



## **Avere OS 4.7.3.2.C8 Release Notes**

2018-05-21

## Copyright Information

Copyright © 2009-2018 Avere Systems, Inc. All rights reserved. Specifications subject to change without notice.

No part of this document covered by copyright may be reproduced in any form or by any means – graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system – without prior written permission of the copyright owner.

The product described in this document may be protected by one or more U.S. patents, foreign patents, or pending applications.

## Table of Contents

Upgrading to Avere OS 4.7

New in Avere OS 4.7.3.2.C8

New in Avere OS 4.7.3.2.C7

New in Avere OS 4.7.3.2.C6

New in Avere OS 4.7.3.2.C5

New in Avere OS 4.7.3.2.C4

New in Avere OS 4.7.3.2.C3

New in Avere OS 4.7.3.2.C2

New in Avere OS 4.7.3.2.C1

New in Avere OS 4.7.3.2

New in Avere OS 4.7.3.1

Contact Support - Avere Global Services

# Upgrading to Avere OS 4.7

Systems running Avere OS 4.0 or later can upgrade directly to the current version of Avere OS 4.7 without installing any intermediate versions.

If you are upgrading a system using Avere OS 3.2.1.2.C12 or earlier, please upgrade to 3.2.1.2.C15 before installing a 4.6 or 4.7 build. Contact Avere Global Services if you have questions.

When upgrading an Avere Cluster from version 4.6 or earlier, note that AWS-hosted vFXT clusters might require additional privileges (described below).

## Existing AWS vFXT Clusters Might Need Updated Role

A change in Amazon Web Services security might require you to manually update the cluster role for existing AWS vFXT clusters. The cluster role now must include permission for the `ec2:DescribeSubnets` action.

This change is only required for existing vFXT clusters that were created with cluster creation software released before version 4.7.3.1. Both the Cluster Manager software and the command-line `vfxt.py` script that were distributed with Avere OS 4.7.3.1 include the new configuration automatically.

Also, you can use the newer `vfxt.py` script with an Avere OS version 4.6 image to create a 4.6 vFXT cluster that automatically includes the new action.

To edit the cluster role:

1. From the AWS console, go to the IAM service section.
2. Click the **Roles** link in the left sidebar.
3. Find the cluster role by searching the list for the name of your cluster.
4. Select the cluster role and click **Edit Policy**.
5. Add the DescribeSubnets statement to the list of EC2 actions:

```
...
    "Effect": "Allow",
      "Action": [
        "ec2:AssignPrivateIpAddresses",
        "ec2:UnassignPrivateIpAddresses",
        "ec2:DescribeInstance*",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSubnets",
        "ec2:ReplaceRoute",
        "ec2:CreateRoute",
        "ec2>DeleteRoute"
      ],
      "Resource":
    "*"https://download.averesystems.com/software/avereos_4.7_release_notes_4732.pdf
  ...
```


6. Save the edited policy.

# New in Avere OS 4.7.3.2.C8

The Avere Privacy Policy has changed. Administrators should review the changed policy before using (or continuing to use) Avere OS support data upload features.

The current document can be read at <http://www.averesystems.com/privacy-policy>.

If you have not yet accepted the terms, a pop-up dialog appears when you load the **Support** tab or the **Cluster > Support** configuration page on the **Settings** tab of the Avere Control Panel. This dialog includes a link to the revised policy. Click the **Confirm** button to agree.



Links to the current terms of use and privacy policy statements have been added to the left margin of the **Support** tab.

# New in Avere OS 4.7.3.2.C7

This update includes a variety of fixes and improvements.

## Resolved Issues

### Cloud Object Store

- 23833 Added a lock to prevent a problem reading metadata from cloud storage that could cause a system restart.
- 25155 Snapshots cannot be cancelled while being written to the core filer. Before this change, attempting to cancel a snapshot that was in this state could cause an error.

### Filesystem

- 24038 Improved inode handling; this change prevents a rare situation that could cause a system restart when the file count is greater than the allotted number of cached inodes.
- 24654 Improved memory buffer allocation to avoid a system restart that could affect large clusters.
- 25003 Fixed a race condition related to inode storage that could cause a system restart.

### NFS

- 23996 Fixed a defect that could cause new client TCP connections (NFS, MOUNT, NLM, NSM) to be dropped with a 15 second delay due to an internal resource limit. Before this change, the only way to solve the connection issues was to restart the affected node.

# New in Avere OS 4.7.3.2.C6

The changes listed in this section were added in version 4.7.3.2.C6.

## Resolved Issues

### Cloud Object Store

- |       |  |
|-------|--|
| 24369 | Fixed a race condition that sometimes caused a filesystem service failure during post-processing after a cloud snapshot.   |
| 24633 | The minimum number of TCP connections to a cloud core filer now can be set as low as 32. A previous update changed the minimum to 128, which caused problems in some environments. |

### Filesystem

- |       |  |
|-------|--|
| 24061 | Changed code to prevent a system restart caused by an incorrect policy ID being applied to an inode.   |
| 24986 | Fixed an infrequently seen issue that could cause a FlashMove® or FlashMirror® job to incorrectly assign junction locations at the end of the migration.   |
| 24989 | This change prevents a problem that could cause a node to restart when upgrading to a 4.7 release. The issue occurred only if the number of cached files exceeded the file limit when the upgrade started. |

# New in Avere OS 4.7.3.2.C5

The following fixes were introduced in Avere OS 4.7.3.2.C5.

## Resolved Issues

### Cloud Object Store

- |       |   |
|-------|---|
| 23884 | Fixed an error that could result in a filesystem restart after invalidating a cloud core filer. |
|-------|---|

### General

- |       |  |
|-------|--|
| 19519 | <p>The Avere Secure Proactive Support (SPS) system now provides more options for allowing remote access by Avere Global Services staff. Cluster administrators now can choose from the following levels of remote access:</p> <ul style="list-style-type: none"><li>• Disabled - No SPS access</li><li>• Support - Allows remote access for support tasks and statistics gathering</li><li>• XMLRPC - Allows remote access for any XML-RPC API call</li><li>• Full - Allows remote access for any API calls as well as any shell command</li></ul> <p>Before this change, administrators had two options: to opt in for access equivalent to the "Support" level, or to opt out.</p> |
|-------|--|

- 24218 Fixed a software error that caused an incorrect "invalid license" error when replacing an expired Avere OS feature license.
- 24655 Fixed a memory leak in Local Directories management that could result in a filesystem restart.
- 24852 Added support for new SSD drives now being used in Avere FXT 5000 Series hardware. Note that if you receive any hardware that includes these drives, you must upgrade your Avere OS software to this release, and you cannot downgrade to a previous version. (Avere OS 4.6.2.5.C10 also introduces support for these components.)

## Filesystem

- 23176 Addressed an issue that caused spurious alerts about operations taking too long to complete.

## NFS

- 24127 This change prevents a service restart when the "showmount" feature is enabled in NetApp Clustered Data ONTAP version 8.3 or later.

The showmount feature is unsupported; use the following configuration command to disable it:  
`nfs server modify -vserver ${VSERVER} -showmount disabled`

## SMB/CIFS

- 21487 This change allows Microsoft Previous Versions to display correct timestamps when multiple SMB vservers are defined. Before this change, NAS core filers with a large number of NFS exports or snapshots could behave erratically, and log repeated "no FH entry found" messages.
- 24101 Removes the vulnerability documented in security advisory [CVE-2017-7494](#).
- 24838 Changed code to fix a service restart associated with large directories on a share that has the "hide unreadable" (Access Based Enumeration) option set.

## vFXT

- 23995 Added support for using r4.2xlarge and r4.8xlarge instance types to create Avere vFXT clusters.

# New in Avere OS 4.7.3.2.C4

Avere OS 4.7.3.2.C4 included two changes to prevent system restarts, and other stability enhancements. Note that Avere OS 4.7.3.2.C4 blocks downgrading to earlier software if used with SMB ACLs on cloud core filers - see the description of item 23827 in [Resolved Issues](#), below.

## Resolved Issues

### Filesystem

- 24031 Corrected an issue that could cause a faulty inode reference when recycling an inode to free cache space. Before this change, the problem could lead to a service restart.

## SMB/CIFS

- 23827 A problem was fixed that could cause a filesystem service restart under some circumstances when updating ACLs on cluster junctions backed by cloud core filers. The restart was triggered by writes to discretionary ACLs with more than 20 access control entries (ACEs). This change supports ACLs with up to 200 ACEs, and larger ACLs return a client IO error.
- NOTE:** This fix will permanently update any SMB ACL used with cloud storage at the time the ACL is modified; after this update, you *cannot downgrade* to a release before 4.7.3.2.C4.
- 23881 Disabled a performance enhancement for fetching ACLs (documented [below](#) in item 20205) to investigate a possible performance regression related to ACL fetching.
- 24435 Removed an unneeded least-recently-used optimization for the ACL cache. The optimization was superseded by other memory use improvements, and is suspected to contribute to other performance issues. (The disabled change is item 22177 from [version 4.7.3.1.](#))

## New in Avere OS 4.7.3.2.C3

This critical release fixed a memory issue in the cloud cache that could consume all available cache memory and cause stuck operations in one or more nodes that could not be resolved without a restart. (24168)

## New in Avere OS 4.7.3.2.C2

The following fixes and improvements were originally published in Avere OS version 4.7.3.2.C2.

### Cloud Object Store

- 23904 Eliminated a race condition that could cause a filesystem restart when removing directories on a cloud core filer.
- 24005 Changed code to more efficiently handle large file deletions in cloud storage. Before this change, concurrently deleting very large files could result in a long-running delete operation that appeared to be stuck.

### Filesystem

- 22642 Improved directory tracing and validation during HA failover.
- 23721 Fixed an issue that could cause the number of client writes to hit an internal resource limit and result in a restart.
- 23783 A logic error was fixed that could cause the cloud snapshot process to stall indefinitely after an operation failed to trigger the next steps.
- 23875 A problem was fixed that prevented FlashMove or FlashMirror jobs to a destination path on a NetApp Clustered DataONTAP system. Before this change, the data management job would fail with a "read-only filesystem" error.
- 23982 Improvements were made in identifying and handling items left out of sync from a previous operation when starting a new Data Management job.

### General

- 23767 Changes were made to the network driver to improve performance and eliminate some traffic imbalance issues in link aggregation. Changes included enabling additional queues and software interrupt modulation.



- 23867 Fixed a performance issue caused by incorrect default size settings being applied to some internal caches. Affected caches included the Access Control List (ACL) cache and inode cache.
- 23926 When using the Avere OS XML-RPC interface `cluster.modifyCloudRegion` in a custom Amazon Web Services cloud environment, you must supply the IAM service endpoint.

## NFS

- 23869 Fixed an issue that caused internal numeric identifiers to be reported instead of core filer export names when listing hot files. The error affected both the Avere Control Panel Dashboard and XML-RPC calls.
- 24020 A defect was fixed that could cause cloud NLM lock operations to return grace period errors indefinitely.

## SMB/CIFS

- 23288 Memory utilization by SMB service processes has been reduced to avoid SMB service restarts caused by out-of-memory conditions.

# New in Avere OS 4.7.3.2.C1

The C1 critical release fixed several high-priority issues related to multi-node failures when writing to cloud object store backends. These issues were found in internal testing and are not believed to have affected any customer installations.

# New in Avere OS 4.7.3.2

## New Features and Enhancements

### Suspend a Cloud Core Filer

This release adds the ability to suspend a cloud core filer. Previously, the suspend feature could only be used with NAS hardware core filers.

Suspending a core filer disables client access without permanently removing the filer. Use the **Manage Core Filers** settings page in the Avere Control Panel to suspend or reactivate core filers.

### SSH Security Update

Version 4.7.3.2 includes additional security improvements, including revisions to the supported cipher suites, HMAC digests, and Key Exchange (KEX) algorithms.

The cluster SSH server accepts **only** the following security algorithms:

Ciphers	chacha20-poly1305@openssh.com aes256-gcm@openssh.com aes128-gcm@openssh.com* aes256-ctr aes192-ctr* aes128-ctr*
<i>(continued)</i>	

MACs	hmac-sha2-512-etm@openssh.com hmac-sha2-256-etm@openssh.com umac-128-etm@openssh.com* hmac-sha2-512 hmac-sha2-256 umac-128@openssh.com* hmac-sha1*
KEX Algorithms	curve25519-sha256@libssh.org ecdh-sha2-nistp521 ecdh-sha2-nistp384 ecdh-sha2-nistp256 diffie-hellman-group-exchange-sha256 diffie-hellman-group14-sha1*

\*Note: SHA1 and AES128 are provided for backward-compatibility with outdated SSH clients only. These algorithms will be phased out in future releases. Please update all SSH clients to use more modern algorithms.

## Resolved Issues and Other Changes

### Cloud Filesystem Change Delayed

Based on additional testing, part of the cloud storage filesystem change included in version 4.7.3.1 was removed from Avere OS version 4.7.3.2. Avere Systems is reworking the feature for inclusion in a future release.

The functionality described in the 4.7.3.1 section of this document under the headings [Cloud Filesystem Improvements](#) and [Object Storage Repair Tool](#) is not exactly what is included in Avere OS 4.7.3.2.

Specifically:

- The cloud filesystem metadata changes described for 4.7.3.1 are not applied to existing directories and core filers in 4.7.3.2. New cloud-based files and storage systems use the improved filesystem, but storage buckets created and used with Avere OS versions before 4.7.3.2 are not converted.
- The object storage repair tool relies on the new metadata structure, so in 4.7.3.2 the repair tool can be used only on directories and cloud storage systems that were created using a 4.7.3.\* distribution.
- All other features described in the 4.7.3.1 section of this document (starting with HA failover improvements and cloud NLM) are active in version 4.7.3.2.

### End of Support for Dell EMC Atmos Core Filer

Support for the Dell EMC Atmos storage product was discontinued in Avere OS version 4.7.3.1. Avere Systems recommends that Dell EMC customers migrate to Dell EMC's Elastic Cloud Storage (ECS) option, which is a supported core filer.

### Cloud Core Filer Object Store (FlashCloud)

- |       |  |
|-------|--|
| 23487 | Fixed a bug that could sometimes result in stuck directory removal operations during failover.                     |
| 23515 | Fixed an issue that caused failures when trying to change between HTTP and HTTPS protocols on a cloud core filer.  |
| 23667 | Fixed a race condition that could cause file lookup and create operations to become stuck during a failover event. |

### Filesystem

- |       |   |
|-------|---|
| 20937 | Fixed a bug that caused an incorrect file length to be returned for a file that was simultaneously being removed and written to back-end storage. |
|-------|---|

- 21717 Fixed a bug that caused the error badhandle in response an API call to create a subdirectory GNS junction for a cloud core filer.
- 23224 Added a new custom setting that allows administrators to set a maximum permitted file size manually instead of relying on the size reported by a NAS core filer. This setting can be useful if a core filer misreports the file size it can effectively handle. Contact Avere Global Services if you need to use this setting.
- 23320 Fixed an issue where a cloud core filer was not effectively using the read-ahead mechanism to deliver optimal throughput when pre-filling the cache.
- 23355 Fixed an issue that could impact performance by causing a heavily read file to be invalidated in the cache by active writes.
- 23465 Fixed an issue that could disrupt SMB access to the desktop.ini file because of an internal counter rollover in systems with high uptime. This problem could cause Windows Explorer to become unresponsive.
- 23541 Fixed a problem causing SMB directory listings to fail to complete while the directory was processing active creates or deletes.

## FlashMove/FlashMirror

- 21583 Fixed an issue where new junctions to hierarchical exports that have previously been the source or target in a migration are unusable.
- 21724 Improved error handling for the "overwrite existing content" operation of FlashMove and FlashMirror jobs.
- 23182 Fixed an issue that could cause a node reboot loop after removing a cluster node under certain FlashMove/FlashMirror conditions.

## General

- 21312 Fixed a display bug where the old node name was used in certain dashboard alerts and conditions after using the XML-RPC API call node.rename to update the name.
- 23519 Software was changed to prevent downgrading FXT 5000 series nodes with SSD system drives to Avere OS versions below 4.6.2.5.C1. Earlier software versions do not support SSD system drives.
- 22980 Fixed bugs in the statistics counter archive server that could cause internal memory leaks and error messages when monitoring hot clients.
- 23513 Fixed a memory buffer issue that could cause high latency for inbound client writes while the system was pre-filling read data from a cloud core filer.
- 21156 Fixed a bug that prevented proper support for VLAN names that consist of only numerical digits.
- 22808 Changed configuration code to prevent removing a node from a cluster where high availability has been disabled. Removing a node from a non-HA cluster is unsupported. This change affects both the Avere Control Panel web interface and the XML-RPC API.
- 22825 Updated the code for reporting power supply failure alerts to reduce the likelihood of false positives.
- 23404 Updated links to documentation and other tools on the Support tab of the Avere Control Panel.
- 23545 Updated warnings and tooltips shown when creating a new vserver in the Avere Control Panel.
- 23713 Changed the Avere Control Panel Support Settings page to clarify that administrators should enter a unique identifier for the cluster in the Customer Information section.

## NFS

- 16270 Fixed a bug where the path of the NFS export name was misrepresented in the HotFiles listing.

## SMB/CIFS

- 20205 Improved the performance of file ACL reads by caching the ACLs when a directory is listed.
- 22638 Fixed an issue that caused user/group name information download from Active Directory to fail, causing an incomplete username mapping.
- 23444 Prevent out-of-memory conditions by limiting outstanding read and write operations to 256 total. Before the change, this issue could cause SMB filesystem service restarts.
- 23541 Fixed a problem that prevented SMB directory listings from completing while creates and/or deletes were being processed in the same directory.
- 23638 Fixed an issue that caused poor performance in internal SMB path-name resolution when top-level parent directories were actively changing.
- 23651 Fixed a spurious error message introduced in Avere OS version 4.7.3.1 that occurred when SMB protocol access was disabled and then re-enabled on a vserver. The error message reported that there was an error logging in to the AD domain, but SMB access was re-enabled despite the error.

## vFXT

- 23502 Improved performance by optimizing interrupt handling for disk and network devices in Google Compute Engine.

# New in Avere OS 4.7.3.1

Avere OS 4.7.3.1 is a significant update from 4.6 releases. It includes improvements to SMB infrastructure and for working with cloud object storage, as well as security upgrades and expanded functionality for the optional FlashMove® and FlashMirror® data migration features.

## New Features and Enhancements

### Cloud Filesystem Improvements

**Note:** *Some of this functionality is disabled in version 4.7.3.2.*

This release includes several significant improvements for filesystem interaction with cloud object storage.

#### Faster File Creation

File creation times for both NFS and SMB operations to cloud storage have been made more efficient. Improvements to algorithms and to locking strategies eliminate bottlenecks that previously caused create operations and large remote synchronization processes to run slowly.

#### Directory Metadata Hardening

**Note:** *In Avere OS version 4.7.3.2, the filesystem conversion functionality is disabled but newly created storage buckets and cloud core filers use the improved format.*

This release introduces a more robust strategy for storing directory metadata, which allows directory contents to be recovered easily if directory structure becomes corrupt.

As part of this change, a new naming format was implemented for cloud objects. When you install Avere OS 4.7.3.1, existing files are automatically updated to the new format as a low-impact background task.

**Note:** This format change is irreversible, and nodes that use Avere OS 4.6.\*.\* cannot access objects in the updated format. Avere OS 4.7.3.1 can read objects in either format, but will attempt to convert any old-format volumes to the new format.

Client access to files on cloud core filers might be affected during software upgrades to Avere OS 4.7.3.1 because any un-updated nodes will be unable to access files that have been reformatted. Consider upgrading cluster software during a low-demand time or in a maintenance window.

## Object Storage Repair Tool

**Note:** This functionality has limited application in version 4.7.3.2.

Avere OS v4.7.3.1 introduces a utility that helps Avere Global Services repair inconsistencies between a cloud object store and the Avere cluster directory that references the files. This utility detects and repairs missing segments, incorrect link counts, and inconsistent pointers; if unrepaired, the files could become unreadable.

This tool can be used by Avere Systems support staff only. Contact Avere Global Services for additional information.

## HA Failover Improvements

HA Failover Improvements were made to the infrastructure that supports Avere's high availability feature (HA). The HA system replicates data on multiple nodes in an Avere cluster so that cached data is seamlessly available even if a disk fails or a cluster node becomes unreachable. Failover events can occur during FXT upgrades, filesystem service restarts (planned and unplanned), OS reboots, and when cluster nodes are added or removed.

Changes were made to improve the speed, efficiency and robustness of the failover process, primarily in two specific failover scenarios:

- the powering down or losing power for physical nodes
- the pulling of SSD and HDD drives

We also eliminated some rare problems that could cause a kernel error when a disk or node suddenly became unavailable. These could interrupt a proper failover, leaving data unavailable to clients for periods of 90 seconds to several minutes. (21696, 22224)

## Cloud NLM

Avere OS now supports a Network Lock Manager (NLM) file locking service for NFSv3 clients. The new service responds to NLM calls from NFS clients when accessing data that is stored on a cloud-based core filer (AWS S3, Google GCS, private object).

The new cloud NLM service is used only with cloud core filers; when using a traditional NAS core filer, the Avere cluster relies on that core filer's NLM service to provide locking functionality.

The cloud NLM service facilitates both byte-range locking and share reservations. In this release, a cluster can support approximately 500,000 locks per node for N-1 cluster nodes - that is, a cluster with four FXT nodes could support 1,500,000 locks, since  $(4-1) * 500,000 = 1,500,000$ . To maintain high performance, Avere recommends limiting the number of locks for an individual file to a maximum of 1000 locks.

The cloud NLM service is automatically available on clusters created with Avere OS 4.7.3.1 or later; contact Avere Global Services to enable the feature on clusters that are upgraded from an earlier release. When this feature is enabled, any clients that mounted the Avere cluster using the `nolock` or `local_lock` advanced options will need to remount the cluster with those options turned off in order to use the NLM feature. (20298)

## Data Migration Improvements (FlashMove® and FlashMirror®)

Avere OS version 4.7.3.1 includes significant improvements in the FlashMove® and FlashMirror® data migration features. The changes make data management jobs easier to define, more robust, and easier to debug, and give administrators greater configuration control over the way jobs are handled. They also include a number of improvements to support cloud storage by increasing tolerance for latency between the Avere cluster and a source or destination volume.

## Select Junction for New Job

The Avere Control Panel now allows you to select a namespace junction as the source for a new FlashMove or FlashMirror job. Before this change, administrators had to select the source core filer and export, and then specify the subdirectory to define the job source directory.

Now, you can select a junction in the New Data Management Job wizard instead of manually selecting and specifying the job source directory.

This change also integrates migration functionality into the Namespace page:

- Select a junction in the **Namespace** page and click the **Move/Mirror** button to create a new data management job with that junction as the source directory.
- Active data management jobs are shown on the junction details panel in the **Namespace** page.

In the **New Data Management Job** wizard, the first page allows you to select a namespace junction from any of the vservers in the cluster, or select **Custom source definition** to manually define the source.

If you select a junction, page two of the wizard pre-fills the data source fields and customizes the options available depending on whether the junction supports SMB access or not.

If you want to move data from a path that does not correspond exactly to the path of a junction, choose **Custom source definition**. Note that you must specify the correct SMB or non-SMB options for this source, since both options appear in the wizard. (The wizard does not prevent a user from selecting SMB-related options for an NFS-only source directory.)

The custom source option is similar to the behavior in older Avere OS versions.

## Node Selection

A **Primary node selection** field has been added to the **New Data Management Job** wizard to allow you to specify which cluster node handles the work for the move or mirror operation. The default setting allows the cluster to automatically select the node with the fewest data management jobs running, as was done in previous Avere OS releases. (21585)

Selecting a preferred node does not prevent other nodes from taking over the data management job if the selected node fails.

The screenshot shows the 'Add New Data Management Job' wizard. The 'Job Definition' section includes a 'Namespace junction' dropdown set to '/river-0-po' and a 'Job type' dropdown set to 'Mirror'. The 'Options' section includes: 'Initial state' (Stopped), 'Overwrite mode' (Always), 'Mirror synchronization policy' (Flexible (default)), 'Enable file logging' (unchecked), 'Sparse file checking' (Enabled (default)), and 'Primary node selection' (Auto (Node with fewest active jobs)). At the bottom, there is a 'Notes' field and 'Back' and 'Next' buttons.

*Updated Data Management Wizard in 4.7.3.1*

## Configurable Mirror Synchronization Policy

A new option for FlashMirror allows the administrator to configure whether or not the source and destination volumes can become out of sync. (22614)

If you set the Mirror synchronization policy to **Strict**, a failure writing a change to the destination core filer causes the system to retry the write indefinitely. Under this policy, no operation is considered complete until it has reached both the source and the destination core filer.

The other synchronization policy, **Flexible**, is the same behavior as in earlier software releases. If the destination becomes unavailable, the data management system keeps a record of changes that have not been written to the mirror and periodically attempts to repair the situation. Flexible is the default policy.

Each policy has advantages and disadvantages:

- Strict synchronization limits the difference between the primary storage system and its mirror. However, if the mirror volume becomes unavailable, it is possible for the Avere cache to fill with retrying sync operations and become unavailable for client operations.
- Flexible synchronization allows client operations to proceed unaffected by synchronization overhead, but can result in a larger discrepancy between the primary storage system and the mirror site if there are connectivity problems or other errors.

Neither policy affects the initial synchronization between the mirror source and destination. The policies only apply to updates after the mirror has been established and the content is synchronized.

### *Upgrade and Default Setting*

When cluster software is upgraded to Avere OS 4.7.3.1, existing FlashMirror jobs are set to the Flexible policy, which provides the same behavior as previous Avere OS releases. Flexible also is the default policy for new FlashMirror jobs.

## Jobs Proceed Despite Minor Errors

Changes have been made to the way FlashMove and FlashMirror operations handle errors so that the overall operation continues even if certain files or directories can't be moved.

Before this change, an error in the move or mirror migration process could cause the operation to pause and wait for manual intervention, or to stall while retrying the failed operation. Now, the problem is logged, but the file is skipped and the overall migration operation continues.

When the move or mirror job has transferred all of the files it could, but encountered errors on other files, the job's status is listed as paused instead of completed. In this state, all files have been scanned, but they have not been permanently copied to the destination. Administrators can view the error log and resolve the problems, then resume the job at its current state. Before this change, the job had to be stopped, which means that its progress is abandoned and all of the file transfers must be done again from the beginning.

If files are skipped, the Avere Control Panel dashboard shows an alert that includes a link to the error report that lists the problems. The Data Management page also has a link to the error log.

The system skips files only for common errors related to individual files or directories - for example, a conflicting filename, unexpected type, or noncompliant name. If problems are found with the system infrastructure, or the job is attempting to write to a full disk or read-only filesystem, the job pauses immediately. The job also pauses if more than 200 file errors have been logged. (20204)

## Parameters Can Be Modified on a Stopped Job

This release adds the ability to update certain settings on a data management job without aborting the job. The job must be stopped. Before this change, a FlashMove or FlashMirror job had to be aborted and recreated in order to change any of these settings.

The **Actions** menu on the Data Management tab now includes a **Modify** button for stopped jobs.

Settings that can be changed on a stopped job include:

- Overwrite mode
- Mirror synchronization policy
- File logging (enable/disable)
- Sparse file checking
- Primary node selection
- Notes
- SMB Administrator username

## Improved Flexibility for Moving or Copying Subdirectories

This release removes the restriction that made FlashMove/FlashMirror jobs unable to migrate a parent directory if a subdirectory had previously been migrated. Now, a directory can be the source or target of a data migration job regardless of its migration history. (14574)

## FlashMove/FlashMirror Job Stability Improvements

This release includes reworked procedures for handling errors encountered in a FlashMove or FlashMirror job. In addition to providing better recovery for commonly encountered problems, this change allows administrators to use the actions Abort, Stop, Reverse, Transition, and Modify during the job even if core filers or core filer exports are not reachable. (20745, 22288, 22589, 22792)

These changes reduce the likelihood that a FlashMove or FlashMirror job will become unresponsive.

As part of this work, several new alert messages were added to the system for data migration jobs. The system dashboard also might show new background tasks for data management cleanup.

## Accurate Item Count in Job Status Panel

The number of items listed in a Data Migration job detailed status panel now reflects the number of files and directories in the tree. Before this change, the item counter reflected the number of items copied, and it was possible for one file to be copied multiple times if it changed during the move. In the previous system, the status message could show 130 file moves for a directory containing 100 files, which caused user confusion. (21683)

## Job Status Icons

Status icons were added to the data management job list. The icons indicate the general status of the job:

- Green check mark - The job is proceeding normally.
- Yellow exclamation point - The job is proceeding but with some abnormalities. This icon might indicate an out-of-sync mirror, a job that has skipped files due to errors, or other recoverable situations.
- Red X - The job is paused or stopped.

Status icons do not appear for jobs that complete normally or are aborted.

## Command-Line Configuration

New XML-RPC calls have been added to Avere OS 4.7.3.1 for configuring these new features. Read [Data Management XML-RPC API Changes](#) for details.

## Cluster Security Enhancements

Several improvements to FXT administrative security are included in the v4.7.3.1 release. These improvements include updating third-party software packages that are used in Avere OS, as well as changing security-related settings to avoid remote denial of service (DoS) and cryptographic attacks.

Note that some of these changes might require software updates on the systems used to access the Avere Control Panel for configuring the cluster.

Avere OS version 4.7.3.1. includes upgrades and configuration changes to OpenSSL, OpenSSH, and Apache.



Make sure that systems used to administer the Avere cluster (through web access to the Avere Control Panel, or using the XML-RPC API) meet the security requirements described below. (These requirements do not affect client systems that access the cluster cache; they apply only to workstations that are used to configure the cluster.)

## SSH Ciphers

Note: The list of accepted cipher suites and HMAC digests was updated in version 4.7.3.2. Refer to the current list in the [SSH Security Update](#) item in the 4.7.3.2 section of this document.

## HTTPS Configuration

All HTTPS endpoints (used for web administration and XML-RPC API access, among others) now explicitly reject connections that use insecure legacy protocols and cipher suites.

Web browsers, OpenSSL, and Python installations on administrative workstations might need to be upgraded.

- Settings should correspond to the Modern TLS protocol described in these articles:
  - [https://wiki.mozilla.org/Security/Server\\_Side\\_TLS](https://wiki.mozilla.org/Security/Server_Side_TLS)
  - <https://mozilla.github.io/server-side-tls/ssl-config-generator/>

The Mozilla wiki page includes a [list of web browsers](#) that have the minimum support necessary for the Modern configuration.

## TLS Version 1.2

**Only** Transport Layer Security (TLS) protocol version 1.2 is accepted. Insecure legacy protocols, including SSL, TLS v1.0, and TLS v1.1 are rejected.

TLS v1.2 support was introduced in Python v 2.7.9, Python v 3.4, and OpenSSL v 1.0.1; however, Avere recommends upgrading administrative machines to the latest versions of Python and OpenSSL to include additional security updates and bug fixes.

## Apache Settings

The Apache web server now is configured to permit only the following security standards:

- Elliptic Curve Diffie-Hellman Exchange (ECDHE) key exchange
- ECDSA and RSA signature algorithms
- AES256, AES128 ciphers (including AES-GCM mode) with SHA256 or SHA384 message authentication
- CHACHA20 cipher with POLY1305 message authentication

## Browser Input Security

Changes were made to disable web browser autocomplete services for security-sensitive fields in the Avere Control Panel. After this change, browsers will not offer to cache passwords and other strings that could possibly lead to unwanted exposure.

## System Logging Enhancements

### API Event Logging

System logging procedures were overhauled to provide more useful and thorough event logging for Avere OS login events and XML-RPC calls. A larger set of XML-RPC calls are now logged, and log messages are categorized and secured appropriately. Filesystem events are still logged as in previous releases.

### Remote Logging Changes

This overhaul also allows administrators to configure the types of events that will be sent to the remote syslog server.

The Monitoring settings page in the Avere Control Panel now includes the following options:

- FXT filesystem events - Forwards FXT filesystem messages to the remote monitoring server. This option generates a high volume of log entries.
- OS login events - Forwards all login event messages to the remote monitoring server.

- XML-RPC events - Forwards detailed API call logs, including failures, to the remote monitoring server. XML-RPC events now include the username, remote IP address, method name and arguments (with sensitive values masked for security), and results. (Some routine XML-RPC events are not logged.)

All of these options are enabled by default when you specify a remote syslog server; most administrators should disable filesystem event logging to avoid excessive traffic.

### **Command-Line Configuration**

The remote syslog options also can be configured with the XML-RPC API. Details are in [Remote Syslog XML-RPC API](#).

### **Virtual FXT (vFXT) Administration Enhancements**

Avere OS version 4.7.3.1 includes significant enhancements to improve the speed and reliability of creating, destroying and administering Avere vFXT clusters in cloud computing environments.

The enhancements include the following:

- Improvements in the software for individual vFXT nodes that allows them to interact more smoothly with various cloud provider virtual compute environments.
- Improvements to the Avere Cluster Manager web administration tool
- The vfx.py Python API and command-line utility for controlling virtual FXT clusters programmatically

### **SMB Performance Improvements**

Several changes have been made to improve the efficiency of SMB interactions in version 4.7.3.1 of Avere OS. (SMB is sometimes called CIFS for historical reasons.)

### **SMB WAN Optimization**

Code changes were made to optimize Avere OS performance when accessing data from back-end core filers for SMB clients. Specifically:

- The number of requests between the Avere cluster and a core filer when requesting SMB ACLs has been reduced. These optimizations have cut the needed number of round trips to the core filer by half.
- Priorities were changed to ensure that write operations have priority over other types of transactions. Before this change, cloud metadata operations could cause bottlenecks when run across WAN links or with slow back-end filers. (19701)

### **Cache Storage Optimization**

Additional SMB support was added for Avere's "Always Forward" intelligent caching algorithm. This algorithm optimizes the use of cache storage space by designating a preferred node for specific files; other nodes forward client requests to the preferred node instead of caching an additional copy of the same file.

In 4.7.3.1, Avere has augmented Always Forward support for SMB operations.

- ACLs now can be cached at the cluster level, eliminating the need for individual nodes to retrieve ACLs from core filers.
- SMB SETINFO operations are now asynchronous to the core filer.
- SMB2 CREATE operations and SMB2 FIND operations are forwarded to improve performance.

### **SMB Shares Path Can Be Selected**

When adding a new CIFS share in the Avere Control Panel, the **Namespace path** field is now a dropdown list of defined namespace paths rather than simply "/". The **Subdirectory** field needs to be used only when mapping a share to a subdirectory of a namespace path. Previously, the user needed to manually type in a namespace path to the subdirectory field, which was non-intuitive and led to errors. (21913)

## XML-RPC API Changes

The following XML-RPC API calls were changed or added. Read the associated `averecmd --help` output for more details about each interface.

### New XML-RPC API Calls

- `corefiler.modifyAutoExcludeList`
- `monitoring.getSyslogServer`
- `monitoring.getSyslogSettings`
- `node.updateHardwareInfo`

### Data Management XML-RPC API

The following XML-RPC API calls were added for command-line configuration of Data Management features.

- `migration.getErrorReport`
- `migration.modify`

### Remote Syslog XML-RPC API

The XML-RPC call `monitoring.setSyslogServer` now accepts an optional struct parameter to choose which messages to forward by setting the following keys to True or False:

- `auth` (boolean) - Set to **False** to disable forwarding of OS login events.
- `xmlrpc` (boolean) - Set to **False** to disable forwarding of XML-RPC calls and failures.
- `filesystem` (boolean) - Set to **False** to disable forwarding of FXT filesystem messages (high volume).

To forward only cluster security-related events, use `{'auth': True, 'xmlrpc': True, 'filesystem': False}`

By default, all three classes are forwarded if a remote syslog server is configured.

The new XML-RPC call `monitoring.getSyslogSettings` was added to return the new settings.

The following XML-RPC calls are unmodified to maintain backwards compatibility:

- `monitoring.enableSyslogServer`
- `monitoring.syslogServer`
- `monitoring.syslogServerEnabled`
- `monitoring.testSyslog`

### Changed XML-RPC Calls

- `cert.generateCSR` - now supports an optional parameter named `keySize` to specify the length of the newly-generated private key (either 2048 or 4096). The default is 2048.
- `cifs.configure` - the `adminPass` argument is now optional
- `cluster.get` and `cluster.modify` now include an optional parameter named `cloudAdminCredential` in place of the old parameter name `ec2AdminCredential`, which is no longer supported
- `corefiler.create` and `corefiler.createCloudFiler` now accept an optional `filerNetwork` argument to specify which cluster network to use for communication with the new core filer being added

### Deprecated XML-RPC Calls

The calls in this section are no longer valid in Avere OS 4.7.3.1

- `analytics.createCacheReport`
- `analytics.genCacheReportPlot`
- `analytics.getCacheReport`
- `analytics.getCacheReportHistogram`
- `analytics.listCacheReports`
- The `LZ4HC` option for cloud object compression is no longer supported in Avere OS and the configuration option is no longer valid in the API.

# Resolved Issues

## Cloud Core Filer Object Store (FlashCloud)

- 20964 Fixed an issue in the Avere Control Panel that caused the primary cluster network to be selected when attempting to add a cloud core filer that uses a different cluster network.
- 21555 Fixed an issue that could cause a system restart while taking cloud snapshots.
- 22287 Limited callback log spam when modifying core filer objects.
- 22558 Fixed a problem that could lead to a filesystem restart when deleting large quantities of files on a cloud core filer. This change reduces memory resource contention by leaving some of the memory used for filesystem metadata unpinned.
- 22800 Removed the LZ4HC compression option for cloud core filers because enabling it incurs significant performance impact. Any existing core filers with this compression option selected keep the option when upgraded to 4.7.3.1, but LZ4HC cannot be re-selected if removed.
- 23046 Fixed a filesystem restart caused by a filesystem process out-of-memory error involving cloud core filers.
- 23063 Improved snapshot cleanup to limit the number of core filers cleaned at one time. This change also ensured that cleanup activity can be stopped immediately by Avere Global Services during troubleshooting.
- 23110 Fixed a problem that could cause the cluster to ignore the cloud metadata database and attempt to read in a non-existent object from the bucket. The error could cause valid file metadata to be reported as stale.

## Filesystem

- 20931 Fixed a race condition involving client IDs in the HA cache that could cause a filesystem restart.
- 21072 Fixed a filesystem restart caused by a deadlock in multiple threads waiting for the same locks.
- 21233 Code was fixed to properly handle IO errors from core filers without creating alerts on the system Dashboard.
- 21339 Memory leaks were detected in testing and have been repaired.
- 22177 Reworked code to more efficiently maintain the queue of available SMB ACLs and to limit the size of the ACL memory cache.
- 22477 Fixed a race condition observed on systems configured for a cache policy of "Clients Bypassing the Cluster" (or similar) where racing write and getattr operations for a file not in the cache could cause stale file attributes to be cached, incorrectly.
- 22515 Changed procedures to ensure success returning ACLs for commands involving snapshot objects.
- 22777 Fixed a filesystem restart caused by a deadlock in cache operations.
- 22785 Fixed a condition where a filehandle pointer can become null.
- 22924 Fixed a race condition in the directory name lookup cache code that could cause a restart.
- 22894 Changed code to allow data to be discarded from the cluster cache after changes on the backend core filer have altered the root export filehandle or other information. Before this change, the FXT cluster was unable to flush some data back to the core filer; now, it logs the affected files' information and then discards the changes.

## FlashMove/FlashMirror

- 22288 Added logic to prevent a restart caused by core filer ID errors when a mirror is reversed.

## General

16658	Prevented multiple spurious “HAVoter ... missing file” messages from being written to internal log files.
19190	The Management VLAN setting now can be successfully changed on the Avere Control Panel Administrative Network settings page.
21021	Fixed a bug in the alert notification email client code that was resulting in certain alerts not being properly sent via email.
21570	Increased system buffer resource limits to handle a higher quantity of in-flight filesystem operations.
22639	Increased the starting partition size for system storage and made it configurable to avoid issues supporting large hardware clusters.
22723	Fixed a race condition in startup procedures that could cause a filesystem process restart.
22207 22693 22754 22788	Several longstanding performance and reliability issues with cluster status reporting and generation of Conditions and Alerts visible in the Dashboard page were resolved. The cluster now reports events more efficiently, eliminating repeated and unnecessary alerts and conditions.

## SMB/CIFS

19701	Improved performance when setting or retrieving SMB ACLs.
20718	Enhanced performance of file replication procedures when ACLs are copied by improving efficiency in creating NTFS ACL metadata.

## vFXT

20930	Improved the robustness of the system for adding vFXT nodes to existing clusters. Before this change, customers could see an error when trying to add nodes to a busy cluster.
22847	Updated Avere OS to recognize Google Compute Engine vFXTs after Google changed the string returned in response to an SMBIOS query. Avere OS uses the product name portion of the string to configure provider-specific settings and to ensure that unsupported environments are not used.

## Additional Fixes in 4.7.3.1

This table lists internal tracking numbers for problems fixed in Avere OS version 4.7.3.1. If you have a support case that references one of the tracking numbers listed here and would like more information about the changes that were made to address your issue, please contact Avere Global Services.

19526	21075	21952	22298	22559	22734	22829	22934	23036
20687	21235	22005	22478	22566	22743	22835	22937	23057
20747	21388	22030	22502	22579	22745	22846	22943	23064
20823	21529	22047	22516	22643	22774	22878	22978	23074
20829	21696	22197	22521	22688	22790	22881	22986	23124
20929	21812	22207	22528	22717	22792	22913	23010	23152
21070	21825	22261	22545	22725	22826	22915	23027	

# Contact Support - Avere Global Services

Support can be reached by web, phone, or email.

**By web:** <http://www.averesystems.com/support>

**By phone:**

1-888-88-AVERE, Option 2 (Toll-Free)

1-412-894-2570, Option 2

**By email:** [support@averesystems.com](mailto:support@averesystems.com)