



Learn about your AFX system

AFX

NetApp
February 11, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-afx/get-started/ontap-afx-storage.html> on February 11, 2026. Always check docs.netapp.com for the latest.

Table of Contents

| | |
|--|---|
| Learn about your AFX system | 1 |
| Learn about AFX storage systems | 1 |
| Typical application workloads | 1 |
| System design characteristics | 1 |
| Hardware infrastructure | 2 |
| Related information | 2 |
| Details of the AFX storage system architecture | 3 |
| Physical components | 3 |
| Logical components | 3 |
| AFX cluster deployment | 4 |
| Compare AFX storage system to AFF and FAS systems | 5 |
| Configuration differences | 5 |
| Unsupported or restricted Unified ONTAP capabilities | 6 |
| Changes to the command line interface | 6 |
| Related information | 7 |

Learn about your AFX system

Learn about AFX storage systems

The NetApp AFX storage system is based on a next-generation storage architecture that evolves the ONTAP storage model into a disaggregated high-performance NAS solution. AFX supports both file and object workloads with advanced technologies and processing techniques that provide extremely high performance.

Typical application workloads

The NetApp AFX storage system meets the unique demands of NAS and S3 object workloads that require high performance and independent scale. These applications benefit from an advanced design built on high concurrency and parallel I/O. AFX is ideal for organizations deploying and managing several different types of application workloads including:

- Training and iterative model refinement associated with deep learning where continuous high bandwidth and access to massive datasets is required.
- Processing diverse data types including text, images, and video.
- Real-time inference applications with low latency where strict response time windows are needed.
- Data science and machine learning pipelines that can benefit from self-service data management by data engineers and data scientists.

System design characteristics

The AFX system has several design characteristics that enable it to operate as a high-performance NAS platform.

Decouple storage and compute capabilities

Unlike other NetApp ONTAP storage systems, the compute and storage elements of an AFX cluster are decoupled and joined through a switched network. Disk ownership is no longer tied to specific nodes which provides several benefits. For example, the compute and storage components of an AFX cluster can be expanded independently.

Automated storage management

The physical aggregates are no longer available to the AFX storage administrator. Instead, AFX automatically manages the virtual capacity allocations for the nodes, as well as the RAID group configuration, when new storage shelves are added to the cluster. This design simplifies administration and provides an opportunity for nonspecialists to manage their data.

Single storage pool for the cluster

Because the storage nodes and shelves are decoupled with NetApp AFX, all storage capacity for the cluster is collected in a single pool known as a Storage Availability Zone (SAZ). The disks and shelves in a SAZ are available to all the storage nodes in an AFX cluster for read and write operations. In addition, all the cluster nodes can participate in disk rebuilds in the event of a failure. Refer to [FAQ for AFX storage systems](#) for more details.

High performance

NetApp AFX provides high and sustained bandwidth with ultra-low latency and so is designed for high performance NAS and object workloads. AFX uses the latest modern hardware as well as storage shelves capable of handling a high ratio of nodes to disks through its unique architecture. Scaling storage nodes beyond the typical 1:1 (node:shelf) ratio maximizes the possible performance profile of the disks to their edge limits. This design provides efficiency and storage density for your most critical applications.

Independent and massive scale

Based on the decoupled storage nodes and shelves, an AFX cluster can be independently and nondisruptively expanded based on your application needs. You can add storage nodes to get more CPU and throughput or add shelves to get more storage capacity and disk performance. The NetApp AFX architecture brings new possibilities for the maximum size of your cluster. For the latest limits for the AFX cluster based on your ONTAP release, refer to the [NetApp Hardware Universe](#).

Zero copy data mobility

NAS and object clients access volumes at the ONTAP cluster. You can relocate volumes across the nodes nondisruptively to achieve your capacity and performance balancing goals. With Unified ONTAP, a volume move is performed using SnapMirror technology which can take time and additional temporary capacity. But with AFX, a data copy operation is no longer needed within the shared Storage Availability Zone (SAZ). Instead, only the volume metadata is moved which dramatically improves performance. Refer to [FAQ for AFX storage systems](#) for more details.

Enhanced HA functionality

NetApp AFX offers a number of enhancements for high availability (HA) configuration and processing. AFX removes the need to directly connect HA partner nodes and instead allows HA pairs to communicate over the internal cluster network. This design gives administrators the option of deploying HA pairs in separate racks or rows in a datacenter for added fault tolerance. In addition, the AFX zero copy mobility extends to HA failover scenarios. When a node fails, its volumes will failover to the HA partner to commit any remaining writes to disk. Then ONTAP balances the volumes evenly across all surviving nodes in the cluster. This means you no longer need to consider storage failover performance in the initial design of your data placement.

Hardware infrastructure

The NetApp AFX storage system delivers a unified hardware and software solution that creates a simplified experience specific to the needs of high-performance NAS customers.



You should review the [FAQ for AFX storage systems](#) for more information about hardware interoperability and upgrade options.

The following hardware components are used with AFX clusters:

- AFX 1K controllers
- NX224 shelves
- Cisco Nexus 9332D-GX2B or Nexus 9364D-GX2A switches

Related information

- [NetApp Hardware Universe](#)
- [NetApp AFX](#)

Details of the AFX storage system architecture

The AFX architecture is composed of several hardware and software components. These system components are organized in different categories.

Physical components

When first getting started with AFX, it's helpful to begin with a high-level view of the physical components as they're installed in your data center.

Controller nodes

AFX controller nodes run a specialized personality of the ONTAP software designed to support the requirements of the AFX environment. Clients access the nodes through multiple protocols, including NFS, SMB, and S3. Each node has a complete view of the storage, which it can access based on the client requests. The nodes are stateful with non-volatile memory to persist critical state information and include additional enhancements specific to the target workloads.

Storage shelves and disks

AFX storage shelves use Non-volatile Memory Express over Fabrics (NVMe-oF) to connect high-density SSDs. The disks communicate over an ultra-low latency fabric using RDMA over Converged Ethernet (RoCE). The storage shelves, including the I/O modules, NICs, fans, and power supplies, are fully redundant with no single point of failure. Self-managed technology is used to administer and control all aspects of the RAID configuration and disk layout.

Cluster storage switch network

Redundant and high-performance switches connect the AFX controller nodes with the storage shelves. Advanced protocols are used to optimize performance. The design is based on VLAN tagging with multiple network paths, as well as tech-refresh configurations, to ensure continuous operation and ease of upgrade.

Client training environment

The client training environment is a lab environment with customer-provided hardware, such as GPU clusters and AI workstations. It's typically designed to support model training, inference, and other AI/ML related work. Clients access AFX using industry standard protocols such as NFS, SMB, and S3.

Client network

This internal network connects the client training environment to the AFX storage cluster. The network is provided and managed by the customer although NetApp expects to offer field recommendations for requirements and design.

Logical components

There are several logical components included with AFX. They are implemented in software along with the physical components of the cluster. The logical components enforce a structure that determines the use and configuration of the AFX systems.

Common storage pool

The Storage Availability Zone (SAZ) is a common pool of storage for the entire cluster. It's a collection of disks in the storage shelves that all the controller nodes have read and write access to. The SAZ offers a provisioning model with no fixed restrictions regarding which storage shelves can be used by the nodes; volume placement across the nodes is automatically handled by ONTAP. Customers can view free space and storage usage as properties of the entire AFX cluster.

FlexVolumes, FlexGroups, and buckets

FlexVolumes, FlexGroups, and S3 buckets are the *data containers* exposed to the AFX administrators based on the client access protocols. They operate identically to Unified ONTAP. These scalable containers are designed to abstract away many of the complex internal storage details, such as data placement and capacity balancing.

Data layout and access

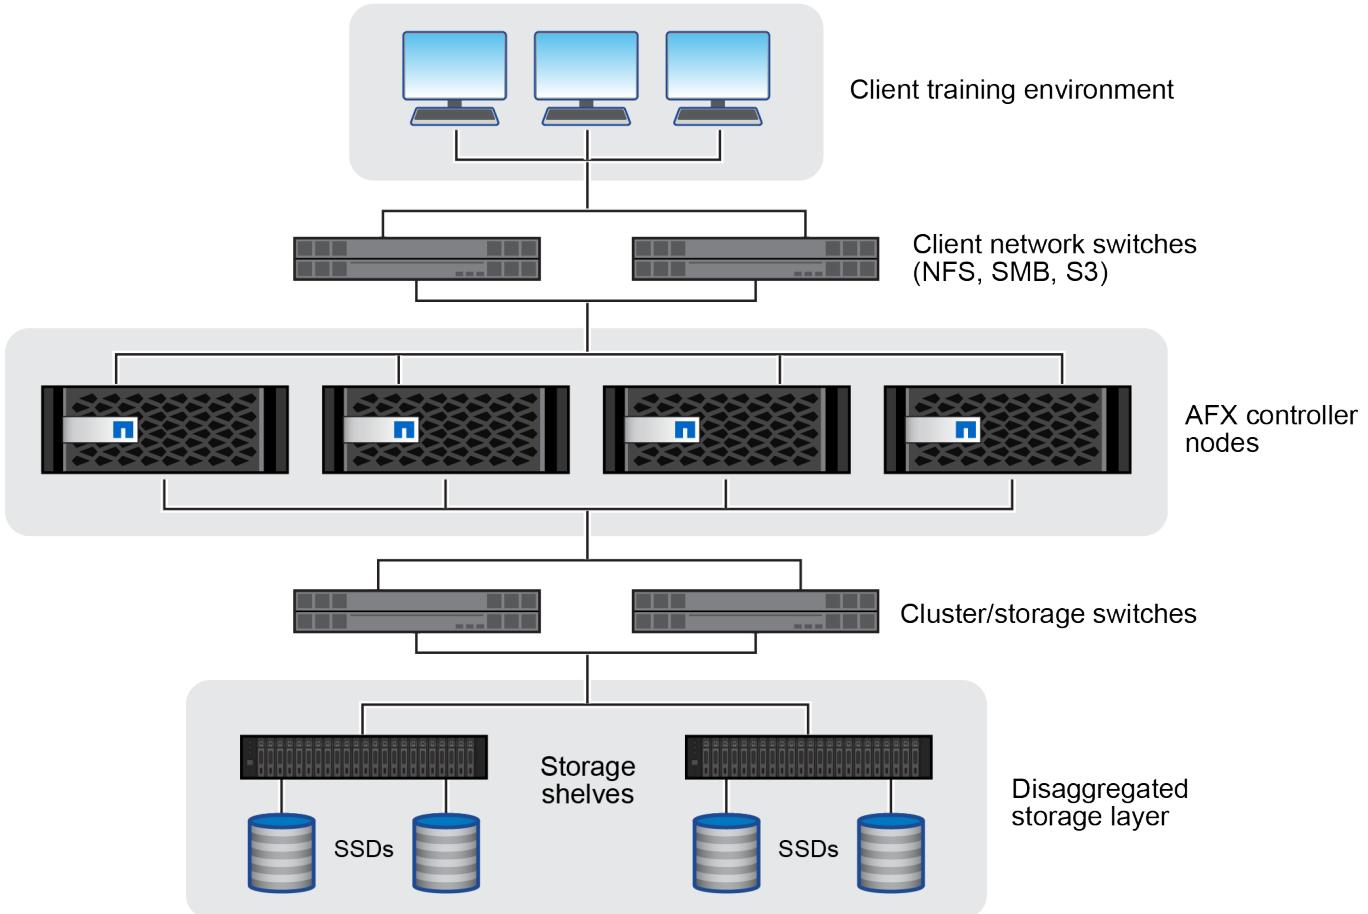
The data layout and access is tuned for seamless access and efficient utilization of the GPUs. This plays a critical role in eliminating bottlenecks and maintaining consistent performance.

SVMs and multi-tenancy

AFX provides a tenant model that builds on the SVM model available with AFF and FAS systems. The AFX tenant model is the same as Unified ONTAP but has been streamlined for simplified administration in a NAS and S3 object environment. For example, configuration options for SAN as well as aggregates and RAID groups have been removed.

AFX cluster deployment

The following figure illustrates a typical AFX cluster deployment. The AFX cluster includes controller nodes which are decoupled from the storage shelves and connected through a shared internal network. Outside the AFX cluster boundary, clients access the cluster through a separate client network.



Compare AFX storage system to AFF and FAS systems

NetApp AFX systems run a customized personality of ONTAP that differs from the ONTAP personality (referred to as Unified ONTAP) that runs on AFF and FAS storage. You should be aware of how AFX systems are similar to and different than FAS and AFF systems. This provides a valuable perspective and can be helpful when deploying AFX in your environment.

 The AFX documentation includes links to various topics at the Unified ONTAP doc site for details about features that behave in the same way regardless of the ONTAP personality. The additional content provides more depth that can be helpful as you administer your AFX storage system.

Configuration differences

There are a few areas where the AFX configuration differs from AFF and FAS systems.

Advanced Capacity Balancing

The advanced capacity balancing feature, controlled using the `-gdd` CLI parameter, is enabled by default for all FlexGroup volumes.

Unsupported or restricted Unified ONTAP capabilities

NetApp AFX is optimized for high-performance NAS and object workloads. Because of this, there are differences with AFF and FAS storage systems. The following features are not available with the NetApp AFX; the list is organized by major feature or functional area. You should also review the updates and changes for AFX in [What's new](#) based on your ONTAP release.

Block and SAN

- SAN administration and client access
- LUNs and NVMe namespaces
- Thick provisioning of volumes

Aggregates and physical storage

- MetroCluster
- Physical node-owned aggregates
- RAID management
- NetApp Aggregate Encryption (NAE)
- Aggregate-level deduplication
- SyncMirror (aggregate mirroring)
- FabricPool tiering
- Load-sharing mirrors

Data replication (SnapMirror)



All data replication is supported in both directions between Unified ONTAP and AFX with the same versioning restrictions described in [Compatible ONTAP versions for SnapMirror relationships](#) (with a few minor exceptions).

- No replication of a volume from an AFF or FAS system that contains a LUN or NVMe namespace
- FlexGroup volumes can only be replicated from AFX to Unified ONTAP version 9.16.1 or later (because of the need for Advanced Capacity Balancing)

Manageability

- ONTAPI API (ZAPI)
- REST APIs for unsupported features (such as MetroCluster)
- Some initial limitations on REST APIs for performance statistics
- AIQ Unified Manager support
- Grafana Harvest version 25.08.1 and later
- NetApp Trident version 25.10 and later

Changes to the command line interface

The ONTAP CLI available with AFX generally mirrors the CLI available with AFF and FAS systems. But there are several differences, including:

- New AFX commands related to:
 - Displaying the capacity of the storage availability zone
 - Boot media
- No SAN-related commands
- Aggregate management commands are no longer required
- Aggregate show now displays the entire Storage Availability Zone (SAZ)

Related information

- [AFX system characteristics](#)
- [Details of the AFX architecture](#)
- [FAQ for AFX storage systems](#)
- [Additional AFX cluster administration](#)
- [Additional AFX SVM administration](#)

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. 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—with prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

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

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.