InfoScale™ Access 7.1 Release Notes

Linux

September 2016

VERITAS™

Veritas InfoScale Access Release Notes

Last updated: 2016-09-12

Document version: 1.0 Rev 2

Legal Notice

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

http://www.veritas.com

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

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

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1

Overview of Veritas InfoScale Access

This chapter includes the following topics:

About this release

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

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

Important release information

Review these Release Notes for the latest information before you install the product.

The hardware compatibility list contains information about supported hardware and is updated regularly. You can use any commodity hardware that is certified and mentioned in the hardware compatibility list. For the latest information on supported hardware, visit the following URL:

http://www.veritas.com/docs/000019707

Changes in this release

Table 1-1 shows the major new features and enhancements added in the 7.1 version of InfoScale Access.

Table 1-1 New features and enhancements in this release

Feature	Description
Ease of use	Web UI A web-based centralized operation, dashboards, workflows, alerts, and usage reports using the Management Server console of Veritas InfoScale Operations Manager. CLI Shell A menu-driven text interface. This interface provides a single point of administration for the entire cluster.
File system features	Cloud as a tier for scale-out file systems Move data to and from the public cloud using a tiering mechanism. Flexible Storage Sharing (FSS) Network sharing of direct-attached storage (DAS) local storage, cluster wide. Replication Periodic file system replication over IP networks. File system replication provides access to data for Disaster Recovery (DR) on the remote site. The user has an ability to consistently replicate a set of files to the DR site. SmartIO Read-write back caching on solid-state drives (SSDs), that helps in I/O acceleration. Snapshots Helps in recovering from data corruption. If files, or an entire file system, are deleted or become corrupted, you can replace them from the latest uncorrupted snapshot.
OpenStack plugin	Block access - OpenStack Cinder plugin File access - OpenStack Manila plugin

Table 1-1 New features and enhancements in this release *(continued)*

Feature	Description
Protocol support	NFS version 3 and NFS version 4 All the nodes in the cluster can server the same NFS shares both read and write (Active-Active). CIFS (SMB 3 and 2) All the nodes in the cluster can serve the same CIFS shares both read and write (Active-Active using CTDB mode). S3 Amazon object storage protocol that helps in accessing data using RESTful APIs. **Note:** InfoScale Access provides concurrent access to data over multiple protocols. In the InfoScale Access 7.1 release, concurrent access to data using multiple protocols is supported only for NFS and S3.

Chapter 2

Software limitations

This chapter includes the following topics:

Flexible Storage Sharing limitations

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

For FSS, only local direct-attached storage (DAS) disks are supported

For FSS, only local direct-attached storage (DAS) disks are supported. DAS disks can be either hard disk drives (HDDs) or solid state drives (SSDs).

Unable to create a file system on an FSS pool after the cluster is renamed (3787522)

In a FSS environment, when a disk is exported to the cluster, a host prefix is added to the disk name. The host prefix identifies the host to which the disk is connected. The disk names with the host prefix are permanently stored in the metadata. If a cluster is renamed, the host names are changed so the disk names are incorrect.

After the FSS disks have been formatted and given their unique names, the cluster name should not be changed. The cluster name can be changes when the cluster is reinstalled or upgraded.

Workaround:

There is no workaround. Do not change the cluster name during an upgrade.

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

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

Limitations related to installation and upgrade

The following limitations are related to installation and upgrade.

Licensing messages displayed for Beta version of InfoScale Access

The Beta version of InfoScale Access has a temporary license key. If you install a Beta version of the product, after 60 days you start seeing licensing messages such as the following:

```
WARNING V-365-1-1 This host is not entitled to run Veritas Storage
Foundation/Veritas Cluster Server.
As set forth in the End User License Agreement (EULA) you must complete one of the
two options set forth below. To comply with this condition of
the EULA and stop logging of this message, you have 0 days to either:
- make this host managed by a Management Server
(see http://go.veritas.com/sfhakeyless for details and free download), or
- add a valid license key matching the functionality in use on this host using the
command 'vxlicinst' and validate using the command 'vxkeyless set NONE'.
```

For Beta, these licensing messages cannot be avoided. The two options described do not apply for InfoScale Access.

The GA version of InfoScale Access has a permanent key.

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.
```

Limitations in the Veritas InfoScale Operations Management Server console

The InfoScale Access Insight Add-on is supported only on the Linux platform of Veritas InfoScale Operations Management Server. The InfoScale Access Insight Add-on cannot be deployed on a Windows server.

InfoScale Access IPv6 limitations

The following InfoScale Access modules are not supported for IPv6:

- NIS

The following IPv6 functionality is not supported for CIFS:

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

FTP create_homedirs limitation

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

See the *InfoScale Access Command-Line Administrator's Guide* for more information.

Samba ACL performance-related issues

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

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

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

InfoScale Access language support

InfoScale Access supports only English.

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

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

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

Limitations on using InfiniBand NICs in the InfoScale Access cluster

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

  Note: The case is observed on Mellanox card.

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

Note: The case is observed on Mellanox card.

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

Limitation on using InfoScale Access in a virtual machine environment

InfoScale Access must be installed on physical machines. InfoScale Access is not supported in a virtual machine environment.

File system limitation

The following issue relates to the InfoScale Access file system.

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

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

Chapter 3

Known issues

This chapter includes the following topics:

InfoScale Access known issues

The following known issues relate to the InfoScale Access commands.

Backup issues

This section describes known issues related to backup.

Backup or restore status may show invalid status after the BackupGrp is switched or failed over to the other node when the SAN client is enabled (3606322)

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

Workaround:

There is no workaround.

CIFS issues

This section describes known issues related to CIFS.

Cannot enable the quota on a file system that is appended or added to the list of homedir (3853674)

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

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

Workaround:

To solve this issue, follow these steps:

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

If the FC cables are pulled out from the node hosting the Management Server console, the CTDB server goes into the FAULTED state. CTDB service cannot be restored to ONLINE state when you use the "support-service-autofix" command (3845085)

The join state of winbind service is lost, so the winbind service cannot start. To restore the join state, you need to re-enter the Active Directory (AD) credentials. Currently the `support-service-autofix` command does not handle this case. Thus, the CTDB service cannot come ONLINE through the `support-service-autofix` command.

Workaround:

The CTDB server state can be recovered by using the `CIFS> server stop` command and the `CIFS> server start` command. This brings the CIFS server in the ONLINE state.

Deleting a CIFS share resets the default owner and group permissions for other CIFS shares on the same file system (3824576, 3836861)

When you delete a CIFS share, the owner and the group on the file system revert to the default permissions. The default values for both the owner and the group are

set to root. This behavior may be an issue if you have more than one CIFS share on the same file system. Deleting any of the shares also resets the owner and the group for the other shares on the file system.

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

Workaround:

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

```
CIFS> share modify
```

Deduplication issues

This section describes known issues related to deduplication.

Removing lost+found files for a mount point that has deduplication enabled may cause issues with deduplication (3472414)

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

Workaround:

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

To enable the duduplication job:

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

InfoScale Access Insight Add-on issues

The following known issues relate to the InfoScale Access Insight Add-on.

The disk used for I/O fencing cannot be distinguished in the Management Server console (IA-1235)

On an InfoScale Access cluster, if you enable I/O fencing during the installation and configuration process, and add the enclosure for discovery, the disks used for I/O fencing are identifiable. In the **Overview** tab of the enclosure, the InfoScale

Access Insight Add-on shows that some capacity is configured. On the **Physical Disks** tab, you cannot identify which disks are used for I/O fencing.

Workaround:

The configured space is for I/O fencing in the **Overview** tab of the enclosure.

There is no validation when you type an unavailable user name in the CIFS: Add Permission window (IA-1810)

When trying to add a permission to an unavailable user name in the **CIFS: Add Permission** window, no error message is displayed.

Workaround:

There is no workaround.

Data delays to sync up in the Management Server console if you perform any operation in CLISH (IA-1788)

If you perform any operation in CLISH, there is a delay seeing the update in the Management Server console.

Workaround:

There is no workaround.

Management Server console file system capacity (IA-1635)

After you export the document, the column name is not the same as it displays on the report. When you export the document, the size displays in bytes in the document.

Workaround:

There is no workaround.

Management Server console storage pool capacity (IA-1633)

The Management Server console storage pool often shows the volume column as 0 (zero).

Workaround:

There is no workaround.

Host consuming shares values are missing for the file system (IA-1632)

The host consuming shares always show its value as "–" in the file system column of the Host Consuming Shares Report.

Workaround:

There is no workaround.

CLISH and the Management Server console do not allow decimal values for the file system size

When you create a file system, you are not allowed to type decimal values both on the CLISH and the Management Server console.

Workaround:

It is by design.

The input description in the Create Pool wizard does not show in the Management Server console (IA-1497)

The input description in the **Create Pool** wizard does not show in the Management Server console.

Workaround:

There is no workaround.

In the Storage Pools tab, there is no column to show whether the created pool is isolated or not (IA-1484)

When you go to the **Storage Pools** tab, you cannot find any column to identify whether the pool is isolated or not.

Workaround:

There is no workaround.

Ascending or descending does not work correctly as per disk name in the Physical Disks table (IA-1340)

If the disks are named as *words_integer*, and the *integer* is larger than 10, ascending or descending does not work correctly in the **Physical Disks** table.

Workaround:

There is no workaround.

Extra columns show in the exported sheet (IA-1335)

If you export the sheet from the **Physical Disks** tab, it shows text that does not belong to the tables in the Management Server console.

Workaround:

There is no workaround.

RAID LEVEL displays in the Assert Summary (IA-1191)

If you go to **Enclosure** > **Storage Pools**, and click on one pool, then in the right panel, you can find extraneous information about the RAID Level in the **Assert Summary**.

Workaround:

There is no workaround.

InfoScale Access Insight Add-on should not be installed if the Storage Insight Add-on of the required version is not installed on the Veritas InfoScale Operations Manager server (IA-1840)

If you run the preinstall of the Storage Insight Add-on, you are not allowed to install InfoScale Access Insight Add-on.

Workaround:

There is no workaround.

Some special operations may make the enclosure-related tabs appear in the Properties windows (IA-1338)

Some operations may make the enclosure-related tabs appear in the **Properties** window when you leave the **Properties** window open.

Workaround:

There is no workaround.

You cannot grow or shrink the file system when it is offline (IA-1894)

You cannot grow or shrink the file system when it is offline.

Workaround:

There is no workaround.

You cannot create a snapshot with the same name of an existing snapshot (IA-1893)

You cannot create a snapshot with the same name of an existing snapshot.

Workaround:

There is no workaround.

When you export the FSS disks, and then open the export link again, the disks show as unexported (IA-1886)

If you check the updated status immediately after exporting disks, you can go to **Settings** > **Device** > **Enclosure Configurations**, and right-click on **Enclosure** and select **Refresh Configurations**. The updated status of the disks is reflected after you launch the wizard.

Workaround:

There is no workaround.

Enclosure does not recover to the Healthy status after it goes into the At Risk status (IA-1828)

When problems occur to an enclosure, and they are solved, the status of the enclosure does not recover to a **Healthy** status.

Workaround:

There is no workaround.

When you create CIFS shares, some options in Export Options may conflict among themselves (IA-1899)

When you create CIFS shares, some options in **Export Options** may conflict among themselves.

Workaround:

There is no workaround.

You are not allowed to perform multiple operations in parallel

When you start multiple operations in parallel, only the first operation may succeed, and others fail.

Workaround:

There is no workaround.

You may not see all the disk names in the Summary page (IA-1905)

If you select many disks in the **Create Storage Pool: Select Disks** page, you may not see the all the disk names.

Workaround:

There is no workaround.

After the nodes restart, the disk names are duplicated in the Physical Disks tab (IA-1918)

When one of the cluster nodes restart, you might find duplicated disks in the Management Server console.

Workaround:

Re-add the enclosure in the Management Server console.

Installation and configuration issues

The following issues relate to InfoScale Access installation and configuration.

Installer should install sysstat, hal, and python-paramiko as required operating system rpms to enable CLISH commands to operate correctly (IA-2810, IA-2928)

Because the necessary operating system rpms, `sysstat`, `hal`, and `python-paramiko` were not installed, CLISH commands can fail.

For example, the `lshal` command might fail as follows:

```
0 11:14:09 cmd /usr/bin/lshal | /bin/grep 'linux.sysfs_path =.*/net/.*'
| /bin/grep -v 'virtual' |/bin/awk -F \' '{print $2}'|
/bin/grep -iw eth0 2>&1 0 11:14:09 cmd exit=1 (duration: 0.01 seconds)
```

Workaround

Install the rpms manually on each node in the cluster.

Use the following yum commands to install the rpms:

```
yum install hal

yum install python-paramiko

yum install sysstat
```

Response file generated by the installer has an incorrect variable (IA-1863)

The response file generated by the installer located in `/opt/VRTS/install/logs/installaccessXXX/` contains an incorrect variable:*{opt}{vr}=1*. This variable may register the incorrect license on your node.

Delete the *{opt} {vr} =1* variable in your response file manually before using it.

After you restart a node that uses RDMA LLT, LLT does not work, or the gabconifg –a command shows the jeopardy state (IA-1796)

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

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

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

Workaround:

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

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

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

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

For adding an InfoScale Access node in an RDMA LLT environment:

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

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

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

If you run the installer from one of the cluster nodes, the installer log is placed in /var/tmp/, instead of /opt/VRTS/install/logs/ (IA-1848)

Each installer has a log named: `installaccess-`*`timestamp`*`/uninstallaccess-`*`timestamp`*. If you run the installer from the local management console of the cluster node, when the installer is running, the log is placed in the temporary location `/var/tmp`. After the installer is finished, the log is still in `/var/tmp`.

Workaround:

Retrieve the installer log from the `/var/tmp` directory.

Error message Failed to update /boot/initramfs-2.6.32-504.el6.x86_64.img shows in the install log of package VRTSvxvm on RHEL 6.6 (IA-1897)

In the installer logs of RHEL6.6, the package `VRTSvxvm` has a separate log named `install.VRTSvxvm.node_ip`. An error message shows that `Failed to update /boot/initramfs-2.6.32-504.el6.x86_64.img`. It is because the default version of the OS package `nss-softokn-freebl.x86_64` and `nss-softokn-freebl.i686` is `3.14.3-17.el6`, but the `VRTSvxvm` required version is `3.14.3-23.el6_7`.

Workaround:

Update the OS package version of `nss-softokn-freebl.x86_64` and `nss-softokn-freebl.i686` to `3.14.3-23.el6_7`. After that, you can use the installer.

Running individual InfoScale Access scripts may return inconsistent return codes (3796864)

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

Configuring InfoScale Access with the installer fails when the SSH connection is lost (3794964)

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

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

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

Workaround:

To work around this issue:

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

Excluding PCIs from the configuration fails when you configure InfoScale Access using a response file (3686704)

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

Workaround:

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

Installer does not list the initialized disks immediately after initializing the disks during I/O fencing configuration (3659716)

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

Workaround:

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

After reinstalling the InfoScale Access cluster after an uninstallation, the nas_init script is not executed (IA-2534)

After an uninstallation of an InfoScale Access cluster, the `nas_init` script is not executed due to stale entries in the `chkconfig` file that were not cleaned up.

Workaround:

Make sure there are no other *nas_* scripts in the /etc/rc.d/ directory

1 Before starting the reinstallation of the InfoScale Access cluster, run the following commands:

```
# find /etc/rc.d -name *nas_*
/etc/rc.d/rc4.d/K31nas_init
/etc/rc.d/init.d/nas_always
/etc/rc.d/init.d/nas_init
/etc/rc.d/init.d/nas_scan
/etc/rc.d/init.d/snas_notify
/etc/rc.d/init.d/nas_prevcs
/etc/rc.d/rc3.d/S05snas_notify
/etc/rc.d/rc3.d/S69nas_init
/etc/rc.d/rc6.d/K31nas_init
/etc/rc.d/rc5.d/S05snas_notify
/etc/rc.d/rc5.d/S69nas_init
/etc/rc.d/rc0.d/K31nas_init
/etc/rc.d/rc2.d/K31nas_init
/etc/rc.d/rc1.d/K31nas_init
#
```

2 Remove all the links.

```
# for file in `find /etc/rc.d -name *nas_*`; do rm -rf $file >
/dev/null 2>&1; done
```

3 Confirm that there are no `nas_*` files in the `/etc/rc.d/` directory.

Networking issues

This section describes known issues related to networking.

Network DNS set nameserver accepts the same IP address multiple times (IA-905)

The CLI accepts multiple nameserver names, but utilizes one only.

Workaround:

Do not use multiple nameservers in the CLI.

CVM service group goes into faulted state unexpectedly (3793413)

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

Workaround:

Reboot the node on which this issue occurs.

Creating network interface bond when other nodes of the cluster are down results in inconsistent network configuration state (3790781)

When a few nodes on the cluster are down, the `Network> bond create` command fails to synchronize the required configuration files to the unavailable nodes. When the nodes are rebooted, the network group has an invalid configuration.

Workaround:

Create a bond only when all nodes of the cluster are up and running.

In a mixed IPv4 and IPv6 VIP network setup, the IP balancing does not consider IP type (3616561)

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

Workaround:

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

The netgroup search does not continue to search in NIS if the entry is not found in LDAP (3559219)

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

Workaround:

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

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

VIP and PIP hosted on an interface that is not the current IPv6 default gateway interface are not reachable outside the current IPv6 subnet (3596284)

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

Workaround:

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

NFS issues

This section describes NFS issues.

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

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

Workaround:

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

Random-write performance drop of NFS-Ganesha with Linux clients (IA-1304)

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

Workaround:

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

Open request during the lock reclaim frequently fails with NFS4ERR_NO_GRACE in the node restart process (IA-1225)

When you restart the cluster node, all the clients need to reclaim the existing locks (if any). Sometimes the reclaim operation fails due to error `NFS4ERR_NO_GRACE` when you open the file.

Workaround:

There is no workaround.

Latest directory content of server is not visible to the client if time is not synchronized across the nodes (IA-1002)

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

Workaround:

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

NFS share show may list the shares as faulted for some time if you restart the cluster node (IA-1838)

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

Workaround:

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

NFS-Ganesha shares faults after the NFS configuration is imported(IA-849)

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

Workaround:

Restart the NFS service.

NFS-Ganesha shares may not come online when the number of shares are more than 500 (IA-1844)

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

Workaround:

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

Kerberos configuration is not done on newly added cluster nodes using the add node feature (IA-963)

Kerberos configuration is not done on newly added cluster nodes using the `Cluster> add` node command.

Workaround:

There is no workaround.

NFS-Ganesha I/O workload fails for shares that are exported with netgroups (IA-1911)

NFS-Ganesha I/O workload fails for shares that are exported with netgroups.

Workaround:

There is no workaround.

Exporting a single path to multiple clients through multiple exports does not work with NFS-Ganesha (3816074, 3819836)

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

Workaround:

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

For the NFS-Ganesha server, bringing a large number of shares online or offline takes a long time (3847271)

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

Workaround:

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

NFS client application may fail with the stale file handle error on node reboot (3828442)

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

Workaround:

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

ObjectAccess issues

This section describes ObjectAccess issues.

Object access service faults when you upload a 10-GB file with a chunk size of 5 MB (IA-1943)

The object access service faults when you upload a 10-GB file with a chunk size of 5 MB.

Workaround:

If the file size is large, use a chunk size that is less than or equal to 100 MB.

Object access service fails to assemble the uploaded chunks after failover (IA-1944)

During a failover, if the multipart upload is in progress, the object access service fails to assemble the uploaded chunks in the bucket.

Workaround:

After the failover is done, start the multipart upload process again.

OpenStack issues

The following issues are related to OpenStack.

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

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

Workaround:

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

Replication issues

This section describes known issues related to replication.

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

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

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

Workaround:

There is no workaround.

The System> config import command does not import replication keys and jobs (3822515)

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

Workaround:

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

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

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

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

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

5 Create the job from the new source.

Replication job with encryption fails after job remove and add link with SSL certificate error (3839319)

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

`SSL certificate error.`

Workaround:

Follow these steps to solve this issue:

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

Replication job status shows the entry for a link that was removed (3797560)

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

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

Workaround:

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

Creating or deleting NIC bonds stops the replication services (3680071)

The `Network> bond create` and `Network> bond remove` operations involve bringing down the interfaces and then bringing them back up. These operations stop the replication services.

Workaround:

Perform these operations while the replication services are offline. After changing the bond configuration, start the replication services using the following command:

```
Replication> start service
```

The job uses the schedule on the target after replication failover (3668957)

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

Workaround:

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

Storage issues

The following issues relate to the InfoScale Access Storage commands.

Snapshot mount can fail if the snapshot quota is set (IA-1542)

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

Workaround:

Remove the oldest checkpoint and retry.

Sometimes the Storage> pool rmdisk command does not print a message (IA-1733)

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

Workaround:

Check the command output log files to obtain the command's return status.

The Storage> Pool rmdisk command sometimes can give an error where the file system name is not printed (IA-1639)

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

Workaround:

There is no workaround.

If you run the Cluster> show command when a slave node is in the restart, shutdown, or crash state, the slave node throws an exception (IA-900)

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

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

Workaround:

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

If duplicate PCI IDs are added for the PCI exclusion, the Cluster> add node name command fails (IA-1850)

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

Workaround:

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

Reinstalling some pools is visible in the output of the Storage> fss disk list command, but it is not visible in the output of the Storage> disk list detail command (IA-1409)

If you do not clean the cluster, uninstall the InfoScale Access stack from the cluster, and then install the stack again, then due to an unclean reinstallation, even though the pool does not exist, the previous pool tag may remain on the disk.

Workaround:

Contact Technical Support, and request that Technical Support run the following command:

```
# vxdisk rmtag site=pool name disk name
```

The Storage> pool create command fails to create a pool with the local or FSS disk as expected, and CLISH does not show an error message (IA-1660)

You use the `Storage> pool create` command to create a pool with shared SAN LUNs. If you try to create a pool with a local or an FSS disk by using the `Storage> pool create` command, then it fails. So as per expectation, the `Storage pool create` command fails to create a pool with the local or FSS disk, but CLISH does not display an error message.

Workaround:

Do not use FSS disks with the `Storage> pool create` command. For FSS disks, use the `Storage> fss pool create` command instead.

If a file system is used as homedir or anonymous_login_dir for FTP, this file system cannot be destroyed (IA-1876)

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

Workaround:

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

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

The Storage> pool mvdisk command should not allow you to move the disk from a SAN pool to an FSS pool (IA-1912)

Do not use the `Storage> pool mvdisk` command to move disks between a SAN pool and an FSS pool.

Workaround:

Use the `Storage> fss pool rmdisk fsspool1` command and then use `Storage> fss pool adddisk fsspool2` to move a disk from `fsspool1` to `fsspool2`.

Removing a disk from the FSS pool can sometimes unexpectedly fail and an error message displays (IA-1879)

You see the following error message:

```
Storage> fss pool rmdisk fsspool disk_0 ACCESS fss ERROR V-288-349
Pool fsspool is SAN pool
```

There is no operational effect as the disk can still be used; however, the disk cannot be removed from the pool.

Workaround:

There is no workaround. Contact Technical Support if you experience this issue.

The output from the Storage> fss pool list command is empty (IA-1880)

Sometimes after you delete an FSS file system or after you remove the disks from an FSS pool, the output from the `Storage > fss pool list` command disappears. For example:

```
Storage> fss pool list
Pool                    Isolated              List of disks
==================      ==================    ==================
```

Workaround:

There is no workaround. Contact Technical Support if you experience this issue.

The FSS disk names do not show the host name prefix for the disks owned by the management console (IA-1866)

The command `Storage> fs list filesystemname` can be run only on the management console. The disks names are missing their prefixes in the display. For example:

Mirror 03 is on node 1. The command output displays `disk_0` and `disk_1` for mirror 03 when the management console is on node 1.

1. Mirror 01:

   List of pools: `fsspool`

   List of disks: `Drive0418b_03_disk_0 Drive0418b_03_disk_1`

2. Mirror 02:

List of pools: `fsspool`

List of disks: `Drive0418b_02_disk_0 Drive0418b_02_disk_1`

3. Mirror 03:

List of pools: `fsspool`

List of disks: `disk_0 disk_1`

Workaround:

Use the `Storage> fss pool list fsspool` command where the pool name of `fsspool` is obtained from the output of `Storage> fs list filesystemname`.

Not able to enable quota for file system that is newly added in the list of CIFS home directories (IA-1851)

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

Workaround:

Run the following commands from CLISH:

```
Storage> quota cifshomedir disable

Storage> quota cifshomedir enable
```

Destroying the file system may not remove the /etc/mtab entry for the mount point (3801216)

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

Workaround:

Manually remove the `/etc/mtab` entry, or reboot the cluster nodes.

The Storage> fs online command returns an error, but the file system is online after several minutes (3650635)

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

```
access.Storage> fs online fs1

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

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

Workaround:

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

Storage fencing off does not work (3811579)

Because VCS ports are opened, the `vxfenswap` command thinks that there are two nodes in the cluster, but the `llthosts` file only has one node. This causes the `vxfenswap` command to hang and the `vxfen` driver not to unload.

Removing disks from the pool fails if a DCO exists (3452098)

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

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

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

The rollback refresh fails after the file system growby/growto command (3588248)

The rollback refresh fails if it is executed after increasing the file system size.

Workaround:

There is no workaround.

Technical preview issues

The following issues relate to the InfoScale Access technical preview features.

On a scale-out file system, ENOSPC condition is hit even if there is reserved space (3851290)

For the scale-out file system, some reserved space is kept to store the metadata. Once the remaining space is occupied by data, the scale-out file system cannot

use this reserved space. This results in getting an ENOSPC error even if there is free space. In some cases, this free space can be around 18% - 19% of the total space, but the space may vary depending on the metadata usage. There is no external control on how much space to reserve for metadata.

Workaround:

There is no workaround.

Checkpoint creation fails with an ENOSPC error even though there is some space (3817122)

The checkpoint creation fails for the scale-out file system with error ENOSPC even if there is some space. For the scale-out file system, data is managed in different ways than the traditional file system.

Workaround:

There is no workaround.

Removable snapshots for scale-out file systems are not removed if the system runs out of space (3830523)

For scale-out file systems, removable snapshots are not automatically removed if the system runs out of space. When you create a snapshot for a scale-out file system, the option to set the `removable` attribute is allowed. However, InfoScale Access does not automatically remove any snapshots for scale-out file systems even when the `removable` attribute is set to `yes`.

Workaround:

If required, manually remove the scale-out file system snapshots using the following command:

```
Storage> snapshot destroy snapshot_name fs_name
```

Chapter 4

Getting help

This chapter includes the following topics:

Displaying the online Help

You can access the online Help through the Management Server console of Veritas InfoScale Operations Manager by clicking **Help**.

Displaying the man pages

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

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

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

Using the InfoScale Access product documentation

InfoScale Access product documentation is available in the Adobe Portable Document Format (PDF) 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 latest documentation is available on the SORT website.

The following documents are available on the SORT site:

- *InfoScale Access Command-Line Administrator's Guide*
- *InfoScale Access Installation Guide*
- *InfoScale Access Quick Start Guide*
- *InfoScale Access Release Notes*
- *InfoScale Access Third-Party License Agreements*
- *InfoScale Access Troubleshooting Guide*
- *Veritas InfoScale Operations Manager Add-on User Guide*