



Avere OS 4.8.4.3.C1 Release Notes

2018-05-14

Table of Contents

New in Avere OS 4.8.4.3.C1

New in Avere OS 4.8.4.3

New in Avere OS 4.8.4.2

New in Avere OS 4.8.4.1

New Features and Enhancements

FXT 5850 Edge Filer

Add Multiple Nodes at Once

Filters for Hot Files View

SMB2 Support

SMB Security Improvements

Express Snapshots

Snapshot Style Comparison

Snapshots Can Be Taken on FlashMirror Targets

Resolved Issues

Cloud Object Store

Filesystem

FlashMove/FlashMirror

General

NFS

Security

SMB/CIFS

vFXT

Contact Support - Avere Global Services

Copyright Information

Copyright © 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.

New in Avere OS 4.8.4.3.C1

This release includes two critical fixes:

- 25482 – Fixed a system restart that could occur after renaming files multiple times
- 25614 – Prevent a system panic loop during startup

Avere Systems recommends that all customers install this update.

New in Avere OS 4.8.4.3

Avere OS 4.8.4.3 includes one bug fix, which corrects an issue that could cause a system restart when attempting to remove a directory on a cloud core filer that contains more than 200 empty segments. (25373)

New in Avere OS 4.8.4.2

This release includes bug fixes and other enhancements.

Resolved Issues

Filesystem

25544 A logic error was fixed that could cause a system restart after the system failed to get a filehandle.

General

25497 Fixed an error in recently added code that could create excessive spurious alerts. This defect was discovered in internal testing, but can possibly affect any cluster running Avere OS 4.8.4.1.

New in Avere OS 4.8.4.1

New Features and Enhancements

FXT 5850 Edge Filer

This release supports the new Avere FXT 5850 Edge Filer. The FXT 5850 features all Flash storage and has double the DRAM and storage capacity and more than twice the network bandwidth of other FXT 5000 Series models. Read <http://www.averesystems.com/products/fxt-5850> for more details.

Note that Avere OS version 4.8.4.1 is the minimum version that supports the FXT 5850 model. If you want to add an FXT 5850 Edge Filer to an existing cluster, you must first upgrade the cluster's Avere OS

software to version 4.8.4.1 or higher. Earlier software does not recognize an FXT 5850 as an eligible cluster node.

Add Multiple Nodes at Once

Changes made in this release allow you to add two or more nodes to an Avere cluster at one time. The XML-RPC command `node.allowToJoin` now accepts a comma-separated list of node names to support situations where you want to add more than a single node.

This change applies both to FXT Series hardware clusters and to vFXT systems.

The option is not yet available through the Avere Control Panel. Contact Avere Global Services if you need help accessing or using the XML-RPC API. (24421)

Filters for Hot Files View

Filters have been added to the **Hot Files** table on the **Dashboard** tab of the **Avere Control Panel**. The filters allow administrators to set a minimum size or rate for determining which files appear in the list.

The screenshot shows the 'Hot Files' tab selected in the Avere Control Panel. The 'Show' menu is set to 'Ops'. The 'Show at least' field is set to '0 Ops/Sec'. The 'Cluster-wide view' and 'Auto-refresh' checkboxes are unchecked. The table below has columns for Name, Type, FXT Node, and Ops/sec, and currently displays 'No active files'.

The filters are different depending on which option is selected in the **Show** menu.

- For **Operations**, you can filter by a minimum number of operations per second.
- For **Directory Updates**, you can set a minimum number of updates per second.
- For **Bytes Read/Written**, you can set filters for the minimum number of bytes read and for the minimum number of bytes written, and combine the two options with **and** or **or**.

The screenshot shows the 'Hot Files' tab selected in the Avere Control Panel. The 'Show' menu is set to 'Bytes Read/Written'. The 'Show at least' field is set to '0 Byte Reads'. The 'and' radio button is selected, and the 'Show at least' field is set to '0 Byte Writes'. The 'Cluster-wide view' and 'Auto-refresh' checkboxes are unchecked. The table below has columns for Name, Type, FXT Node, Reads, and Writes, and currently displays 'No active files'.

Enter the minimum value in the filter and click **Enter** to see the filtered result. (23555)

SMB2 Support

The Avere cluster now can use SMB2 to communicate with Active Directory domain controllers and several types of NAS core filers. SMB2 is supported on NetApp filers (clustered and non-clustered mode) and Isilon storage systems. Contact Avere Global Services for details about which systems are supported.

SMB1 is still available for all supported core filers. Consult your core filer's documentation to learn how to enable or disable SMB1 and SMB2. (23898)

SMB Security Improvements

Several changes were made to allow administrators to customize security settings in SMB operations.

Two changes were made to allow administrators to disable insecure authentication methods:

- NTLM-encrypted password authentication for users now can be disabled. If NTLM encrypted password authentication is disabled, then only NTLMv2 responses are accepted from clients. (25250)
- NTLMSSP authentication from an Avere cluster to AD servers and NAS core filers now can be disabled. (25244)

These settings can be modified by using the Avere Control Panel or the XML-RPC configuration method `cifs.setOptions`.

An additional change adds SMB server signing configuration to the Avere Control Panel web interface. The **VServer > CIFS** settings page now includes a drop-down control for selecting `auto`, `mandatory`, or `disabled`. (25251)

Express Snapshots

A new setting is available for cloud core filer snapshots. *Express* snapshots perform a snapshot of the cloud core filer contents without flushing all modified data in cache to the core filer. That is, only data that has been written to the cloud bucket is included in the snapshot; recently changed files that are only stored in the Avere cache will not be included.

Express snapshots have a minimal impact on client performance, but they also do not guarantee that the snapshot represents the latest changes.

Standard snapshots are now known as *strict* snapshots. In strict snapshots, the cache is fully flushed to ensure that all data is consistent between the cache and the cloud core filer before the snapshot is taken.

The table below highlights the differences in the two snapshot styles.

Snapshot Style Comparison

Express Snapshot	Strict Snapshot
Snapshot is taken without flushing modified data	All modified data are flushed to the core filer before the snapshot
Minimal impact on client performance	Client performance can slow during snapshot processes
Snapshots can be taken more frequently and under higher load	Snapshots should be scheduled so they do not coincide with periods of high load
Not appropriate if strict point-in-time consistency is needed	Guarantees strict point-in-time data consistency for sensitive applications or database use
Suitable for archiving work load	Suitable for incremental backups with requirement of point of time consistency

A new field labeled **Snapshot Type** appears when creating a cloud snapshot or snapshot policy in the Avere Control Panel. You can choose **Express** or **Strict** (traditional style). Note that you cannot schedule both types of snapshots simultaneously on the same core filer.

Snapshots that were created before the upgrade continue are unchanged and continue as strict snapshots.

Snapshots Can Be Taken on FlashMirror Targets

You can now create snapshots of a cloud core filer that is the destination of a FlashMirror job. Snapshots can run any time the job is active and in a synchronized state. (24912)

Resolved Issues

Cloud Object Store

- 22186 Cloud object store versions numbers are now kept in a replicated persistent database to help guarantee that the latest version is being accessed.
- 22541 Reworked thread creation strategies to prevent a system restart caused by thread memory issues.
- 22947 Corrected a flaw that caused temporary overloads in memory allocations and could lead to filesystem restarts.
- 23324 Modified the strategy for TCP/IP communication with core filers to improve management of multiple cloud core filers.
- 23802 Improved efficiency in parsing the large snapshot database object to avoid system restarts. The snapshot database can grow very large when more than a thousand scheduled snapshots are created and deleted over time. This change also ensures that the snapshot object is only as large as the number of active, undeleted snapshots.
- 23891 Prevents stuck read/write operations to files that exceeded a certain size.
- 23953 Ensure that the correct setting is used for the allowed number of TCP connections to a cloud core filer.
- 24057 Sets a maximum time limit for removing an empty cloud core filer when the "force" option is used. Before this fix, such operations could sometimes fail.
- 24065 Ensures that failover takes place quickly even if a time-consuming file operation is in progress. Before this change, if a node or service failure occurred while a large file was being removed or truncated, the failover was delayed until the file operation completed.
- 24447 This change improves read throughput performance in a mixed read/write workload. Before this, cloud core filer throughput could be poor if a client that was already reading data also began writing.
- 24591 You now can modify the number of TCP connections to a cloud core filer without needing to restart the system.
- 24702 Improved the performance of read operations from cloud core filers while writes are going on.
- 25155, 25163 Changed the Avere Control Panel so that it does not allow a snapshot to be cancelled when the snapshot is in the process of being written to the core filer.
- 25400 Change the default object size ("otherSegmentBytes") for cloud core filers to 8 MB when creating a vFXT cluster and core filer in the same cloud provider.
- 25452 This change ensures that express snapshots always complete. This update fixes an issue that could cause express snapshots to get stuck in a way that also stopped progress on all operations to the affected cloud core filer. (In that situation, the snapshot had to be cancelled and processes restarted.)

Filesystem

- 19180 Fixed a rare race condition that affected setting the failover version while the file system process was restarting.
- 21116 Fixed a filesystem service restart that could occur when a node was removed from a cluster under load.
- 22165 Suppressed repeated "service not found" status messages during cluster rebalance.
- 22750 Corrected a flaw that caused core filer cache policy changes and invalidations to get stuck.
- 23149 Fixed a rare crash in the directory cache.
- 23309 Added the ability to force a flush operation that stopped because it attempted to remove a non-empty directory from the core filer. The alert message on the Avere Control Panel Dashboard now includes information about how to determine whether or not forcing the deletion is safe, and contains a link to the Support page, which shows buttons for retrying the flush operation or forcing it to complete.
- 23407 Fixed an issue that could cause an error message about Local Directories to appear on systems that did not have any core filers using Local Directories.
- 24287 Corrected a flaw in automated recovery from stuck directory removal processes. Before this change, excessive retries could be generated if a "no entity" message was returned when removing a directory as part of a flush.
- 23530 Improved performance by creating a more efficient process for scanning filesystem buffers to determine what needs to be written to the core filer.
- 23553 Fixed a rare problem that could cause cache filesystem errors after deleting certain files.
- 23718 Limits the number of retries allowed when flushing a cached metadata operation. This change allows the system to move on to other operations and return to the problematic task later. Before this change, it was possible for a create operation to stall because space was unavailable and block a remove operation that would free the required space.
- 23725 Fixed a bug where removing a core filer caused errors with later operations to valid junctions.
- 23771 Added XML-RPC commands to rebalance metadata servers (Directory Managers) onto a subset of cluster nodes. This change improves efficiency when removing or replacing multiple nodes in a cluster, because you can rebalance the nodes that will remain in the cluster ahead of time instead of rebalancing the cluster after each removal.
- 23783 This change fixed a bug where certain operations are not woken up properly, causing the cloud snapshot process not to make progress.
- 23885 Corrected a race condition during shutdown of metadata caching that occasionally caused filesystem service crashes during failover.
- 23985 Fixed a problem that could cause a system restart when pruning large directories to free space. This issue only affected systems with Local Directories enabled.
- 23997 Fixed a bad assert that could cause a system restart.
- 24071 Improved efficiency by assessing file metadata operations and not pushing them to backend storage if no net change was made.
- 24080, 24268 Added a feature to track operations by user and include details in data dump reports.
- 24153 Corrected thread locking to prevent a restart in filesystem HA services.
- 24253 Resolved a problem that caused the cluster statistics collection process to sometimes hold a user-level thread lock while reading from disk, which could cause a watchdog-related system restart.
- 24322 Made identifier checks more rigorous to resolve a customer issue where an invalid UUID prevented node removal.
- 24332 Fixed a race condition in byte-range locks used for read-ahead operations.
- 24346 Fixed a system restart during high load after accessing more data than fits into memory.

- 24379 Reduced the amount of time that a lock is held during a node removal to increase concurrency.
- 24404 Improved efficiency of token revokes to avoid high CPU load.
- 24509 Removed an unneeded alert about an NFS barrier during writeback operations. This alert message was shown during routine system and service restart processes, and did not automatically clear.
- 24557 Reduced the startup time for a vFXT cluster by removing an unneeded nvram process.
- 24630 Fixed a race condition that could cause the filesystem process to restart when internode communication is timing out.
- 24655 Fixed a memory leak in the directory manager module.
- 24673 Fixed a rare filesystem restart that was caused by an internal filesystem walker that exited without setting the correct flag.
- 24674 Added a mechanism to limit memory consumption by adjusting token numbers on a system with high load.
- 24895 Removed an inefficiency in handling cache lock releases. Before this change, a system restart could sometimes occur while writing a lot of data to a large file.
- 25036 Improved the display of hot files read and write statistics to show 0 if no operations have taken place since the system started.
- 25041 Updated token management processes to grant tokens more quickly.
- 25176 Fixed a race condition in HA transaction operations.
- 25362 This change fixed latency in client operations that was caused by a socket error in inter-node communication that affected token management.
- 25417 Fixed an uncommon system race condition that could possibly cause hung client write calls if writing a single very large file.
- 25454 Updated code to better handle token cleanup after removing a core filer.
- 25471 Some Linux core filers incorrectly return an EXDEV error for a rename operation. This change allows the Avere cluster to ignore such errors and do a quick retry instead, which typically resolves the problem.

FlashMove/FlashMirror

- 20757 Improved performance and reduced the duration of client outages during the transition from source to destination in a data management job.
- 23875 Fixed a system incompatibility that prevented users from creating data management jobs if the target was a NetApp core filer using Clustered Data ONTAP.
- 24064 Enabled the option to transfer source export policies to destinations that are subdirectories of a junction. If the destination is not on a core filer that supports hierarchical exports, or if the subdirectory is not also a mount point, the export policy is not transferred.
- 24882 Fixed a somewhat rare race condition that could cause a main filesystem process restart during post-processing after completing a data management job.
- 24917 Added a custom setting for ignoring the nanosecond value in file mtime and atime values. This setting can be used when testing with rsync to work around an rsync bug that inconsistently updates the nanosecond portion of mtime and atime metadata.
- 24928 Instituted a migration pre-check that verifies that the namespaces involved in the migration are not damaged before the job starts to copy data.
- 25129 Fixed a race condition that could cause a system restart when initializing a data management job.

- 25422 Fixed an issue where aborting a migration and then very quickly attempting to use the core filer export again could sometimes result in the core filer export's being inaccessible to the Avere cluster.
- 25430 In a FlashMove or FlashMirror job where the destination has existing data and the **Overwrite** mode is set to **Date/time**, it's possible to have files that are hard-linked on the source, but not on the destination. This change enables the mover to fix the hard links on the destination so that they match the source.

General

- 18045 Added the `keyMgmt.testKnipServer()` call to the XML-RPC API.
- 20382 Fixed a problem that caused the rate displayed in the dashboard Hot Clients tab to intermittently show as 0 ops/second.
- 22468 Prevent a spurious error message about an unhealthy state related to a forbidden token manager move.
- 23193 Prevent excessive logging of netgroup host triples when netgroups are downloaded from an external NIS server.
- 23321 Fixed an error that could cause a spurious confirmation dialog to pop up when leaving the Manage Core Filers settings page after adding or removing a core filer. The dialog incorrectly warned that changes might not be saved, even if the operation had already completed.
- 23364 Improved the memory buffer cache flushing to prevent system restarts
- 23445 Fixed input verification of the **admin** user's SSH login key when using the command-line Maintenance Menu to set up a cluster. Before this change, the menu interface suggested a default value of `[none]` for the SSH user key, but would not accept that value.
- 23446 Fixed a bug in certificate name matching that prevented administrators from deleting SSL certificates by using the Avere Control Panel.
- 23447 Added the ability to choose an SSL certificate for the cluster on the Avere Control Panel. The **Cluster > General Setup** page now includes a drop-down Cluster SSL Certificate menu populated with the certificates that have been added to cluster. (Define certificates on the **Cluster > Certificates** settings page.)
- 23639 Fixed a multi-process race condition that could lead to deadlocked management processes. This management deadlock had various symptoms, including loss of filesystems services for public cloud object stores when credentials expired and could not be renewed.
- 23767 Improved network performance and solved a link aggregation imbalance problem by enabling additional queues and software interrupt modulation in the network driver.
- 23792, 22763 Avere OS now recognizes additional valid headers on certificate PEM key files.
- 23830 Reduced memory usage by detaching unjoined pthreads.
- 23846 Starting with Avere OS version 4.7, TLS1.2 and strong HTTPs ciphers are required for remote connections. The XML-RPC client now explicitly checks for the Python minimum version of 2.7.9 and the OpenSSL minimum version v1.0.1g, because those are the minimum required versions to support these security features.
- 23867 Fixed default internal cache sizes, which were being set incorrectly in some situations. Before this change, reduced performance could be seen in the Access Control List (ACL) cache and inode cache.
- 23926 The `cluster.modifyCloudRegion` RPC now requires the IAM service endpoint to be provided when registering a custom region.
- 23983 Boot script resilience was increased to allow software to start in the presence of some non-fatal hardware errors.
- 24053 Fixed a problem that could cause the wrong information to show in the SMB user to NFS user overrides setting (the correct information was stored even though it wasn't shown).

- 24191 Changed code to prevent internal tracing files from being written out if the management process crashes.
- 24334 Corrected the method for initializing disk statistics. Before this change, incorrect initialization could sometimes cause invalid statistics values to be reported.
- 24353 Adjusted logic to prohibit certain internal operations based on system state.
- 24416 Fixed a bug that prevented packet captures from being run from the Avere Control Panel **Support** tab when "capture filter" was set.
- 24459 Adds enforcement of the maximum supported and tested cluster size of 24 nodes
- 25115 The text for specifying a new cloud core filer credential when adding a new core filer was changed to "add a credential set" instead of "create a credential" to avoid user confusion. This interface allows you specify a credential that already exists, but does not create a new one.
- 25299 Added internal instrumentation to the Avere OS management service (mgmtd) to keep track of outstanding requests in one of the subsystems.
- 24004 Enables all assigned cluster addresses during a node reboot or full service restart, instead of only enabling the first cluster address on a node. This fix shortens the service failover period, which was sometimes prolonged by active cluster nodes' trying to communicate with the unavailable addresses. VServer and management address activations are still deferred until cluster services are reestablished.

NFS

- 24372 This fix prevents a system restart caused by referencing a particular internal data structure after its associated memory was freed.
- 24893 This fix allows the Avere cluster to use the '/' export in mount requests to NetApp core filers even though this export is not reported by the NetApp filer.

NetApp Clustered Data ONTAP releases starting with version 8.3 contain a new "showmount" feature that lists all of the ONTAP junctions. NetApp Clustered Data ONTAP releases starting with version 9.2 cause showmount to not report the '/' export, effectively blocking the use of these systems as core filers for an Avere cluster. This change works around that limitation to enable the use of these NetApp systems.

When using a clustered Data ONTAP core filer, make sure of the following things:

- The core filer definition in the Avere cluster must have its **Filer Class** field set to **NetappClustered**. You can check or update the setting in the **Core Filer Details** settings page.
- The showmount feature must be disabled on the NetApp filer. Use this configuration command on your NetApp clustered Data ONTAP system:

```
nfs server modify -vserver ${VSERVER} -showmount disabled
```

This change **does not** add support for the clustered Data ONTAP showmount feature.

If the showmount feature is accidentally enabled, this change blocks the use of the additional exports in the showmount response.

- 25160 NFS Kerberos session setup accesses keytabs in the local filesystem, which can cause extreme latency for NFS client operations. The latency is also visible as core filer latency. This change prevents the keytab accesses from interfering with normal operation processing.

25195 This change addresses an Avere filesystem service restart loop caused by lack of disk space in the support partition when snapshots are scanned for use by the SMB **Previous Versions** feature.

Security

23453 This fix restricts NTP so that it can run only control messages coming from nodes within the cluster.

24056 Updated an on-box debugging utility to not display user credentials when in verbose mode.

25030 This change ensures that self-signed SSL certificates are checked for expiration dates.

SMB/CIFS

19699 This change automatically sets the SMB "archive extended" attribute for newly created files. (It does not change the default behavior for files that were renamed, written, or have changed size.)

SMB servers typically set the "archive extended" attribute for newly created files, files that are renamed, files that are written, and files that change size.

Prior to this change, the Avere SMB server did not set the extended attribute in these cases. This behavior is consistent for both NAS and cloud core filers.

Some applications (such as some backup applications) rely on the archive bit in order to determine whether a file must be copied again. A better practice is to perform a timestamp comparison rather than an archive bit check.

A custom setting is available to cause Avere OS to set the archive bit for all of these file changes. Enabling this setting has a performance impact. Contact Avere Global Services for more details about this setting.

23882 This change improves the speed of SMB directory listings by pre-fetching multiple ACLs from a directory at once after a filesystem service restart. This feature is used for NAS core filers only. It is enabled by default.

24866 Fixed a bug that could cause a filesystem service restart when interpreting filesystem symlinks. The problem involved SMB and affected systems using read caching and core filer verification.

vFXT

23812 Fixed an open socket issue that prevented a vFXT instance from rebooting correctly.

24333 Updated code to ensure that cloud credentials can be refreshed during long-running cloud service processes.

24457 Updated vFXT creation code ensures that AWS tags are consistently applied to all components (disks and instances).

25116 Updated validation strategies for the administrative network configuration to allow an empty DNS domain.

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