CIFS share may become unavailable when the CIFS server is in normal mode

When the CIFS server is running in normal mode and a CIFS share is brought online, if the MSDFS referral IP for that share faults and is not able to come online on any other node, the share becomes unavailable. In such cases, the CIFS share has the wrong definition in the smb.conf file and hence remains inaccessible. This issue can occur during an upgrade or when the NIC hosting the MSDFS IP of the CIFS share gets faulted on all nodes. This issue does not occur when the CIFS server is in the CTDB mode.

(IA-22339)

Workaround:

Stop and start the CIFS server.

CIFS share creation does not authenticate AD users

After cluster configuration, if the CVM of the slave node comes up before the management console, it results in an inconsistency in the `nsswitch.conf` file on all the cluster nodes. This affects the authentication of AD users while configuring CIFS shares.

(IA-21747)

Workaround:

1. Configure the `nsswitch` using CLISH.

   ```
   Network > nsswitch conf
   ```

2. Stop and restart CIFS service.

If you mount or access a CIFS share using the local user without netbios or cluster name, the operation fails

There is a behavior change in the latest SAMBA version. SAMBA checks for the domain name passed by the client while mounting the share. So if local users access or mount a CIFS share, the netbios name or cluster name along with user name, *<netbios_name>**<username>* is required. The netbios name cannot be empty.

(IA-31907)

Workaround:

Mount (or remount if you have upgraded from SAMBA 4.6.11) the CIFS share using the cluster name or netbios name.

During upgrade, the CVM and CFS agents are not stopped

When CIFS is configured in the cluster and you perform an upgrade, the upgrade process may fail to stop the running services. This occurs due to a resource dependency violation.

(IA-28386)

Workaround:

1. Run the following command on the cluster node only if the other node got evacuated during the upgrade process.

   ```
   /opt/VRTS/install/installaccess -start <node_ip>
   ```

2. Run the following command on the cluster:

   ```
   Support> service autofix
   ```

3. Stop the CIFS server.

```
CIFS> server stop
```

4. Restart the CIFS server.

```
CIFS> server start
```

Unable to access CIFS shares using the Share Open command from the Access Appliance Shell menu

```
- [Error] V-58000-500-1003: Unable to open network share.
An internal Samba error occurred. Restarting the appliance may
solve this issue.
```

When Active Directory (AD) is configured, CIFS security is set to **ads**. However, when AD is disabled, the security mode is not changed to **user** and SAMBA services cannot be restarted in the **ads** mode. You must change the security level to **user**.

IA-32560

Workaround:

1 Using SSH connect to the appliance node.

2 Use the Access command-line interface to check the AD status:

```
network> ad show
Name               Value
================== ========
Domain             -
netbios name       ltratest
workgroup          -
Domain Controller -
Domain user        -
Status             Disabled
```

3 If the AD status is Disabled, check the CIFS security using the following command:

```
cifs> server status
CIFS Status on access-3340-01 : OFFLINE
CIFS Status on access-3340-02 : OFFLINE

Homedirfs                 :
Security                  : ads
Domain membership status  : Disabled
Domain                    :
Workgroup                 : WORKGROUP
Domain Controller         :
Domain User               :
Clustering Mode           : normal
```

4 Set the security to user:

```
cifs> set security user
Global option updated. Note: Restart the CIFS server.
```

5 Stop/ and then start the CIFS server:

```
cifs> server stop
> cifs server start
```

6 Retry the command from the Access Appliance shell menu.

CIFS share may become inaccessible after an upgrade from Access Appliance version 7.4.2.400 to 8.0

After upgrade from Access Appliance version 7.4.2.400 to 8.0, some of the CIFS shares may become inaccessible. (IA-28194)

Workaround:

Restart the CIFS server using the following commands:

```
cifs> server stop

cifs> server start
```

Upgrade from Access Appliance version 7.4.2.400 to 8.0 fails if CIFS CTBD mode is configured

If CIFS CTBD clustering mode is configured on the appliance (with `security` = `ads`) and you perform an upgrade from Access Appliance version, the upgrade fails as the CTDB port does not get added correctly. (IA-40498)

Workaround:

After rollback, stop CTDB. Change the security to user and perform the upgrade again. After upgrade, configure AD and start the CIFS server.

Steps to perform before upgrade:

```
cifs> server stop
cifs> set security user
cifs> server start
cifs> server stop
```

Steps to perform after upgrade:

```
network> ad set
network> ad enable
cifs> server start
```

General issues

The following issue relates to all the Access Appliance modules.

Reimaging the appliance from the SSD device fails if a CD with the ISO image is inserted in the CD-ROM

Reimaging from the SSD device fails with the following message if a CD with the ISO image is also inserted in the CD-ROM of the appliance:

Pane is dead. Couldn't open file /tmp/ks_top_customization.cfg.

If you select SSD as the boot device and the CD is also inserted, the installer detects two ISO images with the same label, which results in a conflict.

(APPSOL-157945)

Workaround:

Ensure that the CD with the ISO image is not inserted in the CD-ROM when you select the SSD as boot device from the boot menu options. If the **Pane is dead** error message is displayed, remove the CD from the CD-ROM and try to reimage the appliance from SSD again.

A functionality of Access Appliance works from the master node but does not work from the slave node

This issue occurs when any of the operating system-specific configuration files are not as per the Access Appliance requirements. For example, if the `nsswitch.conf` file on a slave node is not the same as the file on the master node, the slave node does not follow the verification order for authentication. This causes authentication of users from the slave node to fail. This issue applies to all the protocols that are dependent on the `nsswitch.conf` such as, CIFS, NFS, FTP, and iSCSI.

IA-14735

Workaround:

Restart the slave node on which the functionality does not work.

The complete attribute list of adapters and RAIDs do not get displayed in the GET Rest API output

The output of the following Rest APIs do not display all the details of the attribute list:

```
GET /api/appliance/v1.0/hardware/nodes/{nodeName}/raids/{raidName}
GET /api/appliance/v1.0/hardware/nodes/{nodeName}/adapters/{adapterName}
```

(IA-41054)

Workaround:

There is no workaround for this issue.

User account gets locked on a management or non-management console node

If the user account is locked on a management console node because of multiple incorrect login attempts, both SSH and GUI sign-in fail on that node till the account lock period is complete.

If the user account is locked on a non-management console node because of multiple incorrect login attempts, SSH to that specific node is blocked. SSH to all the other nodes and sign-in to the GUI continues to work.

Workaround:

Account lock depends on the password policies and STIG rules. Wait for the lock period to get completed.

Setting retention on a directory path does not work from the Access Appliance command-line interface

If you set retention on any directory path using the `Storage> fs retention set <absolute directory path>` command from the Access Appliance command-line interface, the command fails.

For example :

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

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

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

Workaround:

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

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

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

GUI issues

The following issues relate to the GUI.

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

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

Workaround:

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

When provisioning storage, the Access web interface or the command-line interface displays storage capacity in MB, GB, TB, or PB

When you provision storage using the Access web interface or the command-line interface, the storage pool size and the size configured for Veritas Data Deduplication

are displayed in MB, GB, TB, or PB. However, the storage capacity is actually in MiB, GiB, TiB, or PiB.

(IA-15180)

The following section shows the available storage capacity for a Veritas 3340 Appliance using 10-TB and 4-TB disk capacities.

Table 4-1 With 10-TB drive capacity

	One storage shelf (one Primary shelf)	Two storage shelves (one Primary shelf; one Expansion shelf)	Three storage shelves (one Primary shelf; two Expansion shelves)	Four storage shelves (one Primary shelf; three Expansion shelves)
Maximum storage pool size (total capacity)	635 TiB	1.24 PiB	1.86 PiB	2.48 PiB
Maximum Veritas Data Deduplication pool size (available for application data)	631 TiB	1267 TiB	1267 TiB	1267 TiB

Table 4-2 With 4-TB drive capacity

	One storage shelf (one Primary shelf)	Two storage shelves (one Primary shelf; one Expansion shelf)	Three storage shelves (one Primary shelf; two Expansion shelves)	Four storage shelves (one Primary shelf; three Expansion shelves)
Maximum storage pool size (total capacity)	253.44 TiB	510 TiB	764 TiB	1018.6 TiB
Maximum Veritas Data Deduplication pool size (available for application data)	249 TiB	510 TiB	764 TiB	1018.6 TiB

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

When the external certificates are added to Access Appliance, a web server restart is implicitly performed to start the newly provided certificates. This implicit start of the web server does not work in RHEL 7 because the commands are different in RHEL 6 and RHEL 7.

(IA-9739)

Workaround:

Run the `System> guienable` command to start the server in the Access Appliance command-line interface.

Client certificate validation using OpenSSL ocsp does not work on RHEL 7

Client certificate validation is required for the two-factor authentication (2FA). The validation of certificates is successful in RHEL 6. In RHEL 7, an explicit parameter called *–VAfile* and the signer certificate are required to be passed, which does not happen. Hence, the client validation using the certificate does not work on RHEL 7.

Workaround:

There is no workaround for this issue.

GUI does not support segregated IPv6 addresses while creating CIFS shares using the Enterprise Vault policy

If you create CIFS share from the GUI using the Enterprise Vault policy, and you provide a virtual IP for IPv6, then it displays the IP as **sharename@ipv6**, which is not supported by Access Appliance.

Workaround:

Do not use virtual IPs when you create the CIFS share using the Enterprise Vault policy.

During a rolling upgrade the UI becomes inaccessible

After the upgrade is 51 percent complete, the nodes are restarted. After restarting the nodes, the UI is not accessible for some time. However the UI is accessible later after the upgrade is complete.

(IA-31518)

Workaround:

There is no workaround for this issue. After the upgrade is complete, wait until the UI is up and accessible.

REST endpoint field gives an error message for valid values while registering S3-compatible as a cloud service

If you use the GUI to register S3-compatible as a cloud service provider by navigating to **Settings** > **Cloud storage registration** > **Add** > **Select cloud provider** and select S3 COMPATIBLE, when you enter a valid REST endpoint, you get an error message. This happens because of incorrect field validation. (IA-33995)

Workaround:

Register S3-compatible as a cloud service using CLISH.

If lockdown mode was set using CLISH, switching the lockdown mode to Normal mode using the GUI fails if you set the retention period as 0

If you set the lockdown mode using CLISH, and if you try to change the lockdown mode to Normal mode using the GUI by setting the minimum and maximum retention period as 0, the operation fails. (IA-40824)

Workaround:

Do not set the maximum and minimum retention period as 0 in the GUI when switching the lockdown mode to Normal.

Infrastructure issues

The following sections describes the infrastructure-related issues.

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

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

(IA-9774)

Workaround:

After you add a node, log on to the Access Appliance command-line interface interface using the *admin* user credentials. Execute the `cluster reboot <new_node>` command to restart the newly added node.

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

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

Workaround:

Perform the following steps:

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

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

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

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

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

Workaround:

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

Installation and configuration issues

The following issues relate to Access Appliance installation and configuration.

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

Iptables are enabled by default on the Access Appliance 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 Access Appliance script, which is executed after Veritas Cluster Service (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.

(IA-1796)

Workaround:

For a fresh configuration of Access Appliance 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 server console.

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

1 After you complete adding a 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 server console.

Running individual Access Appliance scripts may return inconsistent return codes

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

(3796864)

Workaround:

There is no workaround.

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.

(3659716)

Workaround:

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

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

In a particular flow, if the node that is in the restart, shutdown, or crash state is running, the system calculates the running node list. It returns 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.

(IA-900)

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 the Access Appliance command-line interface by using the

`Network> pciexclusion add` command. If these unique PCI IDs already exist in the PCI exclusion configuration of Access Appliance, 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.

(IA-1850)

Workaround:

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

Phantomgroup for the VLAN device does not come online if you create another VLAN device from the Access Appliance command-line interface after cluster configuration is done

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

(IA-6671)

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 is ONLINE.

Configuring Access Appliance with a preconfigured VLAN and a preconfigured bond fails

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

(IA-11874)

Workaround:

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

In a mixed mode Access Appliance cluster, after the execution of the Cluster> add node command, one type of unused IP does not get assigned as a physical IP to public NICs

If you configure both IPv4 and IPv6 IPs as unused physical IPs, and execute the `Cluster> add <IPv4_ip of node_to_be_added>` command, then only IPv4 unused IPs are assigned as physical NICs. IPv6 IPs are not assigned to the newly added node.

(IA-13271)

Workaround:

Add the node to the cluster using the `Network> ip addr add` command. Then, manually configure the IPv6 physical IP.

NLMGroup service goes into a FAULTED state when the private IP (x.x.x.2) is not free

The Veritas installer assigns the 172.16.0.2 IP as the NLM master IP in the `/etc/VRTSvcs/conf/config/main.cf` file. It does not select it from the private IP range that is provided by the user. The installer does not check if the 172.16.0.2 IP is pingable or not. If that IP has already been used, the NLMGroup service goes into FAULTED state post-installation.

(IA-14577)

Workaround:

Execute the following commands on any node to replace the IP address on all the cluster nodes:

- **`haconf -makerw`**
- **`hares -modify nlmmasterIP IP_address`**

  Where *IP_address* is the free private IP.
- **`haconf -dump -makero`**

Clear the fault using the following command:

`hagrp -clear NLMGroup`

Bring the service group online on the desired node using the following command:

```
hagrp -online NLMGroup -sys target_node
```

Where *target_node* is the target node

The cluster> show command does not detect all the nodes of the cluster

During the Access Appliance configuration, all the cluster host names entries are not added in the `/etc/hosts` file. This issue occurs when a host name is a sub string of another host name in the cluster. When the `cluster> show` command is executed from the Access Appliance command-line interface, it does not detect all the nodes of the cluster.

(IA-14741)

Workaround:

Manually update the `/etc/hosts` file with cluster host name entries including private IPs mapped to the hosts. Make sure that before you perform the Access Appliance configuration, any host name is not a sub string of another host name in the cluster. For example, hostname001 and hostname00 should not coexist.

When you configure Access Appliance as an iSCSI target, the initiator authentication does not work

This happens because the default algorithm (MD5) used for authentication is disabled in RHEL7. (CXC-7150)

Workaround:

Use a node session algorithm other than MD5 such as SHA1, SHA256 on the iscsi initiator.

Internationalization (I18N) issues

This section describes known issues related to I18N.

The Access Appliance command-line interface prompt disappears when characters in a foreign language are present in a command

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

Workaround:

You can use any one of the following methods:

- Log out and log on to the Access Appliance command-line interface again.
- Press `Ctrl + C`.
- Set the locale to the intended non-English language. Start the Access Appliance command-line interface.
  The supported languages are Chinese, Japanese, and Korean.

MSDP-C issues

This section describes MSDP-C issues.

MSDP-C duplication job fails with OpenStorage WORM lock error after the file system is grown to100%

MSDP-C duplication job fails if you retry a duplication job after performing `fs grow` operation on the file system from the Access GUI/CLISH after the file system is already 100% full. (IA-39570)

Workaround:

If the bucket file system is already 100% full, grow the file system from Access GUI/CLISH. Restart the NetBackup's pdde services using the following commands from the NetBackup bash before starting backup again.

```
/usr/openv/pdde/pdconfigure/pdde stop
/usr/openv/pdde/pdconfigure/pdde start
```

Networking issues

This section describes known issues related to networking.

CVM service group goes into faulted state unexpectedly

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

(3793413)

Workaround:

Reboot the node on which this issue occurs.

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

In a mixed IPv4 and IPv6 set up, the IP balancing does not consider the 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.

(3616561)

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 netgroups, the NFS mount on the NFS client fails.

(3559219)

Workaround:

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

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

The IPs 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 and are 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.

(3596284)

Workaround:

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

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

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

(IA-8304)

Workaround:

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

```
# hastart
```

Unable to import the network module after an operating system upgrade

The Access Appliance 8.0 release supports the NIC name retention feature. You cannot import the network module if you perform an operating system upgrade.

(IA-11777)

Workaround:

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

Network load balancer does not get configured with IPv6

If you configured load balancer using the Access Appliance command-line interface with an IPv6 virtual IP, the load balancer configuration appears to be successful but does not balance the load in the background. This is because the load balancer is not supported with IPv6.

(IA-10977)

Workaround:

There is no workaround.

Unable to add an IPv6-default gateway on an IPv4-installed cluster

If you add an IPv6 address on an IPv4-installed cluster, and then add the default gateway, you get the following error on the Access Appliance command-line interface:

```
Route already exists
```

This error occurs if the IPv6 auto assignment feature is enabled on the node of the cluster.

(IA-12942)

Workaround:

You can disable IPv6 auto assignment by adding the following entries in `/etc/sysctl.con` for all the network interfaces of the nodes that are under the control of Access Appliance:

```
net.ipv6.conf.<network interface name>.autoconf=0
net.ipv6.conf.<network interface name>.accept_ra=0
net.ipv6.conf.<network interface name>.accept_ra_defrtr=0
```

Then restart the node of the cluster.

LDAP over SSL may not work in Access Appliance 8.0

When SSL is configured for the LDAP server, the users, groups, and the netgroups may not be listed. This occurs because an IP is used in place of a common name.

(IA-13320)

Workaround:

Use LDAP in non-SSL mode.

The network> swap command hangs if any node other than the console node is specified

The default value of the `nodename` parameter in the `network> swap` command is the console node. If you specify the name of any other node, the command is executed on the specified node through a remote procedure call. Before the swap operation is performed, the script prompts the user to answer a question and waits for the answer. But the remote procedure call does not take any inputs and the command hangs.

(IA-14635)

Workaround:

There is no workaround for this issue.

LDAP user fails to establish SSH connection with the cluster when FTP is configured

If FTP is configured, then SSH connection may not work for an LDAP user or a user of an LDAP group with master role. This happens because the default user home directory for such users changes during login. (IA-32010)

Workaround:

There is no workaround for this issue.

Unable to configure primary, backup DC server as the command does not allow specifying multiple DC servers

Multiple DC server are not supported in this release. If you try to configure primary, backup DC server, the command does not allow you to enter multiple DC servers. (IA-40806)

Workaround:

There is no workaround for this issue.

NFS issues

This section describes NFS issues.

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.

(IA-1002)

Workaround:

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

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 Access Appliance command-line interface commands.

(IA-2758)

Workaround

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

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.

(IA-5083)

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. The kernel KNFS v4 export mount for Netgroup does not work as expected.

(IA-6672)

Workaround:

Restart the NFS service.

When a file system goes into the FAULTED or OFFLINE state, the NFS share groups associated with the file system do not become offline on all the nodes

There is an online local soft dependency between the NFS share group and the vrts_vea_cfs_int_cfsmount groups due to which the status of the VCS share groups is displayed as online on some nodes even after the vrts_vea_cfs_int_cfsmount goes offline.

(IA-14597)

Workaround:

Bring the NFS share group online manually if it is in the FAULTED or OFFLINE state even after the file system (vrts_vea_cfs_int_cfsmount group) is brought online.

Add and delete NFS client operations fail from the GUI

If you try to add or delete an NFS client using the GUI, the operation fails. (IA-40622)

Workaround:

To add a client to the NFS share, use the following CLISH command:

```
nfs> share add {NewExport_options} {ExistingPath} {NewShareName}
{NewClient}
```

To delete client, use the following CLISH command. You have to provide the share name which was provided when the client was added.

```
nfs> share delete ${ShareName}
```

Multiple NFS shares created on a single file system from CLISH do not get listed in the Restful API output and GUI

If multiple shares are created on a specific file system, only the last created share gets displayed in the GUI and Restful API output. All the shares of that specific file system do not get displayed. (IA-40861)

Workaround:

Use the following command to see all the shares:

```
nfs> share show
```

ObjectAccess issues

This section describes ObjectAccess issues.

When trying to connect to the S3 server over SSLS3, the client application may give a warning

Access Appliance 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.

When trying to connect to the S3 server over SSLS3, the client application may give a warning:

```
SSL3_GET_SERVER_CERTIFICATE:certificate verify failed
```

(IA-5378)

Workaround:

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

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

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

(IA-10245)

Workaround:

Map the file system using the Access Appliance command-line interface for local users.

If you have upgraded to Access Appliance 8.0 from an earlier release, access to S3 server fails if the cluster name has uppercase letters

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

(IA-5628)

Workaround:

Use all lowercase letters to access the S3 server.

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

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.

(IA-5631)

Workaround:

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

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

Bucket creation may fail with time-out error

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

(IA-7432)

Workaround:

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

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

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

(IA-7368)

Workaround:

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

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

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

(IA-7407)

Workaround:

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

Self test failed for storage_s3test

There is not enough free space in the storage pool to create an S3 bucket. (APPCPE-5347)

Workaround:

Ensure that the storage pool used by the object access server has sufficient free storage to create a bucket as per the size and type specified by the **fs_size** and **fs_type** parameters. The S3 self test is executed during an upgrade if the object access service is enabled, and the upgrade might fail if no bucket could be created.

Replication issues

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

When running episodic replication and deduplication on the same cluster node, the episodic replication job fails in certain scenarios

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

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

(3804751)

Workaround:

There is no workaround.

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

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

(3822515)

Workaround:

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

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

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

3 Run the `Replication> episodic config auth` command on the source and the target.

4 Delete the job directory from the new source `/shared/replication/jobs # rm –rf jobname/`.

5 Create the job from the new source.

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

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

(3668957)

Workaround:

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

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

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

(IA-3290)

Workaround:

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

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

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

(IA-7597)

Workaround:

Destroy the job and create it again.

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

The functionality provided by the `bwlimit show` command when you use the Access Appliance command-line interface is not available in the GUI.

The `bwlimit show` command is not supported through the GUI.

(IA- 7295)

Workaround:

You can use the following command to set the bandwidth using the Access Appliance command-line interface:

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

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

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

```
SSL certificate error.
```

(3839319)

Workaround:

Follow these steps to solve this issue:

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

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

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

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

(3797560)

Workaround:

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

Episodic replication job modification fails

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

(IA-3356)

Workaround:

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

If a share is created in RW mode on the target file system for episodic replication, then it may result in there being different number of files and directories on the target file system compared to the source file system

This issues occurs because replication does not work as expected and the target file system is not in the same state as the source file system.

Workaround:

There is no workaround for this issue. It is recommended that you do not use the target file system for any IO operations and use it only for replication.

The promote operation may fail while performing episodic replication job failover/failback

While performing site failover/failback, sometimes the promote operation fails with the following error:

```
replication> episodic job failover force=yes job_2 src_to_tgt
Access Appliance replication ERROR V-493-10-1795 Job failover failed
[Performing final sync for job job_2 before promoting...
Failed to promote job job_2 on 10.221.55.195.].
Please verify job details and start job manually.
```

This issue occurs intermittently and as a result, the target file system becomes offline.

Workaround:

Perform the following tasks from CLISH to perform failover/failback.

- Bring the target file system online.
- Perform failover/failback again.

Discrepancy is observed in the outputs of replication episodic service status and replication episodic job stats <job_name> commands

You may see a discrepancy in the outputs of the `replication episodic service status` and `replication episodic job stats <job_name>` commands. It may happen that the service status is displayed as running but the job stats shows that the replication service is not running. (CXC-7336)

Workaround:

Run the `replication episodic service start` command to restart the replication service.

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

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

(IA-7357)

Workaround:

- On the target, set the following rule.

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

- Save the rule.

```
# service iptables save
```

- Restart the IP tables.

```
# service iptables restart
```

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

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

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

Primary Site Info:

Host name           10.10.2.70
RVG state           enabled for I/O

Secondary Site Info:

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

Workaround:

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

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

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

```
Replication> continuous status test_fs
Name                    value
```

```
====================  ======================
Replicated Data Set    rvg_test_fs
Replication Role       Primary
Replication link       link1

Primary Site Info:

Host name              10.10.2.70
RVG state              enabled for I/O

Secondary Site Info:

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

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

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

Workaround:

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

```
# iptables -F
```

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

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

Workaround:

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

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

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

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

Workaround:

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

```
# iptables -F
```

Unplanned failback fails if the source cluster goes down

If the source cluster goes down, unplanned failback fails with an error:

```
ACCESS Sync_rep ERROR V-493-10-91 Failback operation failed!
```

The following error appears in the logs:

```
Message from Host 10.221.62.53: VxVM VVR vradmin
ERROR V-5-52-449 Secondary rvg1 does not have an active Primary
```

The issue is seen if the virtual IP address used by the replication service is on a different node then the logowner node. This happens because continuous replication requires the IP and logowner to be on the same node.

Workaround:

1. Find the logowner node.

   ```
   # vxprint -Vl rvg_test | grep logowner
   logowner: clat743-01 (default)
   ```

   In this example, *clat743-01* is the logowner node.

2. Run the following command.on the Access CLISH:

   ```
   # /opt/VRTSnas/clish/bin/clish -u admin -c "network ip addr show"
        IP       Netmask/Prefix Device Node       Type     Status
   ```

```
------------ --------------- ------ -----      -------- ------
10.221.54.85 255.255.240.0   eth4   clat743-01 Physical
10.221.54.86 255.255.240.0   eth5   clat743-01 Physical
10.221.54.55 255.255.240.0   eth1   clat743-01 Physical
10.221.54.87 255.255.240.0   eth4   clat743-02 Physical
10.221.54.88 255.255.240.0   eth5   clat743-02 Physical
10.221.54.56 255.255.240.0   eth1   clat743-02 Physical
10.221.54.89 255.255.240.0   eth1   clat743-01 Virtual  ONLINE
                                                        (Con IP)
10.221.54.59 255.255.240.0   eth5   clat743-02 Virtual  ONLINE
                                                        (Continuous
                                                        Replication IP)
```

Observe the node on which the continuous replication IP resides. In this example, it is *clat743-02*.

3. Move the logowner to the node on which the continuous replication IP resides. Run the following command on the node where continuous replication IP is running.

   ```
   # vxrvg -g sfsdg set logowner=on rvg_test
   ```

   In the example, run the command on *clat743-02*.

4. Run the failback command.

   ```
   # replication continuous failback fs_name
   ```

Cloud tiering cannot be configured with continuous replication

If a file system has continuous replication enabled on it, it is not possible to configure cloud tiering. This happens because the label of the object that is set by the source cluster and the label expected by the destination cluster is different. Hence the validation of cloud storage label fails on the destination cluster and the data from cloud becomes inaccessible.

(IA-32769)

Workaround:

There is no workaround for this issue.

After continuous replication failover/failback operations, the virtual IPs in the source may appear offline

After continuous replication failover and failback operations, the virtual IPs in the source may appear offline. If you bring the virtual IPs online using the `network ip addr online` command, the command fails with the following error:

```
ACCESS ip addr ERROR V-493-10-1479 Node <node name> is either still
starting or does not have all the filesystems mounted

ACCESS ip addr ERROR V-493-10-1394 ip addr online command failed
```

(IA-32565)

Workaround:

1. List all the file system service groups that are in OFFLINE/FAULTED/PARTIAL state.

   ```
   hagrp -state | grep vrts_vea_cfs_int_cfsmount
   ```

2. Run the following command for the file system service groups that are not online to identify the empty file system service groups.

   ```
   hagrp -resources vrts_vea_cfs_int_cfsmount<integer>
   ```

   where *<integer>* can be any positive number.

3. If the above command shows empty output then run the following command:

   ```
   hagrp -dep vrts_vea_cfs_int_cfsmount<integer>
   ```

4. If there are dependency links, run the following command:

   ```
   /home/maintenance # hagrp -dep vrts_vea_cfs_int_cfsmount<integer>
   ```

   For example:

   ```
   /home/maintenance # hagrp -dep vrts_vea_cfs_int_cfsmount3
   #Parent Child Relationship
   vrts_vea_cfs_int_cfsmount3 RVGgroup_rvg_test online local firm
   vrts_vea_cfs_int_cfsmount3 cvm online local firm
   ```

5. Delete the dependencies.

   ```
   hagrp -unlink vrts_vea_cfs_int_cfsmount<integer> RVGgroup_rvg_test
   hagrp -unlink vrts_vea_cfs_int_cfsmount<integer> cvm
   ```

6. Delete the groups using the following command:

   ```
   hagrp -delete vrts_vea_cfs_int_cfsmount<integer>
   ```

7. Check if the phantom groups are offline:

   ```
   hagrp -state | grep Phantomgroup_pubeth
   ```

8. If there are any offline groups then run the following command for all the groups that are offline:

   ```
   hagrp -online Phantomgroup_pubeth2 -any
   ```

9. Verify that all the Phantomgroups are online.

   ```
   hagrp -state | grep Phantomgroup_pubeth
   ```

   The virtual IPs should now appear online.

Cannot use a file system to create an RVG if it has previously been enabled as the first file system in an RVG and later disabled

If a file system is enabled as the first file system in an RVG (say rvg1), and it is later disabled, then you cannot use the same file system to create another RVG. It can only be added back to the previous RVG (rvg1) or another existing RVG. (IA-40082)

Workaround:

There is no workaround for this issue.

STIG issues

This section describes known issues related to STIG.

When the STIG option is enabled, it enforces an account lockout for any user that enters three consecutive incorrect passwords

When the STIG option is enabled, the user account is locked after three consecutive failed log in attempts. The account is locked for seven days.

(IA-29667)

To unlock the user account:

1. Log in to the system as an appliance administrator and elevate to root.
2. Check the status using the following command:

   ```
   faillock --user username
   ```

   To view all accounts that are locked, use the `faillock` command.
3. To unlock the user account, use the following command:

   ```
   faillock --user <username> --reset
   ```

The changed password is not synchronized across the cluster

When the STIG option is enabled, the user password expires as per the set password policy. After the password expires, the user is prompted to change the password if the user logs on to the system via SSH. If the user changes the password at the prompt, the password is changed only locally and is not synchronized across the cluster.

(IA-29565)

Workaround:

After the password expires, when prompted, do not change the password at the OS prompt. Instead, log on to the GUI with your credentials and change the password from the GUI.

STIG status is not preserved if the system configuration is imported after reimaging the nodes

If the STIG option is enabled on the cluster and the system configuration is exported to a remote location using the `system config export remote URL` command before reimaging the nodes, the STIG configuration is not preserved after importing the configuration. When you reimage the nodes while preserving the storage, and then import the configuration to restore the settings, the STIG option for the nodes is displayed as **Disabled** when you run the `system> security stig show` command.

(IA-31745)

Workaround:

Enable the STIG option again by running the `system> security stig enable` command.

Storage issues

The following issues relate to the Access Appliance Storage commands.

Snapshot mount can fail if the snapshot quota is set

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

(IA-1542)

Workaround:

Remove the oldest checkpoint and retry.

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

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

(IA-1733)

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.

(IA-1639)

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.

(IA-1851)

Workaround:

Run the following commands from the Access Appliance command-line interface:

```
Storage> quota cifshomedir disable

Storage> quota cifshomedir enable
```

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

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

(3801216)

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.

(3650635)

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 Access Appliance command line when you create a file system, Access Appliance might create a data change object (DCO) on disks other than those specified. If free disks are available in the pool, Access Appliance 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
```

(3452098)

Workaround:

There is no workaround.

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

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

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

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

(3588248)

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.

(IA-3269)

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 Access Appliance cluster being in an inconsistent state.

(IA-3427)

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.

(IA-3091)

Workaround:

Run the `Storage> tier move` command again from the Access Appliance command-line interface.

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. You cannot handle this case from the Access Appliance command-line interface. Manual intervention by Veritas Technical Support is required to preserve the rollback.

(IA-3251)

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.

(IA-3239)

Workaround:

There is no workaround.

The Access Appliance command-line interface does not block uncompress and compress operations from running on the same file at the same time

The Access Appliance command line 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.

(IA-2995)

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> 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.

(IA-3241)

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.

(IA-3432)

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, Access Appliance 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.

(IA-3398)

Workaround:

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

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

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

(IA-5434)

Workaround:

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

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

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

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

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

(IA-4061)

Workaround:

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

Unable to create space-optimized rollback when tiering is present

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

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

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

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

(IA-5690)

Workaround:

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

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

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

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

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

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

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

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

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

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

(IA-7219)

Workaround:

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

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

You can now enable I/O fencing.

File system creation fails when the pool contains only one disk

When there is only one disk in pool, the `fs creation` command fails to create an NLM on the file system. Instead, it tries to create the file system with different options. This happens because NLM requires two disks as it creates a mirrored volume/file system.

(IA-7515)

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.

(IA-7174)

Workaround:

Clear the fault using the following command:

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

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.

(3931869)

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 Access Appliance command-line interface.

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

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

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

Workaround:

To bring the CVM service group online

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

```
# hastop -local
```

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

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

3 Deport all the disk groups using the following command:

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

4 Import all the disk groups using the following command:

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

5 Start Veritas Cluster Server (VCS).

```
# hastart
```

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

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

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

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

(IA-10709)

Workaround:

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

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

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

(IA-11831)

Workaround:

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

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

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

(IA-11893)

Workaround:

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

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

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

(IA-11486)

Workaround:

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

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

1. Install Access Appliance without enabling fencing.
2. Execute the `vxdctl scsi3_pr off` on all the nodes.
3. From the Access Appliance command-line interface, execute the `Cluster> reboot all`.

4 After the system restart, execute the `Storage> fencing on majority` from the Access Appliance command-line interface.

5 Create the pool and the file system.

To enable fencing with non-SCSI3 disks without a restart

1 Stop the cluster services.

```
# hastop -all
```

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

```
# vxdctl scsi3pr off
```

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

4 Restart `vxconfigd` on all the nodes of the cluster.

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

5 Start all the nodes of the cluster.

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

Wait for the disk group to get imported.

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

```
#hastart
```

Now, you can enable fencing.

Error while creating a file system stating that the CVM master and management console are not on the same node

When you create a file system and restart the node, the CVM master and management console may get inconsistent and may not be present on the same node. You get an error message stating that the CVM master and management console are not on the same node.

(IA-14727)

Workaround:

From the Access Appliance command-line interface, run the `storage scanbus` command to bring the CVM master and management console on the same node.

When you configure disk-based fencing, the cluster does not come online after you restart the node

When disk-based fencing is configured, sometimes a node may panic or may get reset in an unclean way, and get stuck. If the other nodes in the cluster restart at this time, they detect a split-brain condition. This happens because the nodes that have restarted see the SCSI reservation keys of the stuck node on the fencing disks and cannot determine whether the stuck node is actually stuck or is inaccessible on the network. The cluster does not come online and the following error message is displayed in the `syslog` file:

```
Preexisting split-brain. Dropping out of cluster.
```

(IA-14752)

Workaround:

Restart the stuck node to be to bring the cluster online.

Refer to the user documentation for steps required to clear preexisting split-brain.

After a node is restarted, the vxdclid process may generate core dump

After a node is restarted, the vxdclid process may generate core dump with the following stack trace:

```
(gdb) bt
#0 0x00007f31a2cec248 in getDgDisks (dgname=0x7f318c06459d "sfsdg", vect=0x7f
#1 0x00007f31a2d37f85 in doVmNotify (a=0x0) at vxbridge/common/vxlist_sf_noti
#2 0x00007f31aa986e25 in start_thread () from /lib64/libpthread.so.0
#3 0x00007f31aa6b434d in clone () from /lib64/libc.so.6
(gdb)
```

(IA-14534)

Workaround:

Execute the following commands on the node in which you see the stack trace to bring up the vxdclid process:

```
/opt/VRTSsfmh/adm/dcliunsetup.sh
/opt/VRTSsfmh/adm/dclisetup.sh
```

The cluster> shutdown command does not shut down the node

When the `cluster> shutdown` command is run on a node from the Access Appliance command-line interface, as part of the shutdown process, all the file system groups are made offline on the node before the node is shut down. If any file system groups remain online, the `cluster> shutdown` command is not executed on the node. The `cluster> shutdown` command hangs. The `/opt/VRTSnas/log/shutdown_output.log` displays the following message:

```
VxVM vxclustadm ERROR V-5-1-9360 waiting for applications to end
```

(IA-14890)

Workaround:

Use the `halt` command to shut down the node.

Audit logging of WORM-enabled file systems do not get enabled if the file systems were offline during upgrade

If any of the WORM-enabled file systems remain in offline state during upgrade, audit logging does not get enabled on those file systems. (IA-40434)

Workaround:

Perform the following steps:

1. After upgrade is completed, check if any file system is offline.

   ```
   storage> fs list
   ```

2. Bring online all file systems which were offline during upgrade.
3. Check if any of the file systems which were offline are WORM-enabled.

   ```
   storage> fs list <fs_name>
   ```

4. For each WORM-enabled file system, perform the following steps:
   - Perform `mount -t vxfs`, and identify raw device path and mount point of each file system.
   - Check if the audit logging is already enabled on the file system.

     ```
     # mkfs.vxfs -m <raw_device_path>
     ```

     If audit logging is enabled skip the next steps and go to the next WORM-enabled file system.
     For example:

```
# mkfs.vxfs -m /dev/vx/dsk/sfsdg/wrmfs
mkfs -t vxfs -o bsize=8192,version=17,inosize=256,logsize=2048,
rcqsize=8192,largefiles,maxlink,worm,aulog,nomaxts
/dev/vx/dsk/sfsdg/wrmfs 8388608
```

aulog means audit logging is enabled on the file system

- If audit logging is not enabled, run the following command:

  /opt/VRTS/bin/fsadm [-t vxfs] [-o aulog] path_of_mountpoint

- Confirm if audit logging is enabled using the following command:

  # **mkfs.vxfs -m <*raw_device_path*>**

System issues

The following issues relate to the Access Appliance 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.

(IA-8725)

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.

Access services are up and running if system is restarted after manually stopping services on both the nodes

IA-40392

If Access services are stopped on both the nodes using the `Cluster> stop all` command from the Access command-line interface and if the nodes are restarted, the Access services do not remain offline as expected but are up and running. The services should remain offline and be running only if the services are restarted manually using the `Cluster> start` command.

Workaround:

After restarting the nodes, stop the services again using the `Cluster> stop all` command.

Phantom service group remains offline if a cluster node is restarted.

After running the `cluster> reboot` *`node-name`* command, the node is restarted and all the services are brought online. But if the system load is high, there might be a race condition between CFSMount and Phantom service groups, which can lead to Phantom group remaining in OFFLINE state after the node is restarted.

(IA-38716)

Workaround:

Run the `Support> services autofix` command to resolve the issue.

Upgrade issues

This section describes known issues related to upgrade.

Unable to roll back the system after an attempt to upgrade from 7.4.2.400 to 7.4.3.300 failed.

On 7.4.2.400, file system deduplication was enabled for some of the file systems. When file system deduplication is enabled, the dedup_schedd_group service group is created. However, when the file system deduplication is disabled, the group is not deleted. File system deduplication is no longer supported on 7.4.3.x. During the upgrade, the group could not be brought online, causing the upgrade and roll back to fail.

(APPCPE-5183)

Workaround:

Delete the dedup_schedd_group service group before upgrading.

During the Access Appliance upgrade, I/O gets paused with an error message

When you upgrade the appliance, if the nmdb service stops running on a node, the I/O goes into PAUSED state and the following error message is displayed:

```
Invalid username or password
```

On all the other node, the shares are accessible.

(IA-14892)

Workaround:

Run the `cifs> server status` command. If the server is in online state and still I/O is paused, then restart the CIFS server from the Access Appliance command-line interface.

```
cifs> server stop
```

```
cifs> server start
```

The required configurations get updated and I/O processes are resumed.

During rolling upgrade, Access Appliance shutdown does not complete successfully

The Access Appliance shutdown operation does not complete successfully during rolling upgrade and the following error message is displayed:

`vxfs failed to stop on <node-name>`

(IA-14910)

Workaround:

1. When this issue occurs, the installer displays the following prompt:

   ```
   Do you want to continue? [y,n,q] (n) y
   ```

   Enter **y** and continue with the rolling upgrade.

2. The installer continues with phase 1 of rolling upgrade on the remaining nodes and exits before performing phase 2 of the upgrade on the cluster with the following message:

   ```
   It is recommended to perform rolling upgrade phase 2 on all the
   cluster systems in the next step. Rerun the installer to do this
   after reboot. It is strongly recommended to reboot the following
   systems:<node-name>.
   ```

3. Restart the node before phase 2 of rolling upgrade. Verify that the node is up and has joined the cluster by executing the following command on the master node:

   ```
   # vxclustadm nidmap
   ```

4. Verify that all the recovery tasks are complete by executing the following command on the master node:

   ```
   # vxtask list
   ```

5. Check if the following keywords are present in the `vxtask list` command output:

   ```
   ECREBUILD/ATCOPY/ATCPY/PLXATT/VXRECOVER/RESYNC/RECOV
   ```

   If the keywords are not present, start the upgrade again using the following command:

   `./installaccess -rolling_upgrade`

6. The installer asks if you want to proceed with phase 2 of rolling upgrade. Enter **y** to continue.
7. If the issues persist after restart, contact Technical Support.

CVM is in FAULTED state after you perform a rolling upgrade

After you perform a rolling upgrade, the CVM is in FAULTED as the `vxglm` module is not able to load. This occurs because of improper linking of the `vxglm` module. The status of the module is displayed as:

```
systemctl status vxglm.service
systemd[1]: Starting Systemd Veritas GLM service...
vxglm[18415]: Starting GLM...
vxglm[18415]: modprobe: FATAL: Module vxglm not found.
vxglm[18415]: ERROR: modprobe error for vxglm. See documentation.
```

(IA-14702)

Workaround:

1. Link the `vxglm` upgraded kernel module using the following command:

   `ln -sf /etc/vx/kernel/vxglm.ko.`*`<module_versoin>`*
   `/lib/modules/`*`<module_version>`*`/veritas/vxglm/vxglm.ko`

   You can get the latest version using the following command

   `rpm -ql `*`<VRTSglm_installed_pkg>`*` | grep vxglm.ko`

2. Execute the `depmod` command.
3. Restart the node on which CVM was in FAULTED state.

If rolling upgrade is performed when NFS v4 is configured using NFS lease, the system may hang

When you perform a rolling upgrade and if NFS v4 is configured using NFS lease, the system may hang with the following message:

```
BUG: soft lockup - CPU#5 stuck for 22s! [vx_glmlist_thre:18580]
```

The stack trace in the kernel log has the following information:

```
queued_spin_lock_slowpath
_raw_spin_lock
__break_lease
wake_up_atomic_t
vx_hlock_putdata
vx_glm_cbfunc
vx_glmlist_thread
vx_glm_cbfunc
vx_osdep_deinit
vx_kthread_init
kthread
insert_kthread_work
ret_from_fork_nospec_begin
insert_kthread_work
```

NFS v4 protocol uses lease per file. This delegation can be taken in read or write mode and can be released conditionally. For CFS, the delegation is released from a specific node while the inode (index node) is being normalized. This can lead to a race condition with another set lease operation on this node and may end in a deadlock. This causes the system to hang.

(IA-28572)

Workaround:

Restart the node which is in hang state.

Stale file handle error is displayed during rolling upgrade

When upgrading from an earlier version, read and write operations to the node where NFS shares are mounted might fail. Some of the NFS shares might be unmounted from the system during the rolling upgrade and the applications accessing these shares might face read and write issues. After the upgrade is complete, the NFS shares are automatically mounted and the file writes are started automatically.

(IA-27274)

Workaround:

There is no workaround for this issue.

The upgrade operation fails if synchronous replication is configured

If synchronous replication is configured and you perform an upgrade, the synchronous replication service groups try to stop the vradmin service in order to go offline while some other procedures try to start the vradmin service. The synchronous replication service groups are not able to go offline and this causes the upgrade operation to hang.

(IA-28466)

Workaround:

1. Login to Access CLISH either using the console IP address (ssh admin@<consoleIP>) or by using the `/opt/VRTSnas/clish/bin/clish -u admin` command.
2. Execute the following command before starting the upgrade:

   ```
   replication> continuous service stop
   ```

3. Start upgrade and wait for the upgrade operation to complete.
4. After the upgrade is complete, login to the Access CLISH and execute the following command:

   ```
   replication> continuous service start
   ```

Rolling upgrade fails when the cluster has space-optimized rollback in online state

Space-optimized rollbacks are always created with a simple file system layout irrespective of the underlying file system layout. Hence, there is a possibility that the CVM/CFS service group may go into a faulted state during rolling upgrade. If the CVM/CFS service group are in faulted state, then the rolling upgrade fails.

(IA-28502)

Workaround:

Bring all the space-optimized rollbacks to offline state before performing rolling upgrade. To list the rollbacks:

```
Storage> rollback list
```

To bring all the rollbacks to offline state:

```
Storage> offline <rollback_name>
```

After upgrading from version 7.4.2 to a later version, the default route entry in the ip rule table on one of the cluster nodes is missing

After upgrading, the default route entry is not present in the ip rule table on one of the cluster nodes. After installing the patch, you need to wait until the reboot is complete to view the ip rules in the table. Some of the cluster nodes become accessible after installing the patch. However, the gateway setting is deleted on one of the cluster nodes because of which the node becomes inaccessible.

(IA-29632)

Workaround:

1. Run the below script on the management console node:

   ```
   /opt/VRTSnas/scripts/lib/rpm_install_post.sh 1
   ```

2. Run the following commands on the node where problem is observed:

   ```
   # hastop -local
   ```

   ```
   # hastart
   ```

Note: After completing the above steps if the Access Storage Unit (STU) is not reachable from NetBackup, restart the deduplication services using the dedupe stop and dedupe start commands.

GUI might fail to start after upgrading from version 7.4.2 to 8.0

When you upgrade an appliance with bonded network interfaces, the GUI may fail to open and display only the flashing Access Appliance logo because the public key is not generated correctly in `cert.pub` file. The following error is displayed in the `isagui_webserver.log`:

```
UnauthorizedError: invalid signature
```

(IA-32619)

Workaround:

1 SSH to the appliance node using the admin credentials.

2 Elevate to maintenance on the node where the management console is online.

3 Back up the `cert.pub` file from the `/var/opt/VRTSnas/sslcerts/cert.pub` location.

4 Run the following command:

```
openssl rsa -in /var/opt/VRTSnas/sslcerts/key.pem -pubout >
/var/opt/VRTSnas/sslcerts/cert.pub
```

5 From the Access Appliance command-line interface run the following commands:

```
system guidisable
```

```
system guienable
```

Storage provisioning might fail after upgrading the Access Appliance from version 7.4.2 to 8.0

After upgrading from 7.4.2 to 7.4.3, storage provisioning for NFS, CIFS, Enterprise Vault Archival, and S3 buckets may fail intermittently.

IA-32599

Workaround:

Restart the following services manually on both the nodes:

```
/usr/bin/systemctl restart sds.service
```

```
/usr/bin/systemctl restart sds-access.service
```

Upgrade may fail if operations such as such as OS reboot, cluster restart, and node stop and shutdown are used during the upgrade

Performing any of the following operations during an upgrade can lead to an upgrade failure and the cluster might go in an inconsistent state. It is recommended that these operations should not be performed until the upgrade completes successfully. (IA-40377)

- Logging into maintenance mode and executing OS commands such as reboot and shutdown.
- Logging into a node and executing commands such as cluster reboot, shutdown, and stop from the command-line interface.
- Executing node maintenance operations from the UI such as shutdown node and start node.
- Operations that will disrupt network connections such as unplugging the network cable.

Workaround:

There is no workaround for this issue.

A wrong upgrade status might be displayed while upgrading from version 7.4.2. to a later version.

The **Software>UpgradeStatus** command might wrongly display the following message:

This appliance node is not in upgrade state.

(APPSOL-153523)

Workaround:

You can ignore the message and run the command again.

Disk layout version (DLV) of file systems are not upgraded if the file systems were offline during the upgrade

During an upgrade, if any of the file systems remain in offline state, the DLV version of the file system is not upgraded to the latest DLV 17 version.

(IA-40432)

Workaround:

1. After the upgrade is complete, bring the file systems that were offline during the upgrade to the online state by logging in to maintenance mode.
2. After the file systems are online, run the following command to check the raw device path of the file systems:

   ```
   mount -t vxfs
   ```

3. For each file system that was in offline state, perform the following operations:
   - Run following command to get current DLV version of the file system:

     ```
     mkfs.vxfs -m raw_dev_path_of_fs
     ```

     For example:

     ```
     mkfs.vxfs -m /dev/vx/dsk/sfsdg/vdd_fs
     mkfs -t vxfs -o bsize=8192,version=17,inosize=256,logsize=2048,
     rcqsize=8192,largefiles, maxlink,nomaxts
     /dev/vx/dsk/sfsdg/vdd_fs 8388608
     ```

   - If the DLV version is not set to 17, upgrade the DLV version in increasing order of 1 from the current DLV version to 17.

```
# /opt/VRTS/bin/vxupgrade UX:vxfs vxupgrade: INFO: V-3-22568:
usage: vxupgrade [-n new_version] [-r rawdev] mount_point #
/opt/VRTS/bin/vxupgrade -n current_dlv + 1 -r raw_dev_path
mount_point
```

For example, if the current DLV version is 12, run the above command with **new_version** set to 13,14,15,16, and 17.

Failed to retrieve the current password policy after upgrading to version 8.0

After successfully upgrading to version 8.0, the `system>password-policy get` command fails to get the current password policy because the `access_password_rules.conf` file does not exist in the `shared/security` directory. (IA-40097)

Workaround:

After the upgrade is complete, go to the Access command-line interface and run the `system>password-policy get` command to check if the password policy is set. If the command is successful, no further action is required. However, if the command displays the **Failed to get the system password policy** error message, run the `system>password-policy set` command with the default values to create the `access_password_rules.conf` file in the `/shared/security` directory:

```
access-clus> system password-policy set minlen=8 ucredit=1
maxclassrepeat=vxdefault dcredit=1 ocredit=1 minclass=vxdefault
lcredit=1 maxrepeat=vxdefault difok=vxdefault passsystem
password-policy set minlen=8 ucredit=1 maxclassrepeat=vxdefault
dcredit=1 ocredit=1 minclass=vxdefault lcredit=1 maxrepeat=vxdefault
 difok=vxdefault pass_min_days=vxdefault pass_max_days=vxdefault
pass_warn_age=vxdefault remember=vxdefault deny=vxdefault
unlock_time=vxdefault fail_interval=vxdefault
```

Upgrade from version 7.4.3.200 to 8.0 fails in post-upgrade self-test

Before starting the upgrade, the objectaccess service was online but an S3 bucket was not created. At the preupgrade stage, the storage pool was created with four disks but the number of columns mentioned during objectaccess configuration were five. As the bucket was not present and objectaccess service was online, the self-test tried to create a bucket but failed as the storage pool did not have the required number of disks.

(IA-41052)

Workaround:

Resolve the objectaccess bucket creation failure by adding a sufficient number of disks and then complete the following steps:

1. Login to Veritas Appliance shell. Elevate to bash with the `support elevate` command. To elevate to bash when the enterprise or compliance mode is enabled, refer to the Accessing the root shell in lockdown mode section of the *Veritas Access Appliance Administrator's Guide*.
2. Run the following command:
   - Ensure that the post upgrade self-test passed:

     ```
     hacli -cmd "python
     /opt/veritas/appliance/selftest/scripts/selftest.py -c
     /inst/patch/appliance/installed/8.0/scripts/selftest_access_cluster.configure
      postupgrade"
     ```

   - Run the following scripts, where *node_name_1* and *node_name_2* are the appliance node names:
     - `hacli -cmd "/inst/patch/appliance/installed/8.0/scripts/hooks/upgrade_is_successful/0030_node_installation_version_after_upgrade.sh -m rolling -n node_name_1 node_name_2"`
     - `/bin/bash /inst/patch/appliance/installed/8.0/scripts/hooks/upgrade_is_successful/0400_cluster_default_os_password_check.sh -m rolling -n node_name_1 node_name_2`
     - `/bin/bash /inst/patch/appliance/installed/8.0/scripts/hooks/upgrade_is_successful/0410_cluster_default_ipmi_password_check.sh -m rolling -n node_name_1 node_name_2`
   - Run the following command:

     ```
     hacli -cmd "/bin/rpm -e --allmatches VRTSaccess-app-update"
     ```

   - Run the following command:

     ```
     hacli -cmd "/bin/rm -rf /log/upgrade/upgrading/upgrade.flag"
     ```

Veritas Data Deduplication issues

This section describes known issues related to Veritas Data Deduplication.

The Veritas Data Deduplication storage server does not come online on a newly added node in the cluster if the node was offline when you configured deduplication

If a node is not online when you configure deduplication, the configuration file present in the node is not in sync with the cluster configuration. The soft link created during configuration is also not present in the node. When the node is made online, it does not get the required configuration details for starting the service.

(IA-14708)

Workaround:

1. Copy the `/opt/VRTSnas/conf/dedupe.yml` file from the master node to the newly added node.
2. Create a soft link, `/etc/pdregistry.cfg`, that points to the `/vx/fs_name/dedupe/etc/pdregistry.cfg` configuration file by using the following command on the newly added node:

   ```
   ln -sf /vx/fs_name/dedupe/etc/pdregistry.cfg /etc/pdregistry.cfg
   ```

   Where *fs_name* is the name of the file system provided during configuration.

The Veritas Data Deduplication server goes offline after destroying the bond interface on which the deduplication IP was online

When you configure Veritas Data Deduplication using a virtual IP configured on a bond NIC, Access Appliance creates a dependency between the virtual IP group and the phantom group as part of the bond creation. When the bond NIC is destroyed, the virtual IPs configured on the NIC are moved to the individual NICs using which the bond was created, but Access Appliance fails to delete the dependency between the virtual IP and the phantom group. As a result, the virtual IP group does not come online and the MSDP resource also remains offline.

(IA- 14536)

Workaround:

1. Log on to the shell.
2. Find the virtual IP group which contains the MSDP resource using the following command:

   ```
   hares -display MSDPRes | grep Group
   ```

3. Find the groups on which the virtual IP group is dependent using the following command:

```
hagrp -dep <virtual ip group>
```

4. Multiple phantom group dependencies get listed. Select the entry which contains the name, `bond`.
5. Delete the dependency between the virtual IP group and phantom group using the following command:

```
hares -unlink <virtual IP group> <phantom group>
```

If you grow the deduplication pool using the fs> grow command, and then try to grow it further using the dedupe> grow command, the dedupe> grow command fails

You can grow the deduplication pool using either the `fs> grow` command or the `dedupe> grow` command. But, if you initially grow the deduplication pool using the `fs> grow` command, and then try to grow it further using the `dedupe> grow` command, the `dedupe> grow` command fails.

(IA-14807)

Workaround:

Use only the `dedupe> grow` command to grow the deduplication pool.

The Veritas Data Deduplication server goes offline after bond creation using the interface of the deduplication IP

After you create a bond using the interface of the deduplication IP, the VCS group goes offline causing the virtual IPs and the Veritas Data Deduplication server to go offline.

(IA-18709)

Workaround:

Bring the VCS group online manually using the following command:

```
<group> = hares -value MSDPRes Group

hagrp -online <group> -any
```

Provisioning for Veritas Data Deduplication is displayed as failed in GUI

If there are any subtasks which are in running state, the provision for Veritas Data Deduplication fails.

(IA-23055)

Workaround:

Check the task details to make sure that no subtask is in running state and check the reason for the failure. Wait for all the subtasks to get completed before retrying the operation.

During reconfiguration of Veritas Data Deduplication with WORM, the specified username and password are not considered

During initial configuration of Veritas Data Deduplication with WORM, the username and password are used to configure the storage server. During reconfiguration, the username and password are taken as input from the user, but since the storage server was already configured, the credentials are not utilized anywhere, and the services are started.

Hence, the username and password provided during reconfiguration are not validated against those provided during initial configuration, and are not stored anywhere.

Workaround:

This behavior is by design and there is no workaround for this issue.

WORM-enabled MSDP does not start after a switch or restart of the deduplication engine

The sshd daemon does not start within the container used by the MSDP engine.

(IA-41835)

Workaround:

To resolve the issue, contact Veritas Support and refer them to article 100053136.

Access Appliance operational notes

This section contains the topics that explain important aspects of Veritas Access Appliance 8.0 operations that may not be documented elsewhere in the documentation.

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

Access services do not restart properly after storage shelf restart

If the Veritas Access 3340 Appliance loses connectivity to an attached Primary or Expansion storage shelf, the underlying storage connectivity is lost and the VxVM disk group goes into a deported state. This issue occurs whenever a storage shelf intentionally or unintentionally restarts. To correct this issue, you need to restart the Access services.

To restart the Access services after the appliance storage shelves restart

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

```
ltrcluster> storage scanbus
```

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

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

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

Chapter 5

Getting help

This chapter includes the following topics:

Displaying the Online Help

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

Displaying the man pages

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

Access Appliance 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 Access Appliance product documentation

The latest version of the Access Appliance 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 for Access Appliance on the SORT site:

- *Access Appliance Administrator's Guide*
- *Access Appliance Cloud Storage Tiering Solutions Guide*
- *Veritas Access Command Reference Guide*
- *Access Appliance Release Notes*
- *Veritas Access RESTful API Guide*
- *Access Appliance Solutions Guide for Enterprise Vault*
- *Access Appliance Solutions Guide for NetBackup*
- *Access Appliance Troubleshooting Guide*
- *Access Appliance Command Reference Guide*
- *Access Appliance Hardware Installation Guide*
- *Access Appliance Initial Configuration Guide*
- *Access Appliance Product Description*
- *Access Appliance Safety and Maintenance Guide*
- *Access Appliance Third-party Legal Notices Guide*
- *Access Appliance Upgrade Guide*