



Avere OS 5.1.3.1 Release Notes

2019-03-12

Table of Contents

New in Avere OS 5.1.3.1	1
Resolved issues	1
Cloud object store	1
Filesystem	1
General	2
NFS	2
SMB/CIFS	3
New in Avere OS 5.1.1.2	4
Security updates	4
SSL and TLS updates	4
Permitted cipher suites	4
Update private object storage to maintain compatibility	4
SSH server updates	5
Resolved issues	6
Cloud object store	6
Filesystem	7
General	8
Migration	9
NFS	10
Security	11
SMB/CIFS	11
vFXT	12
Contact Microsoft Customer Service and Support	13

Copyright Information

Copyright © 2019 Microsoft Corporation. All rights reserved.
Specifications are 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 5.1.3.1

This release includes several bug fixes. It is applicable for all customers.

Resolved issues

Cloud object store

- 25845 Fixed a defect in a cache file system database that could cause a process restart in the system that writes data to the cache disks.
- 26556 Changed code to avoid an out-of-memory condition that could cause an assert to be tripped under heavy load.
- 26754 Removed a rare race condition that could sometimes result in objects being written to cloud storage with invalid names.
- 26756 Increased the frequency of checks for upcoming cloud credential expirations.

Filesystem

- 25047 Fixed an error that could cause the filesystem to restart because a filesystem service got into an unexpected state.
- 25963 Eliminated a reference leak on a filesystem object, which could cause a system panic during a restart.
- 26255 Fixed a race condition in the directory name lookup cache component, which could cause a system restart under heavy file recycling conditions.
- 26532 Changed FlashMirror code to ensure that all files that are marked out-of-sync on both the primary node and the mirror node are copied if there is a failover event.
- 26640 Divided high availability file consistency checking after an HA write failure into multiple chunks to avoid a system restart because of the watchdog timer.
- 26654 Fixed an issue that could cause an incorrect “access” error to be returned for a core filer that has the optional extended user credential groups feature enabled. The problem was caused when the target directory did not yet have proper tokens.
- 26678 Moved flush request processing to separate tasks to prevent a possible process restart.
- 26796 Sped up the file scan process to reduce the time needed for HA file validation after cluster startup.
- 26852 Updated code to prevent a system restart that could happen if an optional directory repair tool is used on a directory at the same time the directory is being removed.
- 26891 Fixed an issue that caused the filesystem to restart when it tried to process an operation that required metadata files and logs to be opened and ready.

General

- 26604 Corrected a problem that created invalid support uploads because a placeholder cluster name from the Avere Control Panel was incorrectly passed as the unique cluster name. Always set a unique cluster name that will identify your cluster to support personnel.

- 26666 The network driver for vFXT nodes in Microsoft Azure has been changed to handle network hardware resets better. These resets occur as part of regular maintenance of Azure servers, such as network card firmware updates. Before this change, a network hardware reset could cause the vFXT node to lose network connectivity and need to be rebooted. After the fix, the driver handles the reset transparently and no action should be required from the user.

- 26723 Fixed a problem where bootloader messages and prompts were not being shown on the serial-over-LAN console for some 2000-series, 3000-series and 4000-series FXT models.

- 26854 Changed behavior so that when a hardware cluster node is wiped, the previously set password is preserved. Before this change, a node that was restored to factory settings did not have a node management password until after it joined a cluster.

- 26856 Fixed a connectivity problem in clusters with VLANs configured to use alternate default routers. The problem was caused by the system incorrectly assigning IPs and routes for the affected VLANs.

- 26870 Fixed an issue with a cluster recovery tool that put the cluster in a state requiring engineering intervention to complete the recovery.

- 26883 Fixed a regression that caused the serial port to connect at 9600 baud instead of 115200 baud. The issue was introduced in Avere OS 5.1.1.2 because of an upgrade to FreeBSD 11.2.

NFS

- 24155 A change was made to avoid timeouts in cache-to-cache communication among cluster nodes by overriding custom settings that exceed internal connection resources.

- 26753 Addressed a filesystem service restart and associated core file during upgrade to V4.8.9.1 or later release from a release prior to V4.8.9.1. The defect involved Open Network Computing Remote Procedure Call (ONCRPC) communication among nodes within the cluster.

SMB/CIFS

26220 A defect was fixed that involved CIFS ACL data management jobs (both FlashMove and FlashMirror) from cloud core filers. The defect could cause the data management job to fail, or cause a filesystem service restart (with a core file).

26522 The outbound NTLMSSP default setting for new vservers created with Avere OS version 5.1.2.1 is different from the setting for the previous release (5.1.1.2).

The current default values are:

- `disable_outbound_ntlmssp` - yes
- `ntlm_auth` - no

The previous release had `ntlm_auth` set to yes.

Administrators might choose to enable outbound NTLMSSP to work around problems with Kerberos domain join. Consider security risks before changing this setting. The setting can be changed by using the XML-RPC command method `cifs.setOptions`.

26567 Addressed two defects involving DNS Service Resource Record lookups for Microsoft AD domains. The first defect involved not switching to TCP after UDP lookups failed due to large responses. The second defect involved failed lookups using EDNS against Windows 2008R2 DNS servers.

26648 Added a custom setting to avoid contacting domain controllers that do not permit access. The failed accesses resulted in log spam for the cluster.

Contact Microsoft Customer Service and Support for details about how to apply the custom setting in the winbind daemon.

26658 Fixed a defect that could cause a vserver to fail to join an AD domain when the domain controller is configured with `RequireSecuritySignature` set to 0. If `RequireSecuritySignature` is set to 1, domain joins can succeed, but customers must verify that all devices joined to the domain are configured to support signing with domain controllers.

The workaround for this defect is to disable this particular signing path, however, this should only be done if the security risks are considered.

Contact Microsoft Customer Service and Support for more information and to learn how to apply the custom setting.

New in Avere OS 5.1.1.2

Avere OS 5.1 includes bug fixes, feature improvements, and significant security enhancements.

Security updates

This release improves security standards in several areas in order to comply with Microsoft security requirements.

Additional changes and bug fixes are documented in the [Resolved issues - Security](#) section later in this document.

SSL and TLS updates

Avere OS now requires the following standards for SSL and TLS:

- TLS1.2 must be enabled
- SSL V2 and V3 must be disabled

TLS1.0 and TLS1.1 may be used for backward compatibility with private object stores; contact Microsoft Customer Service and Support for details.

Permitted cipher suites

Microsoft security requirements permit the following TLS cipher suites to be negotiated:

- ECDHE-ECDSA-AES128-GCM-SHA256
- ECDHE-ECDSA-AES256-GCM-SHA384
- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-RSA-AES256-GCM-SHA384
- ECDHE-ECDSA-AES128-SHA256
- ECDHE-ECDSA-AES256-SHA384
- ECDHE-RSA-AES128-SHA256
- ECDHE-RSA-AES256-SHA384

The cluster administrative HTTPS interface (used for the Avere Control Panel web GUI and administrative RPC connections) supports only the above cipher suites and TLS1.2. No other protocols or cipher suites are supported when connecting to the administrative interface.

Update private object storage to maintain compatibility

The SSL and TLS changes in this version can impact customers using private object storage cloud core filers that are not updated to the current standard.

If the encryption software on your private cloud core filer is incompatible with the above [cipher suites](#) and TLS 1.2, you might lose access to the core filer.

CAUTION

Your Avere cluster might lose connectivity to your private object store core filer after installing this update if the core filer does not use compatible encryption software. Contact Microsoft Customer Service and Support before installing Avere OS 5.1 to avoid possible service disruption.

For compatibility with this version of Avere OS, your private object store cloud core filer must meet these requirements:

- TLS 1.2 must be enabled
- SSL V2 and V3 must be disabled
- The TLS cipher suites in the [permitted cipher suites](#) list above must be supported

Consult your private object storage vendor to confirm that it complies with these standards. Take any additional necessary steps to ensure that the private object environment is correctly updated and configured.

If your core filer cannot be brought into compliance, there is a temporary workaround available through Microsoft Customer Service and Support. Remember to remove this workaround after your core filers have been updated.

In general, it is a good practice to update and audit your private object core filers to maintain security standards.

SSH server updates

To meet Microsoft security standards, the security of the cluster node SSH server has been tightened.

Remote login as the superuser “root” has been disabled. If remote SSH access is required under the guidance of Microsoft Customer Service and Support, you must log in as the SSH “admin” user, which has a restricted shell.

The following SSH cipher suites are available on the cluster SSH server. Make sure that any client that uses SSH to connect to the cluster has up-to-date software that meets these standards.

Ciphers	aes256-gcm@openssh.com aes128-gcm@openssh.com aes256-ctr aes128-ctr
MACs	hmac-sha2-512-etm@openssh.com hmac-sha2-256-etm@openssh.com hmac-sha2-512 hmac-sha2-256
KEX Algorithms	ecdh-sha2-nistp521 ecdh-sha2-nistp384 ecdh-sha2-nistp256 diffie-hellman-group-exchange-sha256

Resolved issues

Cloud object store

- 23802 Improved management of the cloud snapshot database to avoid system panics. Before this change, the snapshot database object could grow very large, even though the number of actively needed snapshots was small. This fix improves database parsing efficiency and also ensures that the snapshot object file is not larger than the number of undeleted snapshots.
- 25227 Fixed an incorrect alert that suggested a cloud bucket or container could not be found after an Avere OS upgrade.
- 25717 Reduced the number of log entries from the certificate verifier component.
- 25837 Added automation to some cloud core filer modifications so that they no longer require service restarts.
- 25845 Fixed a bug in a database that could inappropriately stop the process that writes data to the cache disks.
- 25913 Improved performance of file creates and file lookups in directories by adding bloom filter support.
- 25992 Changed the way Avere OS handles a REST parse error for non-data objects from AWS storage. Now it returns a jukebox code instead of returning I/O error.
- 25995 Fixed a problem with snapshot identifiers to ensure that only the latest snapshot ID is used.
- 26157 Implemented a timer to fix timeout issues with HA barriers.
- 26159 Added a timer to reduce TCP connections on overwhelmed Amplidata core filers.
- 26179 Reduced memory consumption by improving memory management in the cloud module memory system.
- 26187 Fixed two issues related to certificate checking.
- 26235 Fixed a problem that could sometimes cause a system restart because flags were set incorrectly during express snapshot creation.
- 26338 Eliminated an issue that could cause a system restart when reattaching cloud storage to a cluster with a newer Avere OS version than used previously.
- 26341 Fixed an issue where a restart would occur if an object handle snapshot did not exist in the cloud core filer.
- 26366 Fixed a logic error that could cause a loop because a variable was defined in the wrong location.
- 26401 Removed string buffer bugs found under boundary conditions.
- 26469 Fixed an issue that could cause directory segments to become inaccessible after an upgrade, blocking accesses to multiple entries in the directory.

- 26556 Addressed an issue where cloud object segment sizes and number of segments were improperly set, causing a service restart.
- 26647 Disabled a bloom filter to fix a problem where a move from a source to a destination would cause a core loop if the destination already had the file.

Filesystem

- 23719 Added a feature to enable asynchronous writes and improve small file write performance.
- 23849 Fixed an issue that could cause continuous restarts in the migration code.
- 24307 Improved performance of rename actions by 30%.
- 24506 Resolved a problem that affected read-ahead caching on the Avere cluster. Linux-based NFS clients maintain their own client-side file caches, and this cache is populated with the client's own read-ahead. If the file being read is not in the Avere cache, and the NFS client read-ahead calls arrive out of order (that is, the read for the file's offset-zero segment comes in later than other reads), the Avere read-ahead calls for cache population of a "cold" file were not being sent.
- 24562 Added a feature to automatically identify and remove orphaned inodes.
- 24721 Added a custom setting to adjust the threshold that triggers local file replicas to be created.
- 24791 Added tracking for HA operations.
- 25003 Fixed a race condition in detecting the available space for internal metadata structures.
- 25084 Added code to fix a race condition in the token request system that could happen when failover caused changes in a mirror relationship.
- 25376 Fixed resource leaks in the cloud filesystem that caused occasional crashes due to memory shortages.
- 25439 Added more visibility into the status of the HA voter component.
- 25949 Improved token management when recycling files from the cache.
- 25963 Fixed an issue that could cause a service restart when adding or removing nodes or upgrading software.
- 25970 Corrected an incorrectly set flag that could cause long lookup times, timeouts, or system restarts.
- 25980 Fixed a filesystem service restart caused by performing read-ahead operations on a file that was concurrently being reclaimed from the cache due to space constraints.
- 26040 Changed code to correctly initialize the field `_ac1F1ags`.
- 26044 Fixed an error updating a task counter. This issue could cause a system restart when the task was cleaned up.
- 26140 Fixed an issue where temporary files prevented a migration job from completing.

- 26151 Corrected a situation where token management references were lost, causing system restarts.
- 26205 Updated code to also forward fsstat, pathconf, and fsinfo operations when always-forward is set.
- 26225 Changed code to prevent unnecessary token revokes because the token owner was incorrectly identified.
- 26255 Fixed a crash related to a race condition in the Avere Directory Name Cache component when recycling large numbers of files from the cache.
- 26270 Accelerated performance of the version database.
- 26296 Changed filesystem code to allow the Avere system to read a core filer directory even when some directory entries are inaccessible because of permissions. This change allows some core filer exports to be visible to the cluster and others to be ignored.
- 26307 Fixed an issue where retrying certain operations eight times caused a restart.
- 26309 Fixed code to prevent a filesystem restart caused by incorrectly cleaning up lock-related flags from failed operations.
- 26377 Removed a division-by-zero error bug that could be triggered by an inode file size of zero.
- 26414 Fixed a bug where an ENOENT error could cause the node to lose connection to a junction.
- 26452 Updated code for directory entry pruning to avoid problems caused by long queues.
- 26544 Fixed an issue that unnecessarily incremented an iterator on a data structure that was being erased.

General

- 21874 Changed log procedures to avoid repeated messages if HA is not enabled.
- 22141 Updated the help file for the XML-RPC method `alert.override()` to remove the incorrect implication that alert UUIDs are deprecated.
- 22259 Added time stamps and time ranges to traces and XML-RPC output.
- 22584 Fixed a filehandle leak in the Avere OS maintenance process.
- 22914 Added performance metrics to logs and XML-RPC output.
- 23762 Fixed a display issue in the Avere Control Panel related to directory service polling updates.
- 23792 Updated the API error message for proper formatting of private keys.
- 23896 Fixed an issue that caused the Avere OS management web service to restart.
- 24485 The Avere Control Panel now shows which cluster node is primary and which node the session is connected to. This information appears at the right side of the cluster name.

- 24856 Output about cluster upgrade status now displays plain text instead of HTML-encoded text.
- 24964 Updated the hidden alerts function to consistently apply parameter filters.
- 25037 Added a **Users** table to the dashboard that displays the IDs of recently active users and the number of operations they've performed. User activity can be displayed by node or cluster-wide (all nodes).
- 25144 Adjusted permissions for read-only XML-RPC API users to make sure that they can access non-modifying methods.
- 25406 Improved the layout and usability of the Alerts table in the Avere Control Panel dashboard.
- 25569 Fixed a bug that prevented support uploads from including custom settings.
- 25579 Fixed a bug in lock contention handling when reading the configuration database.
- 25583 Custom graphs in the **Analytics** tab now display core filer names in the legend instead of internal names.
- 25840 Fixed an issue caused by locks allocating more memory than needed.
- 25899 Fixed a bug that prevented uploading a new software image from a local machine.
- 25941 IP addresses on the IPMI settings page are now links that open a connection to that IPMI port.
- 25987 Fixed a software alert message to correctly identify software upgrade status.
- 26106 Upgraded to FreeBSD 11.2.
- 26166 Corrected an issue where a memory-related default was not set correctly on FXT 5850 hardware systems.
- 26229 Improved security by removing the ability to reuse cookies to execute XML-RPC commands.
- 26237 Fixed an issue where the SPS agent would stop and not restart.
- 26273 Added a file locking option in the cache policy management section of the Avere Control Panel. Use the **Advanced** settings section to select or clear the **Enable NLM caching** checkbox.
- 26408 Added a filtering option in the **Alerts** table of the Avere Control Panel Dashboard. You can choose to view only conditions, only alerts, or both conditions and alerts.

Migration

- 21725 Fixed an issue where a migration would not release an inode at file close unless it was a destroy operation.
- 23849 Fixed an issue that caused migration code to restart continuously.
- 25499 Changed code to properly repopulate a migrated subdirectory in the file system. This change prevents getting a stale junction after a junction refresh following a data migration.

- 25708 Fixed code to correctly handle cloud snapshot deletion when a snapshot has already been deleted from some nodes.
- 26140 Fixed a problem where some tasks were not marked as complete after a migration.

NFS

- 10363 Fixed several defects involving NFS export policy reference counts. These issues could result in the inability to delete an unused policy. Upon upgrade, this change corrects NFS export policy reference counts that were set incorrectly because of these defects.
- 14237 This ticket increases the number of supported NLM clients per vservers client-facing address from 4096 to 10000. This increased limit is consistent with the 10000 NFS client limit per node when only one client-facing address per node is allocated. This change only applies to new vservers created with a build containing this change. Downgrades to builds not containing this change are prohibited if the new feature has been used, because the new vservers will not process lock operations correctly after downgrade.
- 19883 Addressed NIST security vulnerability CVE-1999-1225.
- 24876 Improved LDAP/AD compatibility of the Avere Control Panel login system. Specifically:
- Removed the requirement to enable both LDAPS and LDAP with StartTLS on the LDAP server. Now, only LDAPS is required. (Enabling LDAP with StartTLS is optional.)
 - Allows custom LDAP and LDAPS ports to be configured. Use the XML-RPC method `dirServices.modify login` to set the variables `LDAPport` or `LDAPSport`.
 - Attribute queries can now be set to `ad` or `rfc2307` instead of the default, `auto`. The `auto` setting searches for both types of attributes, which is slower and can cause timeouts when used with large LDAP trees. Use the XML-RPC method `dirServices.modify login` to set the variable `loginQueryAttributes`.
- 25023 Addressed an internal configuration timing race that could cause the filesystem service to fail during startup after a software upgrade.
- 25491 Improved the NFS core filer RRDNS change procedure.
- 25546 Implemented NLM lock caching for NAS core filers. This feature is available only on new clusters created with a build containing this change. NLM lock caching for NAS core filers is enabled by default for system cache policies that do not allow directly connected clients. Custom cache policies can control the behavior using the XML-RPC methods `cachePolicy.create()` and `cachePolicy.modify()` by specifying the `enableNlmCaching` parameter.
- 25978 This change fixes a defect where the XML-RPC method `nfs.modify()` did not correctly handle changes to a subset of the available options. The workaround for this defect is to specify all of the options in each call to `nfs.modify()`.

- 26560 Added a buffer size check to prevent a possible internal string overflow.
- 26562 Fixed two defects found in an on-box debugging tool. These failures are not thought to have caused any problems at customer sites.

Security

- 21514 The procedure for storing cluster login information used with the Cluster Manager feature has changed. The Cluster Manager allows you to monitor multiple clusters from a single web page.
Before this change, the usernames and passwords used to access the remote clusters were stored in cleartext in an internal configuration file. The configuration file could be uploaded to the Avere support portal as part of troubleshooting, possibly exposing the passwords outside the cluster.
Now the login information is stored in a protected key distribution database that is stored on the cluster and never uploaded for support.
It is strongly recommended that all affected customers change the administrative passwords on clusters monitored in the Cluster Manager. Configure the Cluster Manager to use read-only accounts to connect to the remote clusters.
- 23579 Improved security by denying root access over SSH except from addresses in the local cluster.
- 26081 Updated PHP software to version 5.6.37.
- 26119 Fixed port vulnerabilities identified by the following NIST CVE numbers:

CVE-2017-15710	CVE-2018-1302	CVE-2018-7323	CVE-2018-7331
CVE-2017-15715	CVE-2018-1303	CVE-2018-7324	CVE-2018-7332
CVE-2017-3738	CVE-2018-1312	CVE-2018-7325	CVE-2018-7333
CVE-2018-0495	CVE-2018-6798	CVE-2018-7326	CVE-2018-7334
CVE-2018-0732	CVE-2018-6913	CVE-2018-7327	CVE-2018-7335
CVE-2018-0737	CVE-2018-7320	CVE-2018-7328	CVE-2018-7336
CVE-2018-0739	CVE-2018-7321	CVE-2018-7329	CVE-2018-7337
CVE-2018-1283	CVE-2018-7322	CVE-2018-7330	CVE-2018-7417
CVE-2018-1301			
- 26443 When an HTTP request is redirected to HTTPS, the system now returns the status code 308 (moved permanently) instead of 302 (moved temporarily).

SMB/CIFS

- 18095 Addressed a defect that could result in a CIFS service restart if the user token was unknown.
- 18917 Added interlocks that prevent advanced SMB/CIFS share options associated with POSIX mode bit junctions from being set on CIFS ACL junctions.
- 19493 Fixed an issue that prevented users with non-read access to a file or directory from accessing the file or directory at all.

- 23716 Addressed a problem that affected NetApp Clustered Data ONTAP core filers and prevented adding or modifying a subdirectory junction when the export policy prevented one of the components of the filepath from being created.
- 23843 Addressed a defect in Data Management out-of-sync processing for CIFS ACL migrations to cloud core filers. Out-of-sync processing of Native Identity File/Directory creations now fetches the source ACL from an internal cache rather than failing to fetch from the core filer.
- 23895 Resolved an issue where client calls to fetch SMB ACLs could cause recursive forwarding between FXT nodes due to an internal file-operation-routing loop.
- 25154 Addressed a defect in the receipt of Kerberized NFS client operations that could cause a filesystem service restart.
- 25576 Corrected an error that prevented the **Force directory mode** setting from displaying in the CIFS share details page on the Avere Control Panel.
- 25866 Fixed an issue where an SMB1 client could cause a CIFS service restart by performing repeated authentications.
- 25872 Updated code to prevent a filesystem service restart that could occur when invalid UTF-8 characters were present in NFS directory entry names. Problems were seen when these invalid directory entries were reported in the Data Management or Hot Files features.
- 26190 Fixed a problem that could cause a filesystem service restart when an SMB client lists large directories through a CIFS ACL junction.
- 26462 Addressed a defect that could occur when a NAS core filer returned an error in response to an SMB1/SMB2 FILE_INFO query operation. This defect caused failures with two different symptoms:
- A long running operation on an SMB/CIFS client, which could possibly result in an operation timeout.
 - An eventual restart of the CIFS ACL service because of a memory leak related to the original core filer error response.
- 26506 A problem was fixed that caused an NFS username lookup to sometimes return the UID of a different user that had the same primary GID. This could cause SMB clients to see incorrect file ownership or incorrect ACE identities in access control lists.
- 26522 For increased security, new vservers are now configured with outbound NTLMSSPN disabled. Contact Microsoft Customer Service and Support for more information about the implications of this change.

vFXT

- 24457 Updated vFXT deployment code to ensure that AWS tags are consistently applied.
- 24700 Removed external internet dependencies from vFXT systems running in AWS C2S.
- 26592 Fixed an issue that caused a service restart when attempting to look up network configuration in GCE projects with large numbers of virtual machines.

26595

Fixed an issue that could cause VM instance lookup to fail if the region information returned by the GCE API did not include zones for each region.

Contact Microsoft Customer Service and Support

Microsoft Customer Service and Support can be reached by website, phone, or email.

By web: Use the links under Support Information on <https://www.microsoft.com/avere/contact-us>

By phone: 1-888-88-AVERE, Option 2 (Toll-Free)
1-412-894-2570, Option 2

By email: averesupport@microsoft.com