



AFX documentation

AFX

NetApp
May 15, 2026

Table of Contents

- AFX documentation 1
- Release notes 2
 - What’s new in ONTAP 9.19.1 for AFX storage systems. 2
 - Performance. 2
 - REST API. 2
 - Security 3
 - Storage resource management 3
 - Upgrade 3
 - What’s new in ONTAP 9.18.1 for AFX storage systems. 3
 - Storage resource management enhancements 4
 - What’s new in ONTAP 9.17.1 for AFX storage systems. 4
 - Platforms 4
 - Changes to ONTAP limits and defaults affecting AFX systems 4
 - Changes to ONTAP limits 5
 - Related information 5
- Get started 6
 - Learn about your AFX system 6
 - Learn about AFX storage systems 6
 - Details of the AFX storage system architecture 8
 - Compare AFX storage system to AFF and FAS systems 10
 - Learn about the NetApp AI Data Engine 12
 - Quick start for setting up an AFX storage system 13
 - Install your AFX system 14
 - Installation and setup workflow for AFX 1K storage systems 14
 - Installation requirements for AFX 1K storage systems 15

AFX documentation

Release notes

What's new in ONTAP 9.19.1 for AFX storage systems

Learn about the new capabilities included with ONTAP 9.19.1 that are available with your AFX storage system.



Learn about the new and enhanced support for [AI Data Engine](#).

Performance

Update	Description
Cross File Sequential Read (CFSR)	Cross File Sequential Read can significantly improve read performance for certain media and entertainment workloads. The feature intelligently prefetches sequential video frames to dramatically improve throughput and reduce latency for video editing applications. Review Additional administration for an AFX storage system SVM for more information.
Global deduplication	Global deduplication extends ONTAP deduplication space savings to the entire Storage Availability Zone (SAZ). By eliminating duplicate data blocks across all the volumes in the SAZ, it can significantly improve storage efficiency and reduce total cost of ownership. The feature is enabled by default on AFX systems and operates transparently without the need for administrative configuration. Refer to Additional administration for an AFX storage system SVM for more information.
Dynamic Storage Efficiency	Dynamic Storage Efficiency intelligently optimizes storage efficiency operations (compression, deduplication, and compaction) for large workloads that require high performance and maximum throughput. It automatically detects high-volume write conditions and temporarily suspends inline and background storage efficiency processes to prevent performance degradation during peak loads. When throughput drops below a defined threshold, the storage efficiency operations are resumed and the savings are recovered without manual intervention. The feature is available on AFX storage systems and enabled by default with no configuration required.

REST API

Update	Description
Balanced placement enhancement for Kubernetes storage orchestration	ONTAP 9.19.1 adds six parameters currently available in the <code>/storage/volumes</code> REST API endpoint to the <code>/applications/containers</code> endpoint. This enhancement enables customers using NetApp Trident for Kubernetes storage orchestration to create volumes with ONTAP's native balanced placement capability.

Security

Update	Description
ARP/AI support for eligible SnapMirror primary volumes	ARP/AI supports eligible primary volumes in SnapMirror synchronous relationships with AFX systems. ARP/AI runs on the active read-write primary volume, creates primary-only ARP snapshots, and provides machine-learning-based ransomware detection.
Support for TLS hardware offload	You can offload Transport Layer Security (TLS) encryption and decryption for your entire cluster to supported Ethernet cards. This reduces CPU overhead and improves performance for systems using TLS.

Storage resource management

Update	Description
System Manager support for storage availability zone capacity details.	You can view your storage availability zone capacity details including space thresholds, space consumed by snapshots, and data reduction in System Manager. You can also set space usage thresholds to the percentage at which you want ONTAP to generate warning and critical alerts.
Support for modifying default settings that control automatic QoS throughput ceiling increases	You can modify the default settings that control how QoS policy groups automatically increase the QoS throughput ceiling in response to short-term needs for higher performance. These settings cannot be modified in ONTAP 9.18.1 and earlier.

Upgrade

Update	Description
Enhanced reporting for automated pre-upgrade checks	When executing automated pre-upgrade checks, errors and warning are displayed separately to help you easily identify issues that must be resolved before you upgrade.

Related information

- [ONTAP 9 release highlights](#)

What's new in ONTAP 9.18.1 for AFX storage systems

Learn about the new capabilities included with ONTAP 9.18.1 that are available with your AFX storage system.



Learn about the new and enhanced support for [AI Data Engine](#).

Storage resource management enhancements

Update	Description
Enhanced volume placement	NetApp AFX automatically balances the placement of volumes across all the nodes in a cluster. With previous ONTAP releases, the placement algorithm was based on the number of volumes in the cluster. Every node is assigned the same number of volumes regardless of activity. Beginning with ONTAP 9.18.1, the algorithm has been enhanced to consider the performance of the nodes when placing or moving volumes. This results in improved performance balancing across the nodes in the AFX cluster and makes it much less likely that any single node becomes overloaded.

Related information

- [ONTAP 9 release highlights](#)

What's new in ONTAP 9.17.1 for AFX storage systems

Learn about the new capabilities included with ONTAP 9.17.1 that are available with your AFX storage system.

Platforms

Update	Description
Platforms	<p>The following NetApp AFX storage system components are available along with the related supporting technology. Together this platform delivers a unified hardware and software solution that creates a simplified experience specific to the needs of high-performance NAS and S3 customers.</p> <ul style="list-style-type: none">• AFX 1K controllers• NX224 shelves• Cisco Nexus 9332D-GX2B and Nexus 9364D-GX2A switches

Related information

- [ONTAP 9 release highlights](#)

Changes to ONTAP limits and defaults affecting AFX systems

Learn about the changes to limits and defaults affecting AFX systems. NetApp strives to help its customers understand the most important default and limit changes in each ONTAP release.



You should review the details available at the NetApp Hardware Universe and Interoperability Matrix Tool to determine the capabilities of your AFX cluster.

Changes to ONTAP limits

Feature	Limit change	Changed in release
Nodes per cluster	You can expand your AFX cluster to a maximum of 32 nodes. This allows you to scale your AFX storage system to meet the growing demands of your workloads and applications.	ONTAP 9.19.1 RC

Related information

- [NetApp Hardware Universe](#)
- [NetApp Interoperability Matrix Tool](#)
- [Interoperability Matrix Tool overview](#)

Get started

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

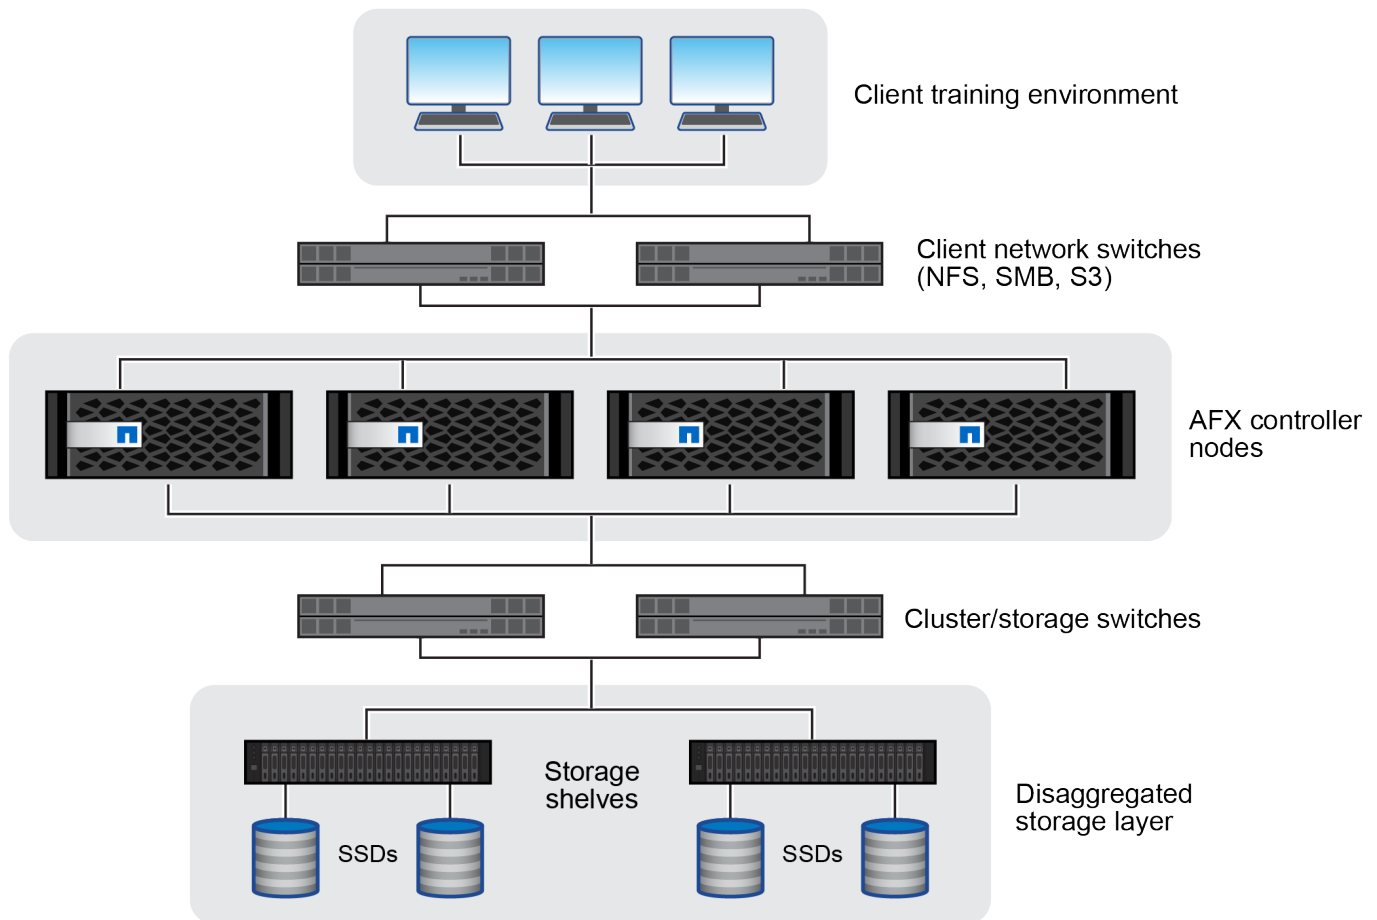
The data layout and access is tuned for seamless access and efficient utilization of the GPUs. This plays a critical role in eliminating bottlenecks and maintaining consistent performance.

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 versions prior to 25.08.1

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](#)

Learn about the NetApp AI Data Engine

NetApp AI Data Engine (AIDE) is a storage-integrated AI data service that runs on compute nodes attached to an AFX cluster. You can optionally install and integrate AIDE with your AFX cluster using one of two deployment models: NetApp-provided data compute nodes (DCNs) or AIDE software on third-party servers. NetApp DCN deployments provide full AIDE capabilities, while third-party server deployments are metadata-focused.

AIDE overview

AIDE is a system designed to accelerate and simplify the preparation and management of your data for AI and ML workloads. It is installed separately and integrates closely with an AFX storage system cluster. AIDE leverages the high-performance and scalable AFX storage foundation while adding specialized services for AI data processing, governance, and workflow automation.

A core strength of AIDE is its ability to centralize and automate the management of metadata across large, distributed datasets. It can be configured to continuously catalog and update the metadata store based on the data in your ONTAP volumes, enabling users to search, classify, and apply governance policies efficiently. The internal processing pipeline generates AI-ready datasets, including the requisite creation of vector embeddings for semantic search and retrieval tasks.

Security and governance are also central to its design. AIDE provides authentication, access control with new RBAC roles, data classification, and policy-driven enforcement for the protection of sensitive information. Features such as automated redaction, masking, and audit logging help you meet applicable regulatory requirements and protect critical data throughout the entire AI lifecycle.

AIDE empowers data engineers and data scientists by providing tools for rapid data discovery, curation, and workflow management. With the dedicated AIDE console and enhanced AFX REST API, users can create curated data collections, perform vector searches, and integrate with external AI and GenAI applications. AIDE reduces the required manual effort and enables teams to focus on extracting insights from their data.

Installation options

There are two options for installing and integrating AIDE with your AFX storage cluster.

Set up AFX and AIDE together

Deploying the AFX storage system and AIDE simultaneously provides a unified and streamlined installation process, ensuring immediate integration and optimal configuration. This turnkey approach simplifies onboarding, guarantees compatibility, and enables you to quickly leverage a fully operational AI data platform which is ideal for new deployments seeking rapid time-to-value.

Integrate AIDE with an existing AFX cluster

If you already have an AFX storage cluster installed and operational, you can add AIDE at any time to enhance the access and value of your data. AIDE is deployed as an add-on to the existing AFX storage environment, connecting to the established cluster storage network without disrupting operations. The integration process automatically discovers available ONTAP volumes and synchronizes metadata. Your data processing workflows are enhanced through AI automation, governance, and curation tools while preserving their existing storage foundation.

Related information

- [Learn about AI Data Engine](#)
- [FAQ for NetApp AI Data Engine](#)
- [Set up your ONTAP AIDE cluster](#)
- [NetApp AI Data Engine](#)

Quick start for setting up an AFX storage system

To initially get up and running with your AFX system, you need to install the hardware components, set up your cluster, and prepare to administer the cluster and SVMs.

1

Will AI Data Engine (AIDE) also be installed with AFX?

You need to decide whether to install AFX only or AFX along with AI Data Engine. Review the [AIDE installation options](#) to help determine the best fit for your needs.

2

Install and set up your hardware

[Install your AFX storage system](#) and optionally [install AIDE third-party servers](#) or [install AIDE NetApp data compute nodes](#) if you're including AIDE as part of the initial configuration.

3

Set up your cluster

Follow the quick and easy process to [set up](#) your ONTAP cluster using System Manager. If you're also installing AIDE with the data compute nodes along with AFX, review [Set up AIDE with a new AFX cluster](#) for additional required configuration.

4

Prepare to administer your cluster

Before deploying AFX in a production environment, it's essential to [prepare](#) by understanding the administrative structure—including storage virtual machines (SVMs), users, roles, and management interfaces—to ensure secure, efficient, and effective cluster management.

Install your AFX system

Installation and setup workflow for AFX 1K storage systems

To install and configure your AFX 1K storage system, you review the hardware requirements, prepare your site, install the switches, install and cable the hardware components, power on the system, and set up your ONTAP AFX cluster.

1

Review the hardware installation requirements

Review the hardware requirements to install your AFX 1K storage system.

2

Prepare to install your AFX 1K storage system

Prepare to install your AFX 1K storage system by preparing the site, checking environmental and electrical requirements, ensuring sufficient rack space, unpacking the equipment, verifying contents to the packing slip, and registering the hardware for support.

3

Install the switches for your AFX 1K storage system

Install Cisco Nexus 9332D-GX2B or 9364D-GX2A switches in the cabinet or rack. Install a pass-through panel kit if using the Cisco Nexus 9364D-GX2A switch.

4

Install the hardware for your AFX 1K storage system

Install the rail kits for your storage system and shelves. Secure your storage system in the cabinet or telecommunications rack. Next, slide the shelves onto the installed rails. Finally, attach cable management devices to the rear of the storage system for organized cable routing.

5

Cable the controllers and shelves for your AFX 1K storage system

To cable the hardware, first connect the storage controller nodes to your network, then connect the controller nodes and storage shelves to the cluster switches.

6

Power on and configure the switches for your AFX 1K storage system

Cable the hardware, then power on and configure the switches for your AFX 1K storage system. Check the configuration instructions for the Cisco Nexus 9332D-GX2B and 9364D-GX2A switches.

7

Power on your AFX 1K storage system

Power on each storage shelf and assign a unique shelf ID before powering on the controller nodes to clearly identify each shelf in the setup.

If you are deploying AFX 1K with AI Data Engine (AIDE), you can install AIDE data compute nodes after you finish installing the AFX system.



- For third-party server deployments, see [requirements for AI Data Engine with AIDE software on third-party servers](#) and [install AIDE software on third-party servers](#).
- For NetApp DCN deployments, see [requirements for AI Data Engine on NetApp-provided data compute nodes](#).

Installation requirements for AFX 1K storage systems

Review the equipment needed and the lifting precautions for your AFX 1K storage controller and storage shelves.

Equipment needed for install

To install your AFX 1K storage system, you need the following equipment and tools.

- Access to a Web browser to configure your storage system
- Electrostatic discharge (ESD) strap
- Flashlight
- Laptop or console with a USB/serial connection
- Paperclip or narrow-tipped ballpoint pen for setting storage shelf IDs
- Phillips #2 screwdriver


Lifting precautions

AFX storage controller and storage shelves are heavy. Exercise caution when lifting and moving these items.

Storage controller weights

Take the necessary precautions when moving or lifting your AFX 1K storage controller.

An AFX 1K storage controller can weigh up to 62.83 lbs (28.5 kg). To lift the storage controller, use two people or a hydraulic lift.




Storage shelf weights

Take the necessary precautions when moving or lifting your shelf.

NX224 shelf

An NX224 shelf can weigh up to 60.1 lbs (27.3 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.



Related information

- [Safety information and regulatory notices](#)

What's next?

After you've reviewed the hardware requirements, you [prepare to install your AFX 1K storage system](#).

= Prepare to install your AFX 1K storage system

:icons: font

:relative_path: ./install-setup/

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

Prepare to install your AFX 1K storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the system to access support benefits.

== Step 1: Prepare the site

To install your AFX 1K storage system, ensure that the site and the cabinet or rack that you plan to use meet specifications for your configuration.

Steps

1. Use [NetApp Hardware Universe](#) to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and switches:
 - 2U for each AFX controller node and NX224 shelf
 - 1U or 2U per switch, depending on switch model.

== Step 2: Unpack the boxes

After ensuring the site and cabinet meet specifications, unpack the boxes and compare the contents to the packing slip.

Steps

1. Carefully open all the boxes and lay out the contents in an organized manner.
2. Compare the contents you've unpacked with the list on the packing slip. If you find any discrepancies, record them for further action.

You can get your packing list by scanning the QR code on the side of the shipping carton.

The following items are some of the contents you might see in the boxes.

Hardware

Cables

- Bezel
- Storage system
- Rail kits with instructions
- Storage shelf
- Cisco Nexus 9332D-GX2B or 9364D-GX2A switch
- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables
- USB-C serial port cable

== Step 3: Register your storage system

After you ensure that your site meets the requirements for your AFX 1K storage system specifications, and you verify that you have all the parts you ordered, register your storage system.

Steps

1. Locate the serial numbers for your storage system.

You can find the serial numbers in the following locations:

- On the packing slip
- In your confirmation email
- On each controller or for some systems, on the system management module of each controller



2. Go to the [NetApp Support Site](#).
3. Decide whether you need to register your storage system:

If you are a...	Follow these steps...
Existing NetApp customer	<ol style="list-style-type: none"> a. Sign in with your username and password. b. Select Systems > My Systems. c. Confirm that the new serial number is listed. d. If the serial number is not listed, follow the instructions for new NetApp customers.
New NetApp customer	<ol style="list-style-type: none"> a. Click Register Now, and create an account. b. Select Systems > Register Systems. c. Enter the storage system's serial number and requested details. <p>Once NetApp approves your registration, you can download the required software. Approval takes up to 24 hours.</p>

What's next?

After you've prepared to install your AFX 1K hardware, you [install the switches for your AFX 1K storage system](#).

= Install hardware

= Install the switches for your AFX 1K storage system

:icons: font

:relative_path: ./install-setup/

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

After you complete your preparation for the AFX 1K storage system installation, you should install the switches in the cabinet or telco rack.

Install Cisco Nexus 9332D-GX2B or 9364D-GX2A switches in the cabinet or rack. Install a pass-through panel kit if using the Cisco Nexus 9364D-GX2A switch.

Before you begin

Make sure you have the following components available:

- The pass-through panel kit, which is available from NetApp (part number X8784-R6).

The NetApp pass-through panel kit contains the following hardware:

- One pass-through blanking panel
- Four 10-32 x .75 screws
- Four 10-32 clip nuts
- For each switch, eight 10-32 or 12-24 screws and clip nuts to mount the brackets and slider rails to the front and rear cabinet posts.
- The Cisco standard rail kit to install the switch in a NetApp cabinet.



Jumper cords are not included with the pass-through kit. Contact NetApp to order the right jumper cables if they are not shipped with your switches.



If the airflow for your switches is configured for port-side intake (burgundy colored fans and PSUs), the network ports for the switches must be installed facing the front of the cabinet, and the exhaust fans must face the rear of the cabinet. With this configuration, you must ensure that you use cables long enough to run from the network ports in the front of the cabinet to the storage ports in the rear of the cabinet.

For more detailed information about these switches, please visit the Cisco website: [Cisco Nexus 9332D-GX2B NX-OS Mode Switch Hardware Installation Guide](#) and [Cisco Nexus 9364D-GX2A NX-OS Mode Switch Hardware Installation Guide](#).




Always wear a grounded wrist strap connected to a verified ground point during installation and maintenance procedures. Failure to follow proper ESD precautions can cause permanent damage to controller nodes, storage shelves, and network switches.

Steps


1. Install the pass-through blanking panel.
 - a. Determine the vertical location of the switches and blanking panel in the cabinet or rack.
 - b. Install two clip nuts on each side in the appropriate square holes for front cabinet rails.
 - c. Center the panel vertically to prevent intrusion into adjacent rack space, and then tighten the screws.

- d. Insert the female connectors of both jumper cords from the rear of the panel and through the brush assembly.




1	Female connector of the jumper cord.
----------	--------------------------------------

2. Install the rack-mount brackets on the switch chassis.
 - a. Position a front rack-mount bracket on one side of the switch chassis so that the mounting ear is aligned with the chassis faceplate (on the PSU or fan side), and then use four M4 screws to attach the bracket to the chassis.




- b. Repeat step 2a with the other front rack-mount bracket on the other side of the switch.
 - c. Install the rear rack-mount bracket on the switch chassis.
 - d. Repeat step 2c with the other rear rack-mount bracket on the other side of the switch.
3. Install the clip nuts in the square hole locations for all four IEA posts.



Mount the two 9332D-GX2B switches in cabinet locations that provide efficient access to controllers and shelves, such as the middle rows.

4. Install the slider rails in the cabinet or rack.

- a. Position the first slider rail at the desired location on the back side of the rear left post, insert screws with the matching thread type, and then tighten the screws with your fingers.



1	As you gently slide the slider rail, align it to the screw holes in the rack.
----------	-------------------------------------------------------------------------------

2	Tighten the screws of the slider rails to the cabinet posts.
----------	--------------------------------------------------------------


- b. Repeat step 4a for the right-side rear post.
- c. Repeat steps 4a and 4b at the desired locations on the cabinet.

5. Install the switch in the cabinet or rack.




This step requires two people: one person to support the switch from the front and another to guide the switch into the rear slider rails.

- a. Position the back of the switch at the desired location on the cabinet.



1	As the chassis is pushed toward the rear posts, align the two rear rack-mount guides with the slider rails
2	Gently slide the switch until the front rack-mount brackets are flush with the front posts.

- b. Attach the switch to the cabinet or rack.



1	With one person holding the front of the chassis level, the other person should fully tighten the four rear screws to the cabinet posts.
----------	------------------------------------------------------------------------------------------------------------------------------------------

- c. With the chassis now supported without assistance, fully tighten the front screws to the posts.
- d. Repeat steps 5a through 5c for the second switch at the desired location on the cabinet.



By using the fully installed switch as a support, it is not necessary to hold the front of the second switch during the installation process.

- 6. When the switches are installed, connect the jumper cords to the switch power inlets.
- 7. Connect the male plugs of both jumper cords to the closest available PDU outlets.



To maintain redundancy, the two cords must be connected to different PDUs.

- 8. Connect the management port on each switch to either of the management switches (if ordered) or connect them directly to your management network.

The management network port is the lower RJ-45 port near the right PSU. Route the CAT6 cable for each switch through the pass-through panel after installing the switches to connect to the management switches or network.

What's next?

After you install the switches in the cabinet or rack, you [install the AFX 1K storage system and shelves in the cabinet or rack](#).

= Install your AFX 1K storage system

:icons: font

:relative_path: ../install-setup/

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

After you install the switches, you should install the hardware for your AFX 1K storage system. First, install the rail kits. Then install and secure your storage system in a cabinet

or telco rack.

Before you begin


- Make sure you have the instructions packaged with the rail kit.
- Understand the safety concerns related to the weight of the storage system and storage shelf.
- Understand that the airflow through the storage system enters from the front where the bezel or end caps are installed and exhausts out the rear where the ports are located.



Always wear a grounded wrist strap connected to a verified ground point during installation and maintenance procedures. Failure to follow proper ESD precautions can cause permanent damage to controller nodes, storage shelves, and network switches.

Steps

1. Install the rail kits for your storage system and storage shelves, as needed, using the instructions included with the kits.
2. Install and secure your controller in the cabinet or telco rack:
 - a. Position the storage system onto the rails in the middle of the cabinet or telco rack, and then support the storage system from the bottom and slide it into place.
 - b. Secure the storage system to the cabinet or telco rack using the included mounting screws.
3. Attach the bezel to the front of the controller.
4. If your AFX 1K storage system came with a cable management device, attach it to the rear of the storage system.



5. Install and secure the storage shelf:
 - a. Position the back of the storage shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

In general, storage shelves and controllers should be installed in close proximity to the switches. If you are installing multiple storage shelves, place the first storage shelf directly above the controllers. Place the second storage shelf directly under the controllers. Repeat this pattern for any additional storage shelves.

- b. Secure the storage shelf to the cabinet or telco rack using the included mounting screws.

What's next?

After you've installed the hardware for your AFX system, review the [supported cabling configurations for your](#)

AFX 1K storage system.

= Cabling

= Supported configurations for your AFX 1K storage system

:icons: font

:relative_path: ./install-setup/

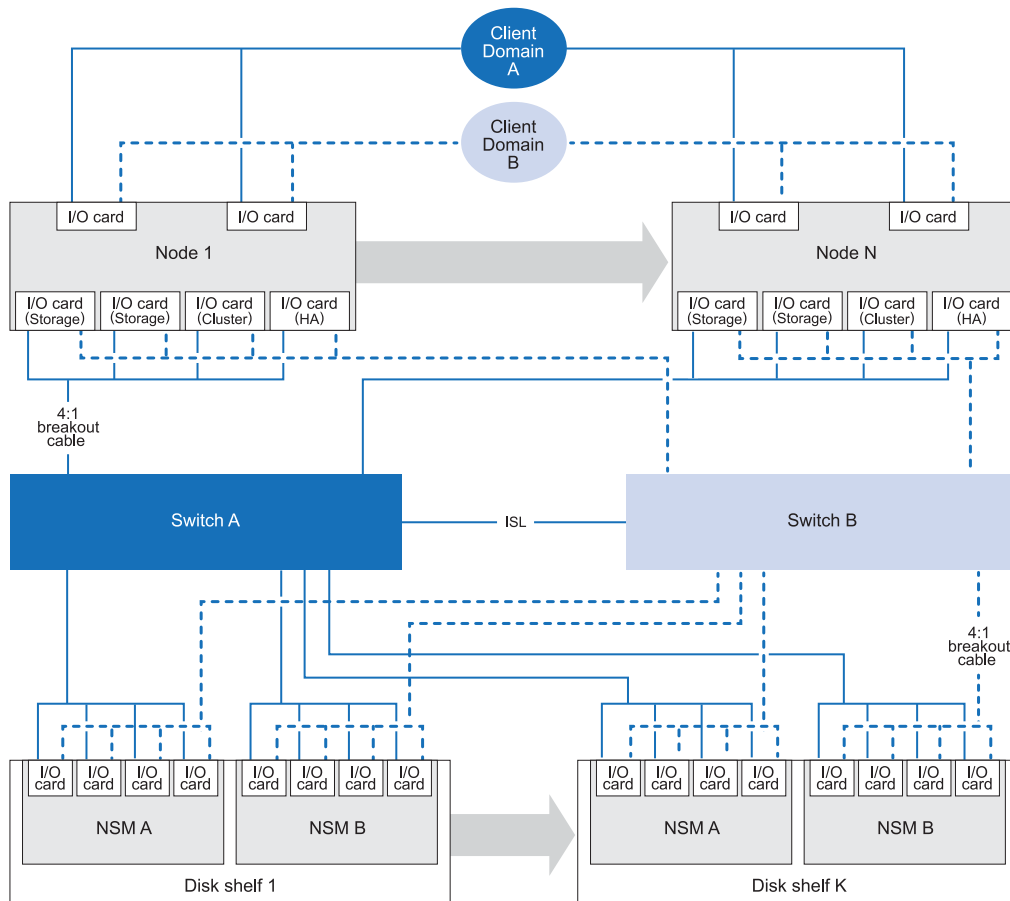
:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

Learn about the supported hardware components and cabling options for the AFX 1K storage system, including compatible storage disk shelves, switches, and cable types required for proper system setup.

== Supported AFX 1K cabling configuration


The initial configuration of the AFX 1K storage system supports a minimum of four controller nodes connected through dual switches to the storage disk shelves.

Additional controller nodes and disk shelves expand the initial AFX 1K storage system configuration. Expanded AFX 1K configurations follow the same switch-based cabling methodology as the schema depicted below.



== Supported hardware components

Review the compatible storage disk shelves, switches, and cable types for the AFX 1K storage system.

Controller Shelf	Disk Shelf	Supported Switches	Supported Cables
AFX 1K	NX224	<ul style="list-style-type: none"> • Cisco Nexus 9332D-GX2B (400GbE) • Cisco Nexus 9364D-GX2A (400GbE) 	<ul style="list-style-type: none"> • 400GbE QSFP-DD breakout to 4x100GbE QSFP breakout cable cables <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  Breakout cables are used for 100GbE connections between the switches, controllers, and disk shelves. </div> <ul style="list-style-type: none"> ◦ 100GbE cables to controller cluster and HA ports ◦ 100GbE cables to disk shelves • 2 x 400GbE cables for ISL connections between switch A and switch B • RJ-45 cables for management connections

What's next?

After reviewing the supported system configuration and hardware components, [review the network requirements for your AFX 1K storage system](#).

= Network requirements for your AFX 1K storage system

:hardbreaks:

:icons: font

:linkattrs:

:relative_path: ./install-setup/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

Record the required information for each network you connect to your AFX 1K storage system.

== Gather network information

Before you begin the installation of your AFX 1K storage system, gather the required network information

- Host names and IP addresses for each of the storage system controllers and all applicable switches.

Most storage system controllers are managed through the e0M interface by connecting to the Ethernet service port (wrench icon).

Refer to the [Hardware Universe](#) for the latest information.

- Cluster management IP address

The cluster management IP address is a unique IP address for the cluster management interface used by the cluster administrator to access the admin storage VM and manage the cluster. You can obtain this IP address from the administrator responsible for assigning IP addresses in your organization.

- Network subnet mask

During cluster setup, ONTAP recommends a set of network interfaces appropriate for your configuration.

You can adjust the recommendation if necessary.

- Network gateway IP address
- Node management IP addresses (one per node)
- DNS domain names
- DNS name server IP addresses
- NTP server IP addresses
- Data subnet mask
- IP subnet for management network traffic.

== Network requirements for Cisco switches

For Cisco Nexus 9332D-GX2B and 9364D-GX2A switch installation and maintenance, be sure to review cabling and network requirements.

=== Network requirements

You need the following network information for all switch configurations.

- IP subnet for management network traffic
- Host names and IP addresses for each of the storage system controllers and all applicable switches
- Refer to the [Hardware Universe](#) for the latest information.

=== Cabling requirements

- You have the appropriate number and type of cables and cable connectors for your switches. See the [Hardware Universe](#).
- Depending on the type of switch you are initially configuring, you need to connect to the switch console port with the included console cable.

What's next?

After reviewing the network requirements, you [cable the controllers and storage shelves for your AFX 1K storage system](#).

= Cable the hardware for your AFX 1K storage system

:icons: font

:relative_path: ./install-setup/

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

After you install the rack hardware for your AFX 1K storage system, install the network cables for the controllers, and connect the cables between the controllers and storage shelves.

Before you begin


Contact your network administrator for information about connecting the storage system to your network switches.

About this task

- These procedures show common configurations. The specific cabling depends on the components ordered for your storage system. For comprehensive configuration details and slot priorities, see [NetApp Hardware](#)


Universe.

- The I/O slots on an AFX controller are numbered 1 through 11.



- The cabling graphics show arrow icons indicating the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it over and try again.



The connector components are delicate and care should be taken when clicking into place.

- When cabling to an optical fiber connection, insert the optical transceiver into the controller port before cabling to the switch port.
- The AFX 1K storage system utilizes 4x100GbE breakout cables on the cluster and storage network. The 400GbE connections are made to the switch ports, and the 100GbE connections are made to the controller and drive shelf ports. Storage and HA/Cluster connections can be made to any non-ISL port on the switch.

For a given 4x100GbE breakout cable connection to the specific switch port, you connect all four ports from a given controller to the switch over this single breakout cable.

- 1 x HA port (slot 1)
- 1 x cluster port (slot 7)
- 2 x storage ports (slots 10, 11)

All "a" ports connect to switch A, and all "b" ports connect to switch B.



Cisco Nexus 9332D-GX2B and 9364D-GX2A switch configurations to the AFX 1K storage system require 4x100GbE breakout cable connections.


== Step 1: Connect the controllers to the management network

Connect the management (wrench) port on each controller to either of the management switches or connect them directly to your management network.

Use the 1000BASE-T RJ-45 cables to connect the management (wrench) ports on each controller to the management network switches.



1000BASE-T RJ-45 cables



Do not plug in the power cords yet.

1. Connect to host network.

== Step 2: Connect the controllers to the host network
Connect the Ethernet module ports to your host network.

This procedure may differ depending on your I/O module configuration. The following are some typical host network cabling examples. See [NetApp Hardware Universe](#) for your specific system configuration.

Steps

1. Connect the following ports to your Ethernet data network switch A.


- Controller (Example)

- e2a
- e3a

100GbE cables



To host network switch A




2. Connect the following ports to your Ethernet data network switch B.

- Controller (Example)

- e2b
- e3b

100GbE cables





== Step 3: Cable the cluster, HA, and storage connections to Switch A

Use the 4x100GbE breakout cable to connect the cluster, HA, and storage ports on the controllers to network switch A. Cluster and HA interconnect cable to connect ports e1a and e7a to switch A. The e1a ports are used for the HA connections, and the e7a ports are used for the cluster connections. For the storage connections, connect ports e10a and e11a to switch A.

Ensure you have the correct cables and connectors for your switches. See [Hardware Universe](#) for more information.

Steps

1. Connect the following controller ports to any non-ISL port on the network switch A.

- Controller
 - e1a (HA)
 - e7a (Cluster)
 - e10a (Storage)
 - e11a (Storage)

100GbE cables



== Step 4: Cable the cluster, HA, and storage connections to Switch B

Use the 4x100GbE breakout cable to connect the cluster, HA, and storage ports on the controllers to network switch B. Cluster and HA interconnect cable to connect ports e1b and e7b to switch B. The e1b ports are used for the HA connections, and the e7b ports are used for the cluster connections. For the storage connections, connect ports e10b and e11b to switch B.

Ensure you have the correct cables and connectors for your switches. See [Hardware Universe](#) for more information.

1. Connect the following controller ports to any non-ISL port on the network switch B.

- Controller
 - e1b (HA)
 - e7b (Cluster)
 - e10b (Storage)
 - e11b (Storage)

100GbE cables




== Step 5: Cable the shelf-to-switch connections
Connect the NX224 storage shelves to the switches.

For the maximum number of shelves supported for your storage system and for all of your cabling options, see [NetApp Hardware Universe](#).

1. Connect the following shelf ports to any non-ISL port on switch A and switch B for module A.

- Module A to switch A connections
 - e1a
 - e2a
 - e3a
 - e4a
- Module A to switch B connections
 - e1b
 - e2b
 - e3b
 - e4b

100GbE cables




2. Connect the following shelf ports to any non-ISL port on switch A and switch B for module B.

- Module B to switch A connections
 - e1a
 - e2a
 - e3a
 - e4a
- Module B to switch B connections
 - e1b
 - e2b
 - e3b

- e4b

100GbE cables



What's next?

After cabling the hardware, [power on and configure the switches](#).

= Power on and configure the switches for your AFX 1K storage system

:icons: font

:relative_path: ./install-setup/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

After you cable your AFX 1K storage system, you need to power on and configure the Cisco Nexus 9332D-GX2B or 9364D-GX2A switches.

Steps

1. Plug the power cords for the switches into the power sources.
2. Connect the ISL cables between the two switches.
 - For Cisco Nexus 9332D-GX2B switches, use ports 31/32 for the ISL connections. See the [Cisco Nexus 9332D-GX2B NX-OS Mode Switch Hardware Installation Guide](#) for more information.
 - For Cisco Nexus 9364D-GX2A switches, use ports 63/64 for the ISL connections. See the [Cisco Nexus 9364D-GX2A NX-OS Mode Switch Hardware Installation Guide](#) for more information.
3. Power on each switch.
4. Configure the switches to support the AFX 1K storage system.
 - For Cisco Nexus 9332D-GX2B switches, see the cluster and storage switches documentation [Configure Cisco Nexus 9332D-GX2B switch](#).
 - For Cisco Nexus 9364D-GX2A switches, see the cluster and storage switches documentation [Configure Cisco Nexus 9364D-GX2A switch](#).

What's next?

After configuring the switches for your AFX 1K storage system, [power on the AFX 1K storage system](#).

= Power on your AFX 1K storage system

:icons: font

:relative_path: ./install-setup/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

After you install the rack hardware for your AFX 1K storage system and install the cables for the controller nodes and storage shelves, you should power on your storage shelves and controller nodes.

== Step 1: Power on the shelf and assign shelf ID

Each shelf has a unique shelf ID, ensuring its distinction in your storage system setup.

About this task


- A valid shelf ID is 01 through 99.
- You must power cycle a shelf (unplug both power cords, wait a minimum of 10 seconds, and then plug them back in) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, and then connecting the power cords to power sources on different circuits.

The shelf automatically powers on and boots when plugged in.

2. Remove the left end cap to access the shelf ID button behind the faceplate.



1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

3. Change the first number of the shelf ID:

- a. Insert the straightened end of a paperclip or narrow tipped ball point pen into the small hole to gently press the shelf ID button.
- b. Gently press and hold the shelf ID button until the first number on the digital display blinks, and then

release the button.

The number blinks within 15 seconds, activating shelf ID programming mode.



If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

- c. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

Each press and release duration can be as short as one second.

The first number continues to blink.

4. Change the second number of the shelf ID:

- a. Press and hold the button until the second number on the digital display blinks.

It can take up to three seconds for the number to blink.

The first number on the digital display stops blinking.

- b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

It can take up to three seconds for the number to stop blinking.

Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.

- a. Unplug the power cord from both power supplies on the shelf.
- b. Wait 10 seconds.
- c. Plug the power cords back into the shelf power supplies to complete the power cycle.

The power supply powers on as soon as you plug in the power cord. Its bicolored LED should illuminate green.

7. Replace the left end cap.

== Step 2: Power on the controller nodes


After you've turned on your storage shelves and assigned them unique IDs, turn on the power to the storage controller nodes.

Steps

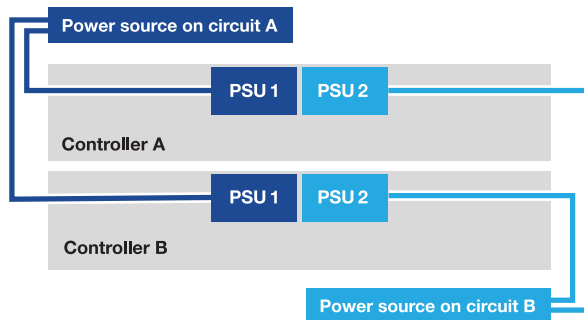
1. Connect your laptop to the serial console port. This allows you to monitor the boot sequence when the controllers are powered on.
 - a. Set the serial console port on the laptop to 115,200 baud with N-8-1.

See your laptop's online help for instructions on how to configure the serial console port.

- b. Connect the console cable to the laptop, and connect the serial console port on the controller using the console cable that came with your storage system.
- c. Connect the laptop to the switch on the management subnet.



2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.



- The system begins to boot. Initial booting may take up to eight minutes.
 - The LEDs flash on and the fans start, indicating that the controllers are powering on.
 - The fans may be noisy at start-up, which is normal.
4. Secure the power cords using the securing device on each power supply.

What's next?

After powering on your AFX 1K storage system, complete the setup based on your deployment option:

- For dedicated AFX 1K storage system deployments, go to [Set up an AFX cluster](#).
- For AFX 1K with AI Data Engine (AIDE) deployments, install AIDE third-party servers or NetApp-provided data compute nodes:
- For third-party server deployments, complete [hardware installation for AI Data Engine on third-party servers](#), then [install AIDE software on third-party servers](#).
- For NetApp DCN deployments, complete [hardware installation for AI Data Engine on NetApp-provided data compute nodes](#).

= Set up your AFX storage system ONTAP cluster

:icons: font

:relative_path: ./install-setup/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

After your AFX hardware has been installed, you can complete the ONTAP cluster setup. This process involves related configuration phases that you should perform sequentially. As part of this, you can optionally add the data compute nodes (DCNs) running the AI Data Engine (AIDE), enabling the AFX storage cluster and DCN cluster to operate as a single system that's managed through the same ONTAP control plane.

== Prepare to set up an AFX storage system

Before setting up your ONTAP cluster, you should review the core AFX requirements and additional AIDE considerations to ensure you have all the necessary information and resources to complete the setup process successfully.

=== Core AFX requirements

There are several core requirements to consider when setting up your AFX cluster.

Hardware installation

Make sure you've completed the hardware installation of your AFX system. For more information, see [Install your AFX system](#).

Basic IP addressing

You need the following IP addressing configuration:

- Cluster management IP address

The cluster management IP address is a unique IPv4 address for the cluster management interface used by the cluster administrator to access the admin SVM and manage the cluster. You can obtain this IP address from the administrator responsible for assigning IP addresses in your organization.

- Network subnet mask

During cluster setup, ONTAP requires a set of network interfaces appropriate for your configuration. You can adjust the recommendation if necessary.

Additional network configuration

You'll also need the following related networking information:

- Network gateway IP address
- DNS domain names
- DNS name server IP addresses
- NTP server IP addresses
- Subnet mask

=== Additional AIDE considerations

When installing and setting up AFX for the first time, you can optionally include NetApp AI Data Engine (AIDE). AIDE supports two deployment models: AIDE software on third-party servers or NetApp-provided data compute nodes (DCNs). If you plan to set up AIDE, complete the prerequisites for your deployment model first:

- For third-party server deployments, complete [hardware installation for AI Data Engine on third-party servers](#), then [install AIDE software on third-party servers](#).

- For NetApp DCN deployments, complete [hardware installation for AI Data Engine on NetApp-provided data compute nodes](#).
- Obtain the addresses and any related network configuration for the compute nodes.

== Perform initial cluster setup

You can connect your laptop device to the AFX cluster and set several global configuration values.

About this task

There are four areas of the AFX cluster that must be initially configured. The first three are required while the last one is optional.

Steps

1. Discover your cluster network.
 - a. Connect your laptop to the management switch and access the network computers and devices.
 - b. Open File Explorer.
 - c. Select **Network**; then right-click and select **Refresh**.
 - d. Select either ONTAP icon; then accept any certificates displayed on your screen.

The System Manager user interface is displayed.

2. Set the administrator password.

Provide and verify the password for the `admin` account. Select **Continue**.

3. Configure the IP addresses for the AFX cluster and controller nodes.

Provide IP addresses and the subnet masks.

4. If you're installing AI Data Engine with your AFX storage system, provide the IP addresses for the three compute nodes.
5. Configure the network services including DNS and NIS. Select **Continue**.

Define the details for your DNS and NTP servers.

6. Optionally set up encryption.

You can define the details for cluster encryption. Select **Continue**.



For information on how to create a cluster for a non-Windows environment, see [Create an ONTAP cluster and join nodes](#).

What's next

You'll be redirected to the System Manager sign-in page. Perform the steps described in [\[Complete cluster setup\]](#).

== Complete cluster setup

After the initial configuration has been performed you can complete the ONTAP cluster setup using System Manager.

About this task

There are three areas of the AFX system ONTAP cluster configured during setup. Complete all three if possible but only the first is required.

Before you begin

You need to have the following information:

- VLAN configuration details.
- NAS and/or S3 configuration details.

Steps

1. Sign in to System Manager using the administrator account you provided during initial cluster setup. Notice the popup window at the top right with three configuration options.
2. Select **VLAN and tagging** and select the network options appropriate for your environment.
3. Select **Network Services** and configure the client access protocols for the default data SVM.
4. Select **Data container** and create a volume or S3 bucket.

What's next

- If you included AIDE as part of AFX cluster setup, complete the additional AIDE configuration steps. Authentication configuration applies to all AIDE deployments. Licensing is required for NetApp DCN deployments that use licensed AIDE features; third-party server deployments are metadata-focused and do not require an AIDE license. See [Set up AI Data Engine](#) for more information.
- You should [Prepare to administer AFX](#) before using your AFX cluster in a production environment.

== Related information

- [Learn about the NetApp AI Data Engine](#)
- [Configure an AFX SVM](#)
- [Prepare to administer AFX](#)

= Prepare to administer your AFX storage system

:icons: font

:relative_path: ./get-started/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

Before deploying AFX in a production environment, it's essential to understand the administrative structure and configuration options. This ensures secure, efficient, and effective management of your AFX cluster.

== Understand storage virtual machines

A storage virtual machine (SVM) is an isolated server or tenant environment within an ONTAP cluster. You can configure an SVM to serve data to the connected clients. You should be familiar with the capabilities and characteristics of the AFX SVMs.

Types of SVMs

An AFX system cluster hosts several different types of SVMs. A **data SVM** is used to serve data to the clients and is the one type an AFX administrator can directly access and configure. There is one data SVM created by default when you set up and initially deploy an AFX cluster, but you can create additional data SVMs if needed. When referring to an SVM in this documentation, a data SVM is implied unless otherwise

noted.

Administrative control

SVMs can be used to establish and enforce isolation of your data and applications. This can be useful when there are many different groups with a larger organization. Administrative control can be delegated to the SVMs to establish policies related to data access, security, and protection.

Accounts and RBAC roles

There are two levels of authentication and authorization with AFX: cluster level and SVM level. In addition to the cluster accounts, every SVM has its own distinct set of users and roles. In most situations, using the cluster level accounts is adequate. But depending on your environment, you might need to configure and use the more restrictive SVM accounts and roles as well. See [Additional AFX SVM administration](#) for more information.

SVM-scoped resources

AFX resources and configurable entities are associated either with the cluster or a specific SVM. There are many resources with an SVM scope, including volumes and buckets as well as the SVM user accounts and RBAC roles.

Dedicated network interfaces

Each SVM has its own dedicated set of network interfaces. For example, separate LIFs are allocated to an SVM for management and client access.

== Two AFX administrative levels

The administrative ONTAP tasks you perform with AFX generally fall into two different categories. Some tasks apply to the ONTAP cluster as a whole, while other tasks apply to a specific SVM. This results in two-tier administrative model.

It's important to note that these levels describe how the administrative tasks are organized and assigned, and not necessarily how the associated security is configured. For example, while a cluster administrator account is needed to perform cluster level administration, it can also be used for SVM administration.



Beginning with ONTAP 9.19.1, SVM administrators can monitor and manage their own SVMs using a dedicated System Manager dashboard. Learn more about the [ONTAP System Manager SVM administrator dashboard](#)

Cluster administrator

The cluster administrator has complete control of the AFX cluster including all the SVMs. The AFX cluster administrative level includes only the tasks that a cluster admin can perform and not any of the SVM-specific administration tasks. See [Administer your cluster](#) for more information.

SVM administrator

An SVM administrator role has control of a specific SVM and so is more restricted compared to the cluster administrator. SVM administration involves performing tasks with objects and resources that have an SVM scope, such as creating a volume. See [Administer your storage VMs and data](#) for more information.

== Three administrative interfaces

Like AFF and FAS systems, AFX has three administrative interfaces. The LIF (or IP address) you need to use varies based on the administrative interface and your environment.



The System Manager user interface is preferred for most administrative tasks. You should use an administrator account unless otherwise indicated.

Interface	Description
System Manager	This is a graphical user interface available through a web browser. It's easy to use and provides access to most of the capabilities customers need. Accessing AFX through System Manager provides the simplest experience for the majority of ONTAP cluster and SVM administration needs.
Command line interface	The ONTAP CLI is accessible using SSH. Depending on your account, you can access the cluster management LIF or SVM management LIF. The CLI is more difficult to use but is more robust. It's preferred, and sometimes required, for advanced administration tasks.
REST API	AFX includes a REST API you can use to automate the administration of your AFX cluster. The API shares many of the same calls available with the Unified ONTAP personality REST API with modifications to support the unique AFX features.

== Learn to search, filter, and sort information in System Manager

The System Manager user interface includes a robust set of features enabling you to access and display the information you need. Learning to use these capabilities will help you to better administer the AFX storage system. See [Search, filter, sort information in System Manager](#) for more information.

== Access the ONTAP CLI

While you can use System Manager for most AFX administration, there are some tasks you can only perform using the ONTAP command line interface.

About this task

You can access the ONTAP CLI through the secure shell (SSH). The CLI has multiple privilege levels that determine the commands and command parameters available to you. The `admin` level is the least privileged and the default when you sign in. You can elevate the privilege of your session to `advanced` if needed using the `set` command.

Before you begin

You'll need the following:

- IP address or domain name of the cluster or SVM management LIF
- Account credentials
- SSH client on your local workstation

Steps

1. Use SSH to connect to your AFX cluster, for example:

```
ssh admin@10.69.117.24
```

2. Provide the account password.

3. Display the command directories at the top of the hierarchy:

?

4. Elevate the privilege level of your session from `admin` to `advanced`:

```
set -privilege advanced
```

== Working with ONTAP HA pairs

As with Unified ONTAP, AFX cluster nodes are configured in high-availability (HA) pairs for fault tolerance and nondisruptive operations. HA pairing provides the ability for storage operations to stay online in the event of a node failure, such as a storage failover. Each node is partnered with another node to form a single pair. This is generally done using a direct connection between the two node's NVRAM modules.

With AFX, a new HA VLAN is added to the backend cluster switches to enable NVRAM modules to stay connected between the HA partner nodes. HA pairs are still used with the AFX system, but there is no longer a need for the partner nodes to be directly connected.

== AFX cluster deployment limitations

There are several limitations, including minimums and maximums, enforced by AFX when configuring and using your cluster. These limits fall into several categories.



Before increasing an operational parameter of your AFX storage system, you should assess how it might impact your applications as well as the cluster as a whole. For example, while expanding the number of controller nodes increases overall capacity, it could also increase traffic in the cluster network. It's important to understand the implication of these configuration changes and how they might impact your environment.

Controller nodes per cluster

Each AFX cluster must have at least four nodes. The maximum number of nodes varies based on the ONTAP release; for example the 9.19.1 RC release allows up to 32 nodes per AFX cluster.

Storage capacity

This is the total capacity across all the SSD disks in the cluster Storage Availability Zone (SAZ). The maximum storage capacity varies based on the ONTAP release.

Cluster switches

You need at least two switches in your cluster storage network. The maximum allowed is determined based on the total number of controller nodes in the cluster.

You should review the details available at the [NetApp Hardware Universe and Interoperability Matrix Tool](#) to determine the capabilities of your AFX cluster. Also review [Changes to ONTAP limits and defaults](#) in AFX release notes section.

== Confirm AFX system health

Before performing any AFX administration tasks, you should check the health of the cluster.



You can check the health of your AFX cluster at any time, including when you suspect an operational or performance issue.

Before you begin

You'll need the following:

- Cluster management IP address or FQDN
- Administrator account for the cluster (username and password)

Steps

1. Connect to System Manager using a browser:

`https://$FQDN_IPADDR/`

Example

`https://10.61.25.33/`

2. Provide the administrator username and password and select **Sign in**.
3. Review the system dashboard and cluster status including cabling. Also notice the *navigation pane* on the left.

[View dashboard and cluster status](#)

4. Display the system events and audit log messages.

[View AFX events and audit log](#)

5. Display and note any **Insight** recommendations.

[Use Insights to optimize AFX cluster performance and security](#)

== Quick start for creating and using an SVM

After installing and setting up the AFX cluster, you can begin performing the administration tasks typical of most AFX deployments. Here are the high-level steps needed to begin sharing data with clients.

1

Display the available SVMs

[Display](#) the list of SVMs and determine if there's one you can use.

2

Optionally create an SVM

[Create](#) an SVM to isolate and protect your application workloads and data if an existing SVM is not available.

3

Configure your SVM

[Configure](#) your SVM and prepare for client access.

4

Prepare to provision storage

[Prepare](#) to allocate and manage your data.

== Related information

- [NetApp Hardware Universe](#)
- [NetApp Interoperability Matrix Tool](#)
- [Interoperability Matrix Tool overview](#)
- [ONTAP user interfaces](#)
- [Set the privilege level in the ONTAP CLI](#)
- [Learn about cluster administration with the ONTAP CLI](#)
- [Types of SVMs in an ONTAP cluster](#)
- [FAQ for AFX storage systems](#)

= Administer your cluster

= Monitor cluster processes

= View the AFX storage system dashboard

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjv8c/source/./install-setup/./media/

You can sign in to System Manager to access the AFX dashboard and display the cluster status. This is a good first step before beginning your AFX administrative tasks or if you suspect an operational issue.

Before you begin

You'll need the following:

- IP address or domain name of the cluster management LIF
- Administrator account credentials

Steps

1. Connect to System Manager using a browser and the cluster management IP address:

```
https://$FQDN_IPADDR/
```

Example

```
https://10.61.25.33/
```

2. Provide the username and password for the administrator account and select **Sign in**.
3. Select **Dashboard** in the left navigation pane and review the tiles on the page including the cluster **Health**.
4. In the navigation pane, select **Cluster** and then **Overview**.
5. Review the cluster name, version, ONTAP personality and the other details.
6. At the top of the overview page, select **Cabling** for a visual display of the cluster hardware and connections.
7. In the navigation pane, select **Events & Jobs** and then **System alerts** to display and review the system alerts.

= View Insights to optimize your AFX storage system

:icons: font
:relative_path: ./administer/
:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

You can use the *Insights* feature of ONTAP System Manager to display suggested configuration updates that align with NetApp best practices. These changes can optimize the security and performance of your AFX cluster.

For example, the Autonomous Ransomware Protection (ARP) feature is available with AFX and provides anti-ransomware protection. Insights will inform you if ARP is not configured.

About this task

Each of the insights is presented as a separate tile or card on the page that you can choose to implement or dismiss. You can also select the associated documentation link to learn more about a specific technology.

Steps

1. In System Manager, select **Analysis** and then **Insights**.
2. Review the available recommendations.

What's next

Perform any of the recommended actions to implement AFX configuration best practices.

= Monitor AFX storage system cluster performance and storage utilization

:icons: font
:relative_path: ./administer/
:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

You can display a high-level overview of the performance and storage utilization of your AFX cluster.

== Storage capacity

Monitor your storage utilization to ensure you have sufficient storage capacity for current and future workloads.

=== Monitor high-level cluster storage utilization

In ONTAP System Manager, you can view a high-level display of the storage utilization for the entire cluster.

Steps


1. In System Manager, select **Dashboard** in the navigation pane.
2. Locate the **Capacity** tile and view the physical storage available and used.
3. Select **History** to access Active IQ to view the historical data.

=== Monitor storage availability zone capacity details

Beginning with ONTAP 9.19.1, you can view the capacity details of your storage availability zone including space thresholds, space consumed by snapshots, and data reduction. You can also set space usage thresholds to the percentage at which you want ONTAP to generate warning and critical alerts.

Steps

1. In System Manager, select **Cluster > Overview** and then select **Availability zone**.

2. To modify the warning or critical space thresholds, select  under the threshold that you want to edit.

== Cluster performance

System Manager provides a detailed overview of the AFX cluster performance.

Steps

1. In System Manager, select **Analysis** and then **Performance**.
2. Review the cluster performance summary at top including latency and throughput.
3. Under the **Top actors** tab, select the desired SVM and then **Enable activity tracking** as needed.
4. Under the **Volume performance** tab, view the performance details of a specific volume.

== Related information

- [Additional AFX cluster administration](#)

= View AFX storage system events and audit log

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

You can review the events and audit log messages generated by AFX to track internal processing and diagnose potential problems. The AFX system can be configured to forward this information, along with other related data, for additional processing and archival.

== Events

The event messages provide a valuable record of system activity. Each event includes a description and unique identifier along with a recommended action.

1. In System Manager, select **Events & jobs** and then **Events**.
2. Review and respond to the recommended actions at the top of the page, such as enabling automatic update.
3. Select the **Events log** tab to display a list of the messages.
4. Select an event message to examine it in more detail, including the sequence number, description, event, and recommended action.
5. Optionally select the **Active IQ suggestions** tab and register with Active IQ to get detailed risk information for the cluster.

== Audit log


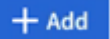
The audit log includes a record of system activity based on the use of access protocols such as HTTP.

1. In System Manager, select **Events & jobs** and then **Audit logs**.
2. Select **Settings** to enable or disable the operations that are tracked.
3. Optionally select **Manage audit destinations**; review [\[Manage notifications\]](#) for more information.

== Manage notifications

There are several types of notifications supported by AFX that you can forward.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Navigate to **Notification management** and select .
3. Select the appropriate action to view or configure the destinations used by AFX. For example, to configure:
 - a. *Event destinations*: select **View event destinations**
 - b. *Audit log destinations*: select **View audit destinations**
4. Select  as appropriate and provide the destination information.
5. Select **Save**.

Related information

- [ONTAP Event, performance, and health monitoring](#)

= View AFX storage system jobs

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

AFX includes an internal platform to run background jobs based on your configuration and administrative actions. These jobs can be long-running AFX components or short-lived processes executed in response to administrative tasks or REST API requests. You can display and monitor the jobs as needed.

Steps

1. In System Manager, select **Events & Jobs** and then **Jobs**.
2. Customize the display as well as search and download the job information as needed.

= Manage networking and security

= Manage AFX storage system cluster networking

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

You need to configure the network of your AFX storage system. The networking environment supports several scenarios including clients accessing data at the SVMs and intercluster communication.



Creating a network resource is an important first step. You also need to perform additional administrative actions, such as editing or deleting network definitions, as needed.

== Create a broadcast domain

A broadcast domain simplifies management of your cluster network by grouping ports that are part of the same layer two network. The storage virtual machines (SVMs) can then be assigned ports in the group for data or management traffic.

There are several broadcast domains created during cluster setup, including:

Default

This broadcast domain contains ports in the “Default” IPspace. These ports are used primarily to serve data. Cluster management and node management ports are also included.

Cluster

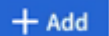
This broadcast domain contains ports in the “Cluster” IPspace. These ports are used for cluster communication and include all the cluster ports from all nodes in the cluster.

You can create additional broadcast domains after your cluster has been initialized. When you create a broadcast domain, a failover group that contains the same ports is automatically created.

About this task

The maximum transmission unit (MTU) value of the ports defined for a broadcast domain are updated to the MTU value set in the broadcast domain.

Steps

1. In System Manager, select **Network** and then **Overview**.
2. Under **Broadcast domains**, select  .
3. Provide the name of the broadcast domain or accept the default.

All broadcast domain names must be unique within an IPspace.

4. Provide the maximum transmission unit (MTU).

The MTU is the largest data packet that can be accepted in the broadcast domain.

5. Choose the desired ports and select **Save**.


== Create an IPspace

An IPspace is an administrative domain for IP addresses and related network configuration. These spaces can be used to support your SVMs through isolated administration and routing. For example, they are useful when clients have overlapping IP addresses from the same IP address and subnet range.



You must have an IPspace before you can create a subnet.

Steps

1. In System Manager, select **Network** and then **Overview**.
2. Under **IPspaces**, select  .
3. Provide the name of the IPspace or accept the default.

All IPspace names must be unique within a cluster.

4. Select **Save**.

What's next

You can use the IPspace to create a subnet.

== Create a subnet

A subnetwork or subnet enforces a logical division of the IP address space in your network. It enables you to allocate dedicated blocks of IP addresses for the creation of a network interface (LIF). Subnets simplify LIF

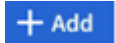
creation by enabling you to use the subnet name instead of a specific IP address and network mask combination.

Before you begin

You must have a broadcast domain and IPspace where the subnet will be defined. Also note:

- All subnet names must be unique within a specific IPspace.
- The IP address range used for a subnet cannot overlap with the IP addresses of other subnets.

Steps

1. In System Manager, select **Network** and then **Overview**.
2. Under the **Subnets** tab, select  **+ Add** .
3. Provide the configuration details, including the name of the subnet, IP address details, and broadcast domain.
4. Select **Save**.

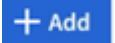
What's next

The new subnet will simplify the creation of your network interfaces.

== Create a network interface

A logical network interface (LIF) consists of an IP address and related network configuration parameters. It can be associated with a physical or logical port and is typically used by the clients to access data provided by an SVM. LIFs provide resiliency in the event of a failure and can migrate among the node ports so communication is not interrupted.

Steps

1. In System Manager, select **Network** and then **Overview**.
2. Under the **Network interfaces** tab, select  **+ Add** .
3. Provide the configuration details, including the name of the interface, interface type, allowed protocols, and IP address details.
4. Select **Save**.

== Related information

- [Manage AFX Ethernet ports](#)
- [Learn about ONTAP broadcast domains](#)
- [Learn about ONTAP IPspace configuration](#)
- [Learn about subnets for the ONTAP network](#)
- [Network architecture overview](#)

= Manage AFX storage system Ethernet ports

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjv8c/source/./install-setup/./media/

The ports used by the AFX system provide a foundation for network connectivity and communication. There are several options available to customize the layer two

configuration of your network.

== Create a VLAN

A VLAN consists of switch ports grouped together into a broadcast domain. VLANs enable you to increase security, isolate potential problems, and limit available paths within your IP network infrastructure.

Before you begin

The switches deployed in the network must either comply with IEEE 802.1Q standards or have a vendor-specific implementation of VLANs.

About this task

Note the following:

- You can't create a VLAN on an interface group port without any member ports.
- When you configure a VLAN over a port for the first time, the port might go down, resulting in a temporary disconnection of the network. Subsequent VLAN additions to the same port do not affect the port state.
- You should not create a VLAN on a network interface with the same identifier as the native VLAN of the switch. For example, if the network interface e0b is on native VLAN 10, you should not create a VLAN e0b-10 on that interface.

Steps

1. In System Manager, select **Network** and then **Ethernet ports**.
2. Select **+ VLAN**.
3. Provide the configuration details, including the id, broadcast domain, and ports across the desired nodes.

The VLAN can't be attached to a port hosting a cluster LIF or to ports assigned to the cluster IPspace.

4. Select **Save**.

Result

You have created a VLAN to increase security, isolate problems, and limit available paths within your IP network infrastructure.

== Create a LAG

A link aggregate group (LAG) is a technique that combines multiple physical network connections into a single logical connection. You can use it to increase the bandwidth and provide redundancy between nodes.

Steps

1. In System Manager, select **Network** and then **Ethernet ports**.
2. Select **Link aggregate group**.
3. Provide the configuration details, including the node, broadcast domain, ports, mode, and load distribution.
4. Select **Save**.

== Related information

- [Manage AFX cluster networking](#)
- [Learn about ONTAP network port configuration](#)
- [Combine physical ports to create ONTAP interface groups](#)

```
= Prepare AFX storage system authentication services
:icons: font
:relative_path: ./administer/
:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/
```

You need to prepare the authentication and authorization services used by the AFX system to the user account and role definitions.


== Configure LDAP

You can configure a Lightweight Directory Access Protocol (LDAP) server to maintain authentication information at a central location.

Before you begin

You must have generated a certificate signing request and added a CA-signed server digital certificate.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Select  next to **LDAP**.
3. Select **+ Add** and provide the name or IP address of the LDAP server.
4. Provide the necessary configuration information, including the schema, base DN, port, and binding.
5. Select **Save**.


== Configure SAML authentication

Security Assertion Markup Language (SAML) authentication enables users to be authenticated by a secure identity provider (IdP) instead of providers using other protocols such as LDAP.

Before you begin

- The identity provider you plan to use for remote authentication must be configured. See the provider documentation for configuration details.
- You must have the URI of the identity provider.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Select  under **Security** next to **SAML authentication**.
3. Select **Enable SAML authentication**.
4. Provide the **IdP URL** and the **Host system** IP address and select **Save**.

A confirmation window displays the metadata information, which has been automatically copied to your clipboard.

5. Navigate to the IdP system you specified and copy the metadata from your clipboard to update the system metadata.
6. Return to the confirmation window in System Manager and select **I have configured the IdP with the host URI or metadata**.
7. Select **Logout** to enable SAML-based authentication.

The IdP system will display an authentication screen.

== Related information

- [Manage AFX cluster users and roles](#)
- [Configure SAML authentication for remote ONTAP users](#)
- [Authentication and access control](#)

= Manage AFX storage system cluster users and roles

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

You can define user accounts and roles based on the authentication and authorization services available with AFX.


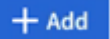


Each ONTAP user needs to have one role assigned. A role includes privileges and determines what actions the user is able to perform.

== Create an account role

Roles for cluster administrators and storage VM administrators are automatically created when your AFX cluster is set up and initialized. You can create additional user account roles to define specific functions that users assigned to the roles can perform on your cluster.


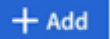
Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. In the **Security** section, next to **Users and roles**, select .
3. Under **Roles**, select .
4. Provide the name of the role and the attributes.
5. Select **Save**.

== Create a cluster account

You can create a cluster-level account to use when performing cluster or SVM administration.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. In the **Security** section, select  next to **Users and roles**.
3. Select  under **Users**.
4. Enter a username and then select the role for the user.

The role should be appropriate for the user. For example, the **admin** role is able to perform the full range of configuration tasks on your cluster.

5. Select the user login method and the authentication method; this will typically be **Password**.
6. Enter a password for the user.
7. Select **Save**.

Result

A new account is created and available for use with your AFX cluster.

== Related information

- [Prepare authentication services](#)
- [Additional AFX SVM administration](#)

= Manage certificates on an AFX storage system

:icons: font

:relative_path: ./administer/





:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

Depending on your environment, you'll need to create and manage digital certificates as part of administering AFX. There are several related tasks you can perform.

== Generate a certificate signing request

To get started using a digital certificate, you need to generate a certificate signing request (CSR). A CSR is used to request a signed certificate from a certificate authority (CA). As part of this, ONTAP creates a public/private key pair and includes the public key in the CSR.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Under **Security** and next to **Certificates**, select 
3. Select  **Generate CSR**.
4. Provide the subject common name and country; optionally provide the organization and organizational unit.
5. To change the default values which will define the certificate, select   **More options** and make the desired updates.
6. Select **Generate**.



Result

You have generated a CSR which can be used to request a public key certificate.

== Add a trusted certificate authority

ONTAP provides a default set of trusted root certificates for use with Transport Layer Security (TLS) and other protocols. You can add additional trusted certificate authorities as needed.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Under **Security** and next to **Certificates**, select 
3. Select the tab **Trusted certificate authorities** and then select  **Add** .
4. Provide the configuration information, including the name, scope, common name, type, and certificate details; you can import the certificate instead by selecting **Import**.
5. Select **Add**.

Result



You have added a trusted certificate authority to your AFX system.

== Renew or delete a trusted certificate authority

Trusted certificate authorities must be renewed annually. If you do not want to renew an expired certificate, you should delete it.

Steps

1. Select **Cluster** and then **Settings**.
2. Under **Security** and next to **Certificates**, select →.
3. Select the tab **Trusted certificate authorities**.
4. Select the trust certificate authority that you want to renew or delete.
5. Renew or delete the certificate authority.

To renew the certificate authority, do this:	To delete the certificate authority, do this:
<ol style="list-style-type: none">a. Select  and then select Renew.b. Enter or import the certificate information and select Renew.	<ol style="list-style-type: none">a. Select  and then select Delete.b. Confirm that you want to delete and select Delete.

Result

You have renewed or deleted an existing trusted certificate authority on your AFX system.

== Add a client/server certificate or local certificate authority

You can add a client/server certificate or a local certificate authority as part of enabling secure web services.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Under **Security** and next to **Certificates**, select →.
3. Select either **Client/server certificates** or **Local certificate authorities** as needed.
4. Add the certificate information and select **Save**.

Result



You have added a new client/server certificate or local authorities to your AFX system.

== Renew or delete a client/server certificate or local certificate authorities

Client/server certificates and local certificate authorities must be renewed annually. If you do not want to renew an expired certificate or local certificate authorities, you should delete them.

Steps

1. Select **Cluster** and then **Settings**.
2. Under **Security** and next to **Certificates**, select →.
3. Select either **Client/server certificates** or **Local certificate authorities** as needed.
4. Select the certificate you want to renew or delete.
5. Renew or delete the certificate authority.

To renew the certificate authority, do this:	To delete the certificate authority, do this:
a. Select  and then select Renew . b. Enter or import the certificate information and select Renew .	Select  and then select Delete .

Result

You have renewed or deleted an existing client/server certificate or local certificate authority on your AFX system.

== Related information

- [Generate and install a CA-signed server certificate in ONTAP](#)
- [Manage ONTAP certificates with System Manager](#)

= Manage storage VMs

= Display the AFX storage system SVMs


:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

You can display the data storage VMs defined in your AFX cluster. Each SVM provides an isolated environment for organizing your data and providing client access.

Steps

1. In System Manager, select **Cluster** and then **Storage VMs**.
2. Hover over the desired SVM and select  to view the primary administrative options including starting and stopping the SVM.
3. Optionally select a specific SVM to view more details including overview, settings, replication, and file system.

Related information

- [Configure an AFX system SVM](#)
- [Understand storage virtual machines](#)

= Create an AFX storage system SVM

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

You can create an SVM to provide isolation and improve security. You might do this for different groups or projects within your organization.

About this task

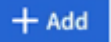
When you create an SVM, you must provide a name and configure at least one protocol for client access. After selecting a client protocol, you will be prompted for the networking configuration as well. You can change the SVM configuration as needed after it has been created.

Before you begin

You'll need the following:

- A minimum of four IP addresses
- Name of an IPspace

Steps

1. In System Manager, select **Cluster** and then **Storage VMs**.
2. Select  .
3. Provide a name for the SVM.
4. Select a protocol for client access and provide the configuration details as appropriate.
5. Add a network interfaces for the SVM including the IP addresses and subnet mask.
6. Under **Storage VM administration**, optionally:
 - a. Enable a maximum capacity and select a value
 - b. Manage an administrator account for the SVM
7. Select **Save**.

Related information

- [Configure an AFX system SVM](#)
- [Manage AFX system cluster networking](#)

= Configure an AFX storage system SVM

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

After you create an SVM, you can update the configuration based on your requirements and clients needs.

About this task

There are four access paths to the SVM configuration as reflected in the tabs on the landing page for a specific SVM. These include:

- Overview

This provides a quick dashboard overview of the current configuration details related to network interfaces and services, protocols, storage, and protection.

- Settings

You can access and update the entire SVM configuration as organized in several areas, such as protocols, services, policies, and security.

- Replication

This page provides a list of the current replication relationships defined for the SVM.

- File system

You can track the activity and analytics for the SVM

Before you begin

You need to decide which SVM you are interested in displaying and updating.

Steps

1. In System Manager, select **Cluster** and then **Storage VMs**.
2. Select the desired SVM and then the **Settings** tab.
3. Review the configuration options on the page; select and update the settings as desired.

= Migrate an AFX storage system SVM

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

You can migrate an SVM from one ONTAP cluster to another. SVM migration with AFX operates the same as with Unified ONTAP, although there are several interoperability considerations and restrictions. Refer to the Unified ONTAP documentation for details about performing an SVM migration.

== Interoperability considerations

Before planning and performing an SVM migration, you should be aware of the interoperability considerations including capabilities and limitations.

=== Use cases

Cluster administrators can relocate an SVM from a source cluster to a destination cluster. You might do this as part of capacity management and load balancing, or to allow for equipment upgrades or data center consolidations. Because the AFX storage system does not support in-place upgrades from Unified ONTAP, SVM migration is an important use case.

You can move your application workloads from a Unified ONTAP cluster to AFX clusters without disruption. In addition, SVMs can be migrated in other ways including from an AFX cluster to a Unified ONTAP cluster as well as among AFX clusters.

=== Version interoperability

The following table describes the allowable SVM migrations based on the ONTAP personality and release of the source and destination cluster.

Direction	Source version	Destination version
Unified to AFX	9.15.1 - 9.17.1	9.17.1
AFX to Unified	9.17.1	9.17.1
AFX to AFX	9.17.1	9.17.1

=== Prechecks

Unified ONTAP includes several prechecks that are also implemented with AFX. In addition, several new prechecks are added to flag features that aren't supported with AFX, including:

- FabricPool (volumes residing on composite aggregates)
- Thick provisioned volumes

=== Volume provisioning

The volumes are provisioned to balance their placement across the Storage Availability Zone (SAZ) of the AFX cluster.

Space guarantee

AFX does not support thick provisioning. A precheck is used to fail a migration if any volume in the SVM being migrated is thick provisioned.

Encryption

An AFX system supports NetApp volume encryption (NVE) but not NetApp aggregate encryption (NAE). Because of this, any NAE volumes at a Unified ONTAP cluster are converted to NVE volumes when migrated to AFX. The following table summarizes the compatibility and conversion.

Source volume	Destination volume
Plain text	Plain text
NVE	NVE
NAE	NVE

=== Additional restrictions

There are additional restrictions you should consider before migrating an SVM.

MetroCluster

The AFX storage system does not support NetApp MetroCluster. This creates a limitation when migrating an SVM. You cannot migrate an AFX SVM to or from an AFF or FAS system (or any NetApp system running the Unified ONTAP personality) that is configured to use MetroCluster. While these migration scenarios are not supported, they are also not explicitly blocked by the AFX prechecks and so you need to be careful not to attempt them.

== Related information

- [ONTAP SVM data mobility](#)
- [Compare AFX storage system to AFF and FAS systems](#)
- [FAQ for AFX storage systems](#)

= Support the cluster

= Manage AutoSupport for an AFX storage system cluster

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

AutoSupport is a NetApp technology you can use to proactively monitor the health of your AFX storage systems. It can automatically send messages to NetApp technical support, your internal support organization, or a support partner.

AutoSupport is enabled by default when you set up an AFX cluster and messages will be sent to NetApp technical support. To send messages to your internal support organization, you need to properly configure your cluster and provide a valid email host. AFX begins sending AutoSupport messages 24 hours after it is active.




You need to sign in to System Manager using a cluster administrator account to manage AutoSupport.

== Test AutoSupport connectivity

After you have set up your cluster, you should test your AutoSupport connectivity to verify that technical support can receive the messages generated by AutoSupport.



Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Next to **AutoSupport** select  and then **Test connectivity**.
3. Enter a subject for the AutoSupport message and select **Send test AutoSupport message**.

== Add AutoSupport recipients

You can optionally add members of your internal support organization to the list of email addresses that receive AutoSupport messages.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Next to **AutoSupport** select  and then **More options**.
3. Next to **Email**, select  and then **+ Add**.
4. Provide the email address for the recipient; for the recipient category, select:
 - **Partner** for your partners
 - **General** for members of your internal support organization
5. Select **Save**.


Result

The email addresses you have added will receive new AutoSupport messages for their specific recipient category.

== Send AutoSupport data

If a problem occurs with your AFX system, you should manually send the AutoSupport data. This can significantly decrease the amount of time it takes to identify and resolve the issue.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Next to **AutoSupport** select  and then **Generate and send**.
3. Provide a subject for the AutoSupport message.
4. Select **Send**.


Result

Your AutoSupport data is sent to technical support.

== Suppress support case generation

If you are performing an upgrade or maintenance on your AFX system, you might want to suppress the generation of AutoSupport support cases until your upgrade or maintenance is complete.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Next to **AutoSupport** select  and then **Suppress support case generation**.
3. Specify the number of hours to suppress the generation of support cases and the nodes you don't want cases generated for.
4. Select **Send**.


Result

AutoSupport cases will not be generated during the time you specified. If you complete your upgrade or maintenance before the specified time expires, you should resume support case generation immediately.

== Resume support case generation

If you have suppressed the generation of support cases during an upgrade or maintenance window, you should resume support case generation immediately after your upgrade or maintenance is complete.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Next to **AutoSupport** select  and then **Resume support case generation**.
3. Select the nodes for which you want to resume AutoSupport case generation.
4. Select **Send**.

Result

AutoSupport cases will be autogenerated for your AFX system as needed.

== Related information

- [Learn about ONTAP AutoSupport](#)
- [Prepare to use ONTAP AutoSupport](#)

= Submit and view support cases for an AFX storage system

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

If you have an issue that requires assistance, you can use ONTAP System Manager to submit a case to technical support. You can also use ONTAP System Manager to view cases that are in progress or closed.

Before you begin

You need to be [registered with Active IQ](#) to view support cases for your AFX storage system.

Steps

1. To create and submit a new support case, in System Manager select:

- a. **Cluster** and then **Support**
 - b. **Go to NetApp Support**
2. To view a previously submitted case, in System Manager select:
- a. **Cluster** and then **Support**
 - b. **View my cases**

Related information

- [View and submit support cases with ONTAP System Manager](#)

= Upgrade and maintain the cluster

= Expand an AFX storage system cluster

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

You can expand the compute capacity of an AFX cluster independent of the storage capacity. The expansion is performed without disruption and increases performance linearly as volumes are rebalanced across the nodes. This feature is a significant benefit as you adjust to the ongoing needs of your AFX system users.

== Prepare to expand a cluster

Before expanding an AFX cluster, you should be familiar with the basic requirements and general approach to troubleshooting.

=== Requirements

You need the credentials for a cluster administrator account and be able to connect to the ONTAP CLI using SSH. When expanding a cluster, you must add an even number of nodes and adhere to the size limitations of your AFX system based on the release.

=== Troubleshooting

There are a few concepts and troubleshooting scenarios you should be aware of as you perform the cluster expansion.

==== Automatic volume rebalancing

Automated Topology Management (ATM) is an internal AFX system component that detects allocation imbalances and rebalances volumes across the cluster nodes. It relies on the Zero Copy Volume Move (ZCVM) technology to relocate volumes using metadata updates instead of copying the data. ZCVM is the default volume move technology available with AFX storage systems.

==== Possible troubleshooting scenarios

There are several scenarios you might need to investigate during the volume moves associated with the expansion of an AFX cluster.

Volumes are not being moved by ATM

This can occur when the cluster is already in balance or when there are no eligible volumes to move.

Confusion about how or when ATM should be active

It may appear that volumes aren't distributed as quickly as expected. ATM attempts to detect and respond to hardware events every five minutes. In the worst case, a rebalance operation is launched 40 minutes after the last one completed.

==== CLI commands

There are several commands you can use to monitor a cluster expansion operation.

- `volume move show`
- `volume move show -instance`

You should contact NetApp support for additional assistance as needed.

== Add nodes to expand a cluster

This procedure describes how to add a pair of nodes to an existing cluster and can be adapted to other deployment environments. You'll need to use both the ONTAP CLI and System Manager administrative interfaces.

Steps

1. Connect to the ONTAP CLI and set advanced privilege level:

```
afx> set advanced
```

2. Display the volume locations of the current nodes; note the number of volumes per node:

```
afx> vol show -fields node,size,constituent-count -is-constituent true -node *
```

3. Display the cluster interconnect IP addresses and save for use in later steps:

```
afx> net int show -role cluster
```

4. Log into the service processor of each node you wish to add to the cluster.
5. From the prompt, type **system console** to access the node's console.
6. Boot the node to display the boot menu prompt:

```
LOADER> boot_ontap menu
```

If the menu does not load, use the **Ctrl+C** technique to access the boot menu.

7. Select one of the boot options from the menu as appropriate; if prompted type **yes** to continue.

If you get sent back to LOADER from here, type **boot_ontap** at the LOADER prompt.

8. Use the cluster setup wizard to configure a node management LIF, subnet, and gateway.

This configuration will be used by System Manager to detect the node to be added to the cluster. Enter the values as prompted, including port, IP address, netmask, and default gateway.

9. Press **CTL+C** to access the CLI.
10. Modify the cluster interconnect addresses so they're routable in your network; use the configuration

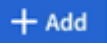
appropriate for your environment:

```
afx> net int show -role cluster
```

```
afx> net int modify -vserver Cluster -lif clus1 -address 192.168.100.201
```

```
afx> net int modify -vserver Cluster -lif clus2 -address 192.168.100.202
```

This step is only needed if the other interfaces do not use the 169.254.x.x addresses that ONTAP auto creates.

11. Repeat the above steps on the other AFX node controller.
12. Access the System Manager using the cluster management IP address.
13. In System Manager, select **Cluster** and then **Overview**; select the **Nodes** tab.
14. Locate the section **Not part of this cluster**; select  .
 - If the nodes were discovered before the cluster interconnect IP addresses are changed, you'll need to re-discover the nodes by exiting the window and navigating back.
 - You can optionally use the CLI to add the nodes instead of System Manager; see the command `cluster add-node`.
15. Provide the configuration details in the **Add nodes** menu; you can add management IP addresses manually or using a subnet.
16. Connect to the ONTAP CLI to monitor the status of the node add operation:

```
afx> add-node-status
```

17. After the operations have completed, confirm the volume placement across all nodes; issue the command once for each node using the appropriate node name:

```
afx> set advanced
```

```
afx> vol show -fields node,size,constituent-count -is-constituent true -node  
NODE_NAME
```

Result

- Adding new nodes to the cluster is nondisruptive.
- Volume moves should happen automatically.
- Performance will scale linearly.

== Related information

- [Prepare to administer your AFX system](#)
- [FAQ for ONTAP AFX storage systems](#)
- [NetApp Support Site](#)

= Upgrade ONTAP on an AFX storage system

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjv8c/source/./install-setup/./media/

When you upgrade your ONTAP software on your AFX system, you can take advantage of new and enhanced ONTAP features that can help you reduce costs, accelerate critical workloads, improve security, and expand the scope of data protection available to your organization.



AFX storage systems do not support [ONTAP revert](#) operations.

ONTAP software upgrades for AFX storage systems follow the same process as upgrades for other ONTAP systems. If you have an active SupportEdge contract for Active IQ Digital Advisor (also known as Digital Advisor), you should [prepare to upgrade with Upgrade Advisor](#). Upgrade Advisor provides intelligence that helps you minimize uncertainty and risk by assessing your cluster and creating an upgrade plan specific to your configuration. If you don't have an active SupportEdge contract for Active IQ Digital Advisor, you should [prepare to upgrade without Upgrade Advisor](#).

After you prepare for your upgrade, it is recommended that you perform upgrades using [automated non-disruptive upgrade \(ANDU\) from System Manager](#). ANDU takes advantage of ONTAP's high-availability (HA) failover technology to ensure that clusters continue to serve data without interruption during the upgrade.

Related information

- [Learn about ONTAP upgrade](#).

== Execute automated pre-upgrade checks before a planned upgrade

You don't have to be in the process of upgrading your ONTAP software to execute the ONTAP automated upgrade pre-checks. Executing the pre-upgrade checks independently of the ONTAP automated upgrade process allows you to see which checks are performed against your cluster and gives you a list of any errors or warnings that should be corrected before you begin the actual upgrade.

Beginning with ONTAP 9.19.1, errors and warning are displayed separately to help you easily identify issues that must be resolved before you upgrade.

[Learn to execute ONTAP automated pre-upgrade checks](#).

= Update firmware on an AFX storage system

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjv8c/source/./install-setup/./media/

ONTAP automatically downloads and updates firmware and system files on your AFX storage system by default. If you want to view the recommended updates before they are downloaded and installed, you can disable automated updates. You can also edit update parameters to show you notifications of available updates before any action is performed.

== Enable automatic updates

When you enable automatic updates for your AFX cluster, recommended updates for storage firmware, SP/BMC firmware and system files are automatically downloaded and installed by default.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Under **Software updates** select **Enable**.

3. Read the EULA.
4. Accept the defaults to **Show notification** of recommended updates. Optionally, select to **Automatically update** or to **Automatically dismiss** recommended updates.
5. Select to acknowledge that your update modifications will be applied to all current and future updates.
6. Select **Save**.

Result

Recommended updates are automatically downloaded and installed on your ONTAP AFX system based upon your update selections.

== Disable automatic updates

Disable automatic updates if you want the flexibility to view recommended updates before they are installed. If you disable automatic updates, you need to perform firmware and system file updates manually.

Steps

1. In System Manager, select **Cluster > Settings**.
2. Under **Software updates**, select **Disable**.


Result

Automatic updates are disabled. You should regularly check for recommended updates and decide if you want to perform a manual installation.

== View automatic updates

View a list of firmware and system file updates that have been downloaded to your cluster and are scheduled for automatic installation. Also view updates that have been previously automatically installed.


Steps

1. In System Manager, select **Cluster > Settings**.
2. Next to **Software updates** select , then select **View all automatic updates**.

== Edit automatic updates

You can select to have recommended updates for your storage firmware, SP/BMC firmware and your system files automatically downloaded and installed on your cluster, or you can select to have recommended updates automatically dismissed. If you want to manually control installation or dismissal of updates, select to be notified when a recommended update is available; then you can manually select to install or dismiss it.

Steps

1. In System Manager, select **Cluster > Settings**.
2. Next to **Software updates** select  and then select **All other updates**.
3. Update the selections for automatic updates.
4. Select **Save**.

Result

Automatic updates are modified based on your selections.

== Update firmware manually

If you want the flexibility of viewing recommended updates before they are downloaded and installed, you can disable automated updates and update your firmware manually.

Steps

1. Download your firmware update file to a server or local client.
2. In System Manager, select **Cluster > Overview**, then select **All other updates**.
3. Under **Manual Updates**, select **Add firmware files**; then select **Download from the server** or **Upload from the local client**.
4. Install the firmware update file.

Result

Your firmware is updated.

= ONTAP revert unsupported with AFX storage systems

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

Reverting an ONTAP cluster is the process of moving all the nodes to the previous major ONTAP release.

NetApp AFX storage systems do not support ONTAP revert. Attempting a revert operation with AFX can result in cluster instability and data loss. You should not attempt a revert operation on an AFX system.

= Additional administration for an AFX storage system cluster

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

In addition to the typical AFX cluster administration, there may be other tasks you need to perform based on your environment. Most of the additional tasks can be performed using System Manager, although in some cases you may need to use the CLI.



The ONTAP features and administration described are common to AFX storage systems and AFF or FAS systems running Unified ONTAP. Links to the relevant Unified ONTAP documentation are included as appropriate.

== Licensing

AFX systems are licensed in a similar way as Unified ONTAP AFF and FAS systems. An AFX cluster includes most features by default for the protocols supported.

=== ONTAP license management

An ONTAP license is a record of one or more software entitlements. All licenses are defined and provided using a NetApp license file (NLF). Refer to [ONTAP licensing overview](#) for more information.

=== Install a license on an AFX system

You can install license files to activate additional features as needed for your AFX storage system.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. Next to **Licenses**, select [→](#).
3. Select the **Features** tab to display the available ONTAP features.
4. To optionally install a licenses, select the **Installed licenses** tab.
5. Select [+ Add](#).
6. Select a local license file and select **Add**.

== Performance

Use common ONTAP features to optimize the performance of your AFX storage system cluster.

=== Modify default settings that control automatic QoS throughput ceiling increases

Adaptive and non-adaptive QoS policy groups are automatically enabled to temporarily increase a workload's throughput ceiling by 50% for a duration of 1 second to respond to short-term demands for higher performance by default. Beginning with ONTAP 9.19.1, you can modify these default settings. In ONTAP 9.18.1 and earlier, these settings cannot be changed.

Refer to [QoS throughput ceilings](#).

== Security

There are several optional security features you can configure and use with your AFX deployment.

=== ONTAP security and data encryption

It's important to protect the security and privacy of your AFX storage system. Refer to [Security and data encryption](#)

=== ONTAP authentication and access control

The AFX storage system provides several options for configuring authentication and access control services. Refer to [Authentication and access control](#) for more information.

=== ONTAP TLS hardware offload

You can offload Transport Layer Security (TLS) encryption and decryption for your entire cluster to supported Ethernet cards. This reduces CPU overhead and improves performance for systems using TLS. Refer to [Configure ONTAP TLS hardware offload](#) for more information.

=== Administer OAuth 2.0 on an AFX system

OAuth 2.0 is the industry standard authorization framework used to restrict and control access to protected resources using signed access tokens.

Steps

1. In System Manager, select **Cluster** and then **Settings**.
2. In the **Security** section, next to **OAuth 2.0 authorization**, select [→](#).
3. Enable OAuth 2.0
4. Select **Add configuration** and provide the configuration details.
5. Select **Save**.

== Related information

- [FAQ for AFX storage systems](#)
- [How to download NLF licenses from the NetApp Support Site](#)
- [Overview of the ONTAP OAuth 2.0 implementation](#)
- [Additional administration for AFX SVMs](#)

= Administer your storage VMs and data

= Manage data

= Prepare to manage your AFX storage system data

:icons: font

:relative_path: ./manage-data/

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

Before managing your AFX data, you should be familiar with the basic concepts and capabilities.



Because many of the concepts and administration procedures available on AFF and FAS systems are the same with AFX storage systems, reviewing the Unified ONTAP documentation can be helpful. Refer to the links in [Related information](#) for more information.

== Terminology and options

There are several terms related to AFX storage you should be familiar with.

FlexVolume

A FlexVol is a type of logical container used in AFX storage systems. FlexVol volumes can be expanded, contracted, moved, and efficiently copied. They can also be partitioned into more manageable units using qtrees and resource usage can be limited using quotas.

FlexGroup

A FlexGroup volume is a scale-out NAS container that provides both high performance and automatic load distribution. Each consists of multiple volumes that share traffic transparently. FlexGroup volumes offer several benefits, including improved scalability and performance as well as simplified management.

FlexCache

FlexCache is an ONTAP caching technology that creates sparse, writable replicas of volumes on the same or different ONTAP clusters. It is designed to improve data access performance by bringing data closer to users, which can result in faster throughput with a smaller footprint. FlexCache is particularly useful for read-intensive workflows and helps to offload traffic from heavily accessed volumes.

S3 bucket

An S3 bucket is a storage container that holds objects or data in the cloud. With ONTAP, an S3 NAS bucket is a mapping between an S3 bucket name and a NAS path, allowing S3 access to any part of an SVM namespace with existing volumes and directory structure.

Data container

In the context of an AFX system, a data container is a generic term and can be either a volume or S3 bucket.

Qtree

A qtree is a logical subdivision within a volume that you can create to manage and organize data. It allows you to specify its properties and security style (NTFS or UNIX) and can inherit export policies from its parent volume or have its own. Qtrees can contain files and directories, and are often used to manage permissions and quotas more granularly within a volume.

Quota

A quota in ONTAP is a limit set on the amount of storage space or number of files that can be used by a user, group, or qtree. Quotas are used to manage and control resource usage within a storage system, ensuring that no single user or application can consume an excessive amount of resources.

NFS session trunking

NFS trunking is a technology that enables NFS v4.1 clients to open multiple connections to different LIFs on the NFS server. This increases the data transfer speed and provides resilience through multiple paths when exporting volumes to trunking-capable clients. The LIFs must be on the same node to participate in the trunk.

To enable trunking, you need to have an SVM configured for NFS and NFSv4.1 should be enabled. It also requires remounting all the NFSv4.x clients after a configuration change which can be disruptive. Support and configuration procedures for NFS trunking are the same for all ONTAP systems. Learn more about [NFS trunking](#)

File system analytics

File System Analytics (FSA) is an ONTAP feature that provides real-time visibility into file usage and storage capacity trends within FlexGroup or FlexVol volumes. It eliminates the need for external tools by offering insights into storage utilization and optimization opportunities. FSA provides detailed views at various levels of a volume's file system hierarchy, including the SVM, volume, directory, and file levels.

== Data migration options

There are several data migration options. The focus is on migrating external data into an AFX cluster.

=== Migrating data from AFF or FAS systems

A fully integrated migration path from AFF or FAS systems (which run the Unified ONTAP personality) to AFX is available using the following technologies:

- SnapMirror
- SVM Migrate
- SVM DR

In addition, FlexCache volumes can be attached between AFX and AFF or FAS systems in either direction.

=== Migrating data from a non-ONTAP source

Data migration from non-ONTAP systems can be performed using file-level copy operations. Fast copy utilities such as [XCP](#) or [Copy and Sync](#) can be used as well as standard utilities such as RoboCopy (for SMB) and rsync (for NFS) as well as third-party tools such as DataDobi.

=== Migration limitations

You can replicate data from AFF or FAS systems to AFX if the source data volume does not contain LUNs or NVMe namespaces. When replicating from AFX to AFF or FAS systems, the minimum supported ONTAP

version for the AFF or FAS system is 9.16.1. This is the first ONTAP release that supports Advanced Capacity Balancing.

== Display an overview of your storage

To get started managing your AFX data, you should display an overview of the storage.

About this task

You can access all the volumes and buckets defined for the AFX cluster. Each of these is considered to be a data container.

Steps

1. In System Manager, select **Storage** and then **Overview**
2. Next to **Volumes**, select → to display a list of volumes.
3. Next to **Buckets**, select → to display a list of buckets.
4. Update or create a data container as needed.

== Related information

- [Learn about ONTAP File System Analytics](#)
- [Additional AFX SVM administration](#)
- [Prepare to administer your AFX system](#)
- [Migrate an AFX system SVM](#)
- [NetApp Interoperability Matrix Tool](#)

= Create and configure a volume on an AFX storage system

:icons: font

:relative_path: ./manage-data/

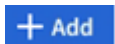
:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

You can create a volume and attach it to an SVM. Each volume can be exposed to clients using one of the access protocols supported by AFX.

About this task

When creating a volume, you need to provide a minimum amount of configuration details. Additional details can be provided during creation or afterwards by editing the volume. You need to select the SVM for the volume if you've created additional SVMs.

Steps

1. In System Manager, select **Storage** and then **Volumes**.
2. Select  and provide the basic configuration including name, capacity, and optimization.
3. Optionally select **More options** for additional configuration related to data protection, SnapLock, and NFS access.
4. Select **Save** to add the volume.

= Manage the AFX storage system volumes

:icons: font

:relative_path: ./manage-data/

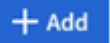
:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

There are several administrative tasks you can perform as part of administering the volumes defined at your AFX cluster.

== Create a qtree

A qtree is a logical subdivision within a volume that you can create to organize and administer data.

Steps

1. In System Manager, select **Storage** and then **Qtrees**.
2. Select  and provide the basic configuration including name, volume, and security style; optionally configure a quota.
3. Select **Save** to add the qtree.

== Create a quota

A quota is a limit set on the amount of storage space or number of files that can be used by a user, group, or qtree. Quotas are used to manage and control resource usage within an AFX system.

Steps

1. In System Manager, select **Storage** and then **Quotas**.
2. Select the **Usage** tab to display a list of the active quotas across all volumes.
3. Select the **Volumes** tab to display a list of the volumes defined in the AFX cluster; select a specific volume to display additional information.
4. To define a quota, select the **Rules** tab.
5. Provide the configuration details, including the quota target, type and limits.
6. Select **Save** to add the quota.

= Create and configure an S3 bucket on an AFX storage system

:icons: font

:relative_path: ./manage-data/

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

You can create a bucket and attach it to an SVM. Each bucket can be exposed to clients using the S3 access protocol supported by AFX.

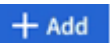
About this task

When creating a bucket, you need to provide a minimum amount of configuration details. Additional details can be provided during creation or afterwards by editing the bucket. You need to select the SVM for the bucket if you've created additional SVMs.

Before you begin

You need to configure the S3 service for the SVM for clients to be able to access the bucket.

Steps

1. In System Manager, select **Storage** and then **Buckets**.
2. Select  and provide the basic configuration including name and capacity.
3. Optionally select **More options** for additional configuration related to data protection, locking, and permissions.

4. Select **Save** to add the bucket.

= Manage the AFX storage system buckets

:icons: font

:relative_path: ./manage-data/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

There are several administrative tasks you can perform as part of managing AFX S3 buckets and client access. S3 configuration and support in AFX is the same as provided with Unified ONTAP. Refer to the Unified ONTAP documentation for details.

It's also important to keep up with the latest changes and updates to ONTAP's support for S3 storage. For example, beginning with ONTAP 9.19.1 you can generate a second access key for each S3 user. This allows you to maintain access to S3 resources if the original access key expires. You can also control conditional operation (write, copy, delete) enforcement.

Related information

- [Learn about ONTAP S3 configuration](#)
- [Regenerate ONTAP S3 keys and modify their retention period](#)
- [Add access rules to the default ONTAP S3 bucket policy](#)

= Monitor and troubleshoot an AFX storage system

:icons: font

:relative_path: ./manage-data/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

The AFX system includes several options to monitor the storage each cluster manages.

== Display NAS clients

You can display a list of the NFS and SMB/CIFS clients currently connected to the AFX cluster.

Steps

1. In System Manager, select **Clients** in the navigation pane.
2. Select the tab **NFS** or **SMB/CIFS** as desired.
3. Customize the display as well as search and download the client information as needed.

== Related information

- [Prepare to manage your AFX data](#)

= Protect data

= Prepare to protect your AFX storage system data

:icons: font

:relative_path: ./protect-data/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

Before protecting your AFX data, you should be familiar with some of the key concepts and capabilities.



Because many of the concepts and administration procedures available on AFF and FAS systems are the same with AFX storage systems, reviewing the Unified ONTAP documentation for [Data protection and disaster recovery](#) can be helpful.

== Terminology and options

There are several terms related to AFX data protection you should be familiar with.

Snapshot

A snapshot is a read-only, point-in-time image of a volume. It is a foundational technology for ONTAP's replication and data protection services.

Consistency group

A consistency group is a collection of volumes that are managed as a single unit. You can create consistency groups to simplify storage management and data protection for application workloads. For example, you can snapshot several volumes in one operation by using the consistency group instead of the individual volumes.

Hierarchical consistency group

Hierarchical consistency groups were introduced with ONTAP 9.16.1 and are available with AFX. With a hierarchical structure, one or more consistency groups can be configured as children under a parent. These hierarchical groups allow you to apply individual snapshot policies to child consistency groups and replicate the snapshots of all the children to a remote cluster as a single unit by replicating the parent.

SnapLock

SnapLock is an ONTAP feature that allows you to protect your files by moving them to a write once read many (WORM) state. This prevents modification or deletion for a specified retention period. SnapLock volumes are created cannot be converted from non-SnapLock volumes after creation based on the retention.

== AFX data protection limitations

You should be aware of the ONTAP data protection limits and restrictions enforced by the AFX storage system.

SnapMirror synchronous (SM-S)

There is a scale limitation when using SM-S. You can have a maximum of 400 relationships across a single AFX system cluster.

== Related information

- [Additional AFX SVM administration](#)
- [Prepare to your administer AFX system](#)

= Create a consistency group on an AFX storage system

:icons: font

:relative_path: ./protect-data/

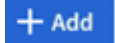
:imagesdir: /tmp/d20260515-4141158-lcjv8c/source/./install-setup/./media/

You can create consistency groups to simplify storage management and data protection for application workloads. A consistency group can be based on existing or new volumes.

Before you begin

If you plan to create one more more new volumes, you should become with the configuration options when creating a ne volume.

Steps

1. In System Manager, select **Protection** and then **Consistency groups**.
2. Select  and choose one of:
 - Using existing volumes
 - Using new NAS volumes
3. Provide the configuration details, including name, volumes, application type, and protection.
4. Select **Add**.

Related information

- [Manage consistency groups](#)
- [Create and configure an AFX volume](#)

= Manage consistency groups on an AFX storage system

:icons: font

:relative_path: ./protect-data/


:imagesdir: /tmp/d20260515-4141158-lcjv8c/source/./install-setup/./media/

You can manage the consistency groups on an AFX system. This can streamline your storage administration.


== Add snapshot data protection to a consistency group

When you add snapshot data protection to a consistency group, local snapshots of the consistency group can be taken at regular intervals based on a pre-defined schedule.

Steps

1. In System Manager, select **Protection** and then **Consistency groups**.
2. Hover over the consistency group you want to protect.
3. Select ; then select **Edit**.
4. Under **Local protection**, select **Schedule snapshots**.
5. Select a snapshot policy.

Accept the default snapshot policy, select an existing policy, or create a new policy.

Option	Steps
Select an existing snapshot policy	Select  next to the default policy; then select the existing policy that you want to use.

Option	Steps
Create a new snapshot policy	<ol style="list-style-type: none"> Select + Add ; then enter the new policy name. Select the policy scope. Under Schedules select + Add . Select the name that appears under Schedule name; then select ▼ . Select the policy schedule. Under Maximum snapshots, enter the maximum number of snapshots that you want to retain of the consistency group. Optionally, under SnapMirror label enter a SnapMirror label. Select Save.

6. Select **Edit**.

== Related information

- [Learn about ONTAP consistency groups](#)

= Create a snapshot on an AFX storage system

:icons: font

:relative_path: ./protect-data/

:imagesdir: /tmp/d20260515-4141158-lciv8c/source/./install-setup/./media/

To back up data on your AFX system, you need to create a snapshot. You can create a snapshot manually or schedule to be created automatically using a consistency group.

== Before you begin

A snapshot is a local, read-only copy of your data that you can use to restore volumes to specific points in time. Snapshots can be created manually on demand or automatically at regular intervals based on a [snapshot policy and schedule](#).

The snapshot policy and schedule specifies the details, including when to create the snapshots, how many copies to retain, how to name them, and how to label them for replication. For example, a system might create one snapshot every day at 12:10 a.m., retain the two most recent copies, name them “daily” (appended with a timestamp), and label them “daily” for replication.

Types of snapshots

You can create an on-demand snapshot of a single volume or a consistency group. You can also create automated snapshots of a consistency group containing multiple volumes. However you cannot create automated snapshots of a single volume.

- On-demand snapshots

You can create an on-demand snapshot of a volume at any time. The volume does not need to be a member of a consistency group to be protected by an on-demand snapshot. If you create a snapshot of a volume that is a member of a consistency group, the other volumes in the consistency group are not included in the snapshot. When you create an on-demand snapshot of a consistency group, all the

volumes in the consistency group are included.


- Automated snapshots

Automated snapshots are created based on the snapshot policy definitions. To apply a snapshot policy to a volume for automated snapshot creation, the volumes needs to be a member of the same consistency group. If you apply a snapshot policy to a consistency group, all the volumes in the consistency group are protected.

== Create a snapshot

Create a snapshot of a volume or consistency group.

Steps

1. In System Manager, select **Protection** and then **Consistency groups**.
2. Hover over the name of the consistency group you want to protect.
3. Select  ; then select **Protect**.
4. If you want to create an immediate snapshot on-demand, under **Local protection**, select **Add a snapshot now**.



Local protection creates the snapshot on the same cluster containing the volume.

- a. Enter a name for the snapshot or accept the default name; then optionally, enter a SnapMirror label.

The SnapMirror label is used by the remote destination.

5. If you want to create automated snapshots using a snapshot policy, select **Schedule snapshots**.
 - a. Select a snapshot policy.

Accept the default snapshot policy, select an existing policy, or create a new policy.


Option	Steps
Select an existing snapshot policy	Select  next to the default policy; then select the existing policy that you want to use.
Create a new snapshot policy	<ol style="list-style-type: none">a. Select  Add ; then enter the snapshot policy parameters.b. Select Add policy.

6. If you want to replicate your snapshots to a remote cluster, under **Remote protection**, select **Replicate to a remote cluster**.
 - a. Select the source cluster and storage VM; then select the replication policy.

The initial data transfer for replication starts immediately by default.

7. Select **Save**.

Steps

1. In System Manager, select **Storage** and then **Volumes**.
2. Hover over the name of the volume you want to protect.
3. Select  ; then select **Protect**.
If you want to create an immediate snapshot on-demand, under **Local protection**, select **Add a snapshot now**.



Local protection creates the snapshot on the same cluster containing the volume.

4. Enter a name for the snapshot or accept the default name; then optionally, enter a SnapMirror label.

The SnapMirror label is used by the remote destination.

5. If you want to create automated snapshots using a snapshot policy, select **Schedule snapshots**.
 - a. Select a snapshot policy.

Accept the default snapshot policy, select an existing policy, or create a new policy.

Option	Steps
Select an existing snapshot policy	Select  next to the default policy; then select the existing policy that you want to use.
Create a new snapshot policy	<ol style="list-style-type: none"> a. Select  Add ; then enter the snapshot policy parameters. b. Select Add policy.

6. If you want to replicate your snapshots to a remote cluster, under **Remote protection**, select **Replicate to a remote cluster**.
 - a. Select the source cluster and storage VM; then select the replication policy.

The initial data transfer for replication starts immediately by default.

7. Select **Save**.

== Related information

- [Create an ONTAP snapshot policy](#)

= Manage snapshots on an AFX storage system

:icons: font

:relative_path: ./protect-data/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

You can manage snapshots on your AFX system. Refer to the Unified ONTAP documentation for details.

Related information

- [Create an ONTAP snapshot policy](#)
- [Protect ONTAP FlexGroup volumes using snapshots](#)

= Create an intercluster SVM peer relationship on an AFX storage system

:icons: font

:relative_path: ./protect-data/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

A peer relationship defines network connections that enable clusters and storage virtual machine (VM) to exchange data securely. You can create a peer relationship between storage VMs on different clusters to enable data protection and disaster recovery using SnapMirror.

Before you begin

You must have established a cluster peer relationship between the local and remote clusters before you can create a storage VM peer relationship. [Create a cluster peer relationship](#) if you have not already done so.

Steps

1. In System Manager, select **Protection > Overview**.
2. Under **Storage VM peers** select **Add a storage VM peer**.
3. Select the storage VM on the local cluster; then select the storage VM on the remote cluster.
4. Select **Add a storage VM peer**.

Related information

- [Learn more about peer relationships](#).

= Manage snapshot replication on an AFX storage system

:icons: font

:relative_path: ./protect-data/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

Snapshot replication is a process in which consistency groups on your AFX system are copied to a geographically remote location. After the initial replication, changes to consistency groups are copied to the remote location based upon a replication policy. Replicated consistency groups can be used for disaster recovery or data migration.

To set up Snapshot replication, you need to establish a replication relationship between your AFX storage system and the remote location. The replication relationship is governed by a replication policy. A default policy to replicate all snapshots is created during cluster set up. You can use the default policy or optionally, create a new policy.



== Step 1: Create a cluster peer relationship

Before you can protect your data by replicating it to a remote cluster, you need to create a cluster peer relationship between the local and remote cluster.

Before you begin

The prerequisites for cluster peering are the same for AFX systems as for other ONTAP systems. [Review the prerequisites for cluster peering.](#)

Steps

1. On the local cluster, in System Manager, select **Cluster > Settings**.
2. Under **Intercluster Settings** next to **Cluster peers** select , then select **Add a cluster peer**.
3. Select **Launch remote cluster**; this generates a passphrase you'll use to authenticate with the remote cluster.
4. After the passphrase for the remote cluster is generated, paste it under **Passphrase** on the local cluster.
5. Select  **Add**; then enter the intercluster network interface IP address.
6. Select **Initiate cluster peering**.


What's next?

You have peered for local AFX cluster with a remote cluster. You can now create a replication relationship.

== Step 2: Optionally, create a replication policy

The snapshot replication policy defines when updates performed on the AFX cluster are replicated to the remote site.

Steps

1. In System Manager, select **Protection > Policies**; then select **Replication policies**.
2. Select  **Add** .
3. Enter a name for the replication policy or accept the default name; then enter a description.
4. Select the **Policy scope**.

If you want to apply the replication policy to the entire cluster, select **Cluster**. If you want the replication policy applied only to the volume in a specific storage VM, select **Storage VM**.

5. Select the **Policy type**.

Option	Steps
Copy data to the remote site after it is written to the source.	<ol style="list-style-type: none"> a. Select Asynchronous. b. Under Transfer snapshots from source, accept the default transfer schedule or select a different one. c. Select to transfer all snapshots or to create rules to determine which snapshots to transfer. d. Optionally, enable network compression.
Write data to the source and remote sites simultaneously.	<ol style="list-style-type: none"> a. Select Synchronous.

6. Select **Save**.

What's next?

You have created a replication policy and are now ready to create a replication relationship between your AFX system and your remote location.

== Step 3: Create a replication relationship


A snapshot replication relationship establishes a connection between your AFX system and a remote location so that you can replicate consistency groups to a remote cluster. Replicated consistency groups can be used for disaster recovery or for data migration.

For protection against ransomware attacks, when you set up your replication relationship, you can select to lock the destination snapshots. Locked snapshots cannot be deleted accidentally or maliciously. You can use locked snapshots to recover data if a volume is compromised by a ransomware attack.

Before you begin

Create a replication relationship with or without locked destination snapshots.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Select a consistency group.
3. Select ; then select **Protect**.
4. Under **Remote protection**, select **Replicate to a remote cluster**.
5. Select the **Replication policy**.

You must select a *vault* replication policy.

6. Select **Destination settings**.
7. Select **Lock destination snapshots to prevent deletion**
8. Enter the maximum and minimum data retention period.
9. To delay the start of the data transfer, deselect **Start transfer immediately**.

The initial data transfer begins immediately by default.



- Optionally, to override the default transfer schedule, select **Destination settings**, then select **Override transfer schedule**.

Your transfer schedule must be a minimum of 30 minutes to be supported.

- Select **Save**.

Steps

- In System Manager, select **Protection > Replication**.
- Select to create the replication relationship with local destination or local source.

Option	Steps
Local destinations	<ol style="list-style-type: none"> Select Local destinations, then select . Search for and select the source consistency group. The <i>source</i> consistency group refers to the consistency group on your local cluster that you want to replicate.
Local sources	<ol style="list-style-type: none"> Select Local sources, then select . Search for and select the source consistency group. The <i>source</i> consistency group refers to the consistency group on your local cluster that you want to replicate. Under Replication destination, select the cluster to replicate to; then select the storage VM.

- Select a replication policy.
- To delay the start of the data transfer, select **Destination settings**; then deselect **Start transfer immediately**.

The initial data transfer begins immediately by default.

- Optionally, to override the default transfer schedule, select **Destination settings**, then select **Override transfer schedule**.

Your transfer schedule must be a minimum of 30 minutes to be supported.

- Select **Save**.


What's next?

Now that you have created a replication policy and relationship, your initial data transfer begins as defined in your replication policy. You can optionally test your replication failover to verify that successful failover can occur if your AFX system goes offline.

== Step 4: Test replication failover

Optionally, validate that you can successfully serve data from replicated volumes on a remote cluster if the source cluster is offline.

Steps

1. In System Manager, select **Protection > Replication**.
2. Hover over the replication relationship you want to test, then select .
3. Select **Test failover**.
4. Enter the failover information, then select **Test failover**.

What's next?

Now that your data is protected with snapshot replication for disaster recovery, you should [encrypt your data at rest](#) so that it can't be read if a disk in your AFX system is repurposed, returned, misplaced or stolen.

= Manage AFX storage system data protection policies and schedules

:icons: font

:relative_path: ./protect-data/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

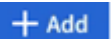
You can use snapshot policies to protect data in your consistency groups based on an automated schedule. The policy schedules within snapshot policies determine how often snapshots are taken.

== Create a new protection policy schedule

A protection policy schedule defines how often a snapshots policy is executed. You can create schedules to run in regular intervals based on a number of days, hours, or minutes. For example, you can create a schedule to run every hour or to run only once per day. You can also create schedules to run at specific times on specific days of the week or month. For example, you can create a schedule to run at 12:15am on the 20th of every month.

Defining various protection policy schedules gives you the flexibility to increase or decrease the frequency of snapshots for different applications. This enables you to provide a greater level of protection and a lower risk of data loss for your critical workloads than what might be needed for less critical workloads.

Steps

1. Select **Protection** and then **Policies**; then select **Schedule**.
2. Select .
3. Enter a name for the schedule; then select the schedule parameters.
4. Select **Save**.

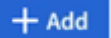
What's next?

Now that you have created a new policy schedule, you can use the newly created schedule within your policies to define when snapshots are taken.

== Create a snapshot policy

A snapshot policy defines how often snapshots are taken, the maximum number of snapshots allowed, and how long snapshots are retained.

Steps

1. In System Manager, select **Protection** and then **Policies**; then select **Snapshot policies**.
2. Select  .
3. Enter a name for the snapshot policy.
4. Select **Cluster** to apply the policy to the entire cluster. Select **Storage VM** to apply the policy to an individual storage VM.
5. Select **Add a schedule**; then enter the snapshot policy schedule.
6. Select **Add policy**.


What's next?

Now that you have created a snapshot policy, you can apply it to a consistency group. Snapshots will be taken of the consistency group based on the parameters you set in your snapshot policy.

== Apply a snapshot policy to a consistency group

Apply a snapshot policy to a consistency group to automatically create, retain, and label snapshots of the consistency group.

Steps

1. In System Manager, select **Protection** and then **Policies**; then select **Snapshot policies**.
2. Hover over the name of the snapshot policy you want to apply.
3. Select ; then select **Apply**.
4. Select the consistency groups to which you want to apply the snapshot policy; then select **Apply**.


What's next?

Now that your data is protected with snapshots, you should [set up a replication relationship](#) to copy your consistency groups to a geographically remote location for backup and disaster recovery.

== Edit, delete, or disable a snapshot policy

Edit a snapshot policy to modify the policy name, maximum number of snapshots, or the SnapMirror label. Delete a policy to remove it and its associated back up data from your cluster. Disable a policy to temporarily stop the creation or transfer of snapshots specified by the policy.

Steps

1. In System Manager, select **Protection** and then **Policies**; then select **Snapshot policies**.
2. Hover over the name of the snapshot policy you want to edit.
3. Select ; then select **Edit**, **Delete**, or **Disable**.


Result

You have modified, deleted or disabled the snapshot policy.

== Edit a replication policy

Edit a replication policy to modify the policy description, transfer schedule, and rules. You can also edit the policy to enable or disable network compression.

Steps

1. In System Manager, select **Protection** and then **Policies**.
2. Select **Replication policies**.
3. Hover over the replication policy that you want to edit; then select .
4. Select **Edit**.
5. Update the policy; then select **Save**.

= Secure data

= Prepare to secure your AFX storage system data

:icons: font

:relative_path: ./secure-data/

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

Before managing your AFX data, you should be familiar with the major concepts and capabilities.

== Terminology and options

There are several terms related to AFX data security you should be familiar with.

Ransomware

Ransomware is malicious software that encrypts files making them inaccessible to the user. There is typically some type of payment demanded to decrypt the data. ONTAP provides solutions to protect against ransomware through features like Autonomous Ransomware Protection (ARP).

Encryption

Encryption is the process of converting data into a secure format that cannot be easily read without proper authorization. ONTAP offers both software-based and hardware-based encryption technologies to protect data at rest. This ensures it cannot be read if the storage medium is repurposed, returned, misplaced, or stolen. These encryption solutions can be managed using either an external key management server or the Onboard Key Manager provided by ONTAP. Refer to [Encrypt data at rest on an AFX storage system](#) for more information.

Digital certificates and PKI

A digital certificate is an electronic document used to prove ownership of a public key. The public key and associated private key can be used in various ways, including to establish identity typically as part of a larger security framework such as TLS and IPsec. These keys, as well as the supporting protocols and formatting standards, form the basis for public key infrastructure (PKI). Refer to [Manage certificates on an AFX storage system](#) for more information.

Internet Protocol Security

IPsec is an Internet standard that provides in-flight data encryption, integrity, and authentication for traffic flowing among network endpoints at the IP level. It secures all IP traffic between ONTAP and clients including higher level protocols such as NFS and SMB. IPsec provides protection against malicious replay and man-in-the-middle attacks on your data. Refer to [Secure IP connections on your AFX storage systems](#)

for more information.

== Related information

- [Additional AFX SVM administration](#)
- [Prepare to administer your AFX system](#)

= Encrypt data at rest on an AFX storage system

:icons: font

:relative_path: ./secure-data/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

You can encrypt your data at the hardware and software level for dual-layer protection. When you encrypt data at rest, it can't be read if the storage medium is repurposed, returned, misplaced, or stolen.

NetApp Storage Encryption (NSE) supports hardware encryption using self-encrypting drives (SEDs). SEDs encrypt data as it is written. Each SED contains a unique encryption key. Encrypted data stored on the SED can't be read without the SED's encryption key. Nodes attempting to read from an SED must be authenticated to access the SED's encryption key. Nodes are authenticated by obtaining an authentication key from a key manager, then presenting the authentication key to the SED. If the authentication key is valid, the SED will give the node its encryption key to access the data it contains.

Before you begin

Use the AFX onboard key manager or an external key manager to serve authentication keys to your nodes. In addition to NSE, you can also enable software encryption to add another layer of security to your data.

Steps

1. In System manager, select **Cluster** and then **Settings**.
2. In the **Security** section, under **Encryption**, select **Configure**.
3. Configure the key manager.

Option	Steps
Configure the Onboard key Manager	<ol style="list-style-type: none">a. Select Onboard Key Manager to add the key servers.b. Enter a passphrase.
Configure an external key manager	<ol style="list-style-type: none">a. Select External key manager to add the key servers.b. Select + Add to add the key servers.c. Add the KMIP server CA certificates.d. Add the KMIP client certificates.

4. Select **Dual-layer encryption** to enable software encryption.
5. Select **Save**.

Related information

- [Encryption](#)

= Secure IP connections on your AFX storage systems

:icons: font

:relative_path: ./secure-data/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

IP Security (IPsec) is an Internet protocol standard that provides data encryption, integrity, and authentication for traffic flowing among network endpoints at the IP level. You can use IPsec to enhance the security of the front-end network between an AFX cluster and the clients.

== Configuring IPsec on an AFX system

The IPsec configuration procedures for AFX storage systems are the same as AFF and FAS systems, with the exception of the supported network interface controller (NIC) cards used with the hardware offload feature. Refer to [Prepare to configure IP security for the ONTAP network](#) for more information.

== Hardware offload feature

Several of the IPsec cryptographic operations, such as encryption and integrity checks, can be offloaded to a supported NIC card on your AFX system. This can significantly improve the performance and throughput of the network traffic protected by IPsec.



Beginning with ONTAP 9.18.1, the IPsec hardware offload feature is extended to support IPv6 traffic.

The following NIC cards are supported for the IPsec hardware offload feature on AFX storage systems beginning with ONTAP 9.17.1:

- X50130B (2p, 40G/100G Ethernet controller)
- X50131B (2p, 40G/100G/200G/400G Ethernet controller)

Refer to the [NetApp Hardware Universe](#) for more information about the supported cards for the ONTAP release running on your AFX system.

== Related information

- [Prepare to configure IP security for the ONTAP network](#)
- [NetApp Hardware Universe](#)

= Additional administration for an AFX storage system SVM

:icons: font

:relative_path: ./administer/

:imagesdir: /tmp/d20260515-4141158-lcjbv8c/source/./install-setup/./media/

In addition to the typical AFX SVM administration, there may be other tasks you need to perform based on your environment. Most of the additional tasks can be performed using System Manager, although in some cases you may need to use the CLI.



The ONTAP features and administration described are common to AFX storage systems and AFF or FAS systems running Unified ONTAP. Links to the relevant Unified ONTAP documentation are included as appropriate.

== Storage management and performance

There are several optional storage management and performance features you can configure and use with your AFX deployment.

=== ONTAP NAS storage management

Network Attached Storage (NAS) provides dedicated file storage which can be accessed by multiple clients in the network. ONTAP supports several NAS protocols. Learn more about [NAS storage management](#).

=== ONTAP FlexCache volumes

FlexCache is an ONTAP remote caching feature. It brings data closer to clients which improves access performance and reduces costs. Creating a FlexCache volume, which initially copies only the metadata from the origin file system, simplifies file distribution and reduces WAN traffic. Learn more about [ONTAP FlexCache volumes](#).

=== ONTAP FlexGroup volumes

A FlexGroup volume consists of several member volumes that share the traffic automatically and transparently. FlexGroup volumes offer several benefits, including high performance and simplified management. Learn more about [FlexGroup volume setup](#).

=== Cross File Sequential Read

Cross File Sequential Read (CFSR) is a feature available on AFX systems beginning with ONTAP 9.19.1 that improves the performance of video editing and related applications. For workflows that store video frames as individual sequential files, typically 50 to 250 MB each, CFSR intelligently predicts and prefetches the next files based on sequential filename patterns. Up to 10 files or 500 MB of data can be prefetched, significantly reducing read latency and ensuring smooth playback of 4K/8K raw video at 30 frames per second.

CFSR is enabled by default on AFX systems running ONTAP 9.19.1 or later, but you need to explicitly enable it on the individual volumes. Storage administrators can configure the feature with the ONTAP CLI or equivalent REST API call. Using the CLI, you can enable CFSR as part of creating a new volume. You can also update an existing volume with the following command:

```
volume modify -volume <volume_name> -aggressive-readahead-mode  
[cross_file_sequential_read|none]
```

Example

```
volume modify -volume myvolume -aggressive-readahead-mode  
cross_file_sequential_read  
volume show -volume myvolume -fields aggressive-readahead-mode
```

The CFSR feature is compatible with both FlexVol and FlexGroup volumes and is file-extension agnostic, supporting decimal and hexadecimal filename numbering schemes. For NFSv4.x protocols, CFSR

integrates with Multi-Protocol I/O (MPIO) read capabilities to further enhance performance by serving subsequent reads directly from memory. CFSR is available exclusively on AFX systems, with limitations on the volumes supported. It's not supported with the following volume types:

- FabricPool-enabled volumes
- FlexCache endpoint volumes
- SnapMirror destination volumes

CFSR is not supported on any other ONTAP platforms (FAS, AFF, or ASA) or in cloud deployments, such as Cloud Volumes ONTAP. It can operate using FlexGroup volumes with GDD enabled or disabled. However, because video frame files are typically much smaller than GDD's minimum segment size, GDD does not stripe these files effectively across FlexGroup constituents. This means GDD has minimal direct impact on CFSR's file prefetching behavior.

== Data protection

There are several optional data protection features you can configure and use with your AFX deployment.

=== Consistency groups

A consistency group is a collection of volumes that are managed as a single unit. Learn more about [ONTAP consistency groups](#).

=== SnapLock

You can protect your files by converting them to a write once read many (WORM) state at a volume level. SnapLock supports two modes. Compliance mode ensures the files cannot be deleted until their retention period expires which addresses government or industry-specific mandates. Enterprise mode allows privileged users to delete files before their retention period expires. Learn more about [ONTAP SnapLock](#).

== Storage efficiency

ONTAP provides several storage efficiency features that can help you optimize your storage usage and reduce costs.

=== Global deduplication

[Deduplication](#) is an [ONTAP storage efficiency](#) feature that reduces physical storage requirements by identifying and removing duplicate data blocks. Traditional deduplication operates at either a volume or aggregate level. This existing capability provides significant storage efficiency for customers across all ONTAP platforms.

Beginning with ONTAP 9.19.1, AFX introduces a new operational scope by extending deduplication to the entire Storage Availability Zone (SAZ). Global deduplication builds on volume-level deduplication by automatically eliminating duplicate blocks across the volumes throughout the SAZ. Global deduplication sessions are triggered automatically, which ensures optimal savings without manual intervention.

Global deduplication is enabled by default on AFX A1K systems and operates transparently in the background with minimal performance impact. No feature license is required and the feature integrates seamlessly with existing storage efficiency technologies including compression and compaction. The savings achieved through global deduplication are reported through the storage efficiency metrics. This feature can be particularly beneficial for application workloads where data is frequently shared or replicated across multiple volumes.

== Security

There are several optional security features you can configure and use with your AFX deployment.

=== FPolicy

FPolicy is a file access notification framework used to monitor and manage file access events on storage virtual machines (SVMs). You can use FPolicy to create policies that define which file operations to monitor, and optionally block, based on criteria you define. FPolicy is commonly used for security auditing, compliance, and data governance. Learn more about [ONTAP FPolicy solutions](#).

== ONTAP event and performance monitoring

You can monitor the health and performance of a cluster. This includes setting up alerts for events and managing notifications for system health alerts. Learn more about [Event, performance, and health monitoring](#).

== Related information

- [FAQ for AFX storage systems](#)
- [Additional administration for AFX clusters](#)

= Maintain the AFX storage system hardware

:icons: font

:relative_path: ./

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

Navigate to the [AFX maintenance documentation](#) to learn how to perform maintenance procedures on your AFX storage system.

= Use the REST API

= Learn about the AFX storage system REST API

:hardbreaks:

:icons: font

:linkattrs:

:relative_path: ./rest/

:imagesdir: /tmp/d20260515-4141158-lcgv8c/source/./install-setup/./media/

The REST API provided with AFX is based on the Unified ONTAP REST API. There are a number of changes which adapt it to the unique characteristics and capabilities of the AFX personality.

== Unsupported features

AFX is a high-performance NAS and S3 storage system. It enables clients to access data using NFS, SMB/CIFS, and S3. Because of this specialization, there are several unsupported features including:

- Metrocluster
- Storage area networking (SAN)
- Disk aggregates

== Removed API endpoints

Several endpoints have been removed from the REST API corresponding to the unsupported features.

```
/cluster/counter/tables
/cluster/metrocluster
/cluster/metrocluster/diagnostics
/cluster/metrocluster/dr-groups
/cluster/metrocluster/interconnects
/cluster/metrocluster/nodes
/cluster/metrocluster/operations
/cluster/metrocluster/svms
/network/fc/fabrics
/network/fc/interfaces
/network/fc/logins
/network/fc/ports
/network/fc/wwpn-aliases
/protocols/nvme/interfaces
/protocols/nvme/services
/protocols/nvme/subsystem-controllers
/protocols/nvme/subsystem-maps
/protocols/nvme/subsystems
/protocols/san/fcp/services
/protocols/san/igroups
/protocols/san/initiators
/protocols/san/iscsi/credentials
/protocols/san/iscsi/services
/protocols/san/iscsi/sessions
/protocols/san/lun-maps
/protocols/san/portsets
/protocols/san/vvol-bindings
/storage/luns
/storage/namespaces
```

== Related information

- [ONTAP automation](#)
- [REST API support for ASA r2](#)

= Your first AFX storage system REST API call

:hardbreaks:

:icons: font

:linkattrs:

:relative_path: ./rest/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

You can issue a simple curl command to get started using the AFX REST API and confirm its availability.

About this task

AFX is one of three ONTAP personalities available from NetApp. You can issue a REST API call to determine the personality of your ONTAP cluster. You can also use System Manager or the CLI to determine the ONTAP personality; see the FAQ page for details.

Before you begin

In addition to having the curl utility available on your local workstation, you'll need the following:

- IP address or FQDN of the AFX system cluster management LIF
- ONTAP credentials for an account with authority to access the ONTAP REST API

Steps

1. Issue the following command at the CLI of your local workstation:

```
curl --request GET \  
"https://$FQDN_IP/api/cluster?fields=disaggregated,san_optimized" \  
--user username:password
```

2. Based on the response, determine the ONTAP personality as follows:

- If "disaggregated" is **true** and:
 - If "san_optimized" is **false** the personality is AFX
 - If "san_optimized" is **true** the personality is ASA r2
- If "disaggregated" is **false** the personality is Unified ONTAP

Related information

- [FAQ for AFX storage systems](#)

= REST API reference for the AFX storage system

:hardbreaks:

:icons: font

:linkattrs:

:relative_path: ./rest/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

The AFX REST API reference contains details about all the API calls. This documentation is helpful when developing automation applications.

Before you begin

You'll need the following:

- IP address or FQDN of the AFX cluster management LIF
- Credentials for a cluster administrator account

Steps

1. Connect your web browser to the cluster management IP address or domain name:

```
https://$FQDN_IP_PORT/docs/api
```

Example

```
https://10.61.25.33/docs/api
```

2. Provide the username and password if prompted.
3. Scroll down to the **Cluster** category and select **GET** next to the endpoint `/cluster` for an example of an individual API call.

Related information

- [ONTAP REST API reference](#)

= Learn more

= Additional resources for AFX storage systems

:icons: font

:relative_path: ./learn-more/

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

There are additional resources you can access to help administer and support AFX as well as learn more about ONTAP and the related NetApp products and services.

== ONTAP documentation

- [Unified ONTAP](#)
- [ASA r2](#)
- [ONTAP automation](#)

== NetApp Support

- [NetApp Support Site](#)
- [NetApp Hardware Universe](#)
- [NetApp Interoperability Matrix Tool](#)
- [NetApp Knowledge Base](#)

= FAQ for AFX storage systems

:hardbreaks:

:icons: font

:linkattrs:

:relative_path: ./

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

This FAQ list provides answers to questions you might have about your AFX storage system. It includes concepts and terminology that can be helpful when exploring AFX in more detail or performing advanced administration tasks.

== General

What is an ONTAP personality?

ONTAP is a robust and versatile storage platform, known for its comprehensive feature set and adaptability to a wide range of storage requirements. While this flexibility makes it an excellent choice for organizations with diverse workloads, some customers can benefit from a more tailored storage solution optimized for the needs of their specific environment.

To address these specialized needs, some NetApp storage systems offer distinct ONTAP personalities, each of which includes a feature set designed to support the unique customer requirements. An ONTAP personality is typically a combination of hardware and software capabilities and purpose-built to deliver an optimized experience for targeted use cases. NetApp provides three ONTAP personalities:

- **Unified ONTAP** - The Unified ONTAP personality delivers a broad set of data management features, supporting NAS, SAN, and S3 protocols for maximum flexibility. This is the NetApp flagship offering, available on AFF and FAS systems as well as virtualized deployments such as ONTAP Select and Cloud Volumes ONTAP.
- **AFX** - The AFX ONTAP personality provides a disaggregated solution engineered to meet the rigorous requirements of high-performance NAS and S3 workloads, including AI/ML applications. AFX systems deliver specialized capabilities for customers requiring scalable, high-throughput file and object storage.
- **ASA r2** - The ASA r2 ONTAP personality provides a disaggregated solution designed specifically for SAN-only environments. [ASA r2 systems](#) streamline the storage experience for block workloads, providing simplified management and performance optimized for SAN customers.

By offering these distinct ONTAP personalities, NetApp enables organizations to select a storage solution aligned with their operational requirements and application workloads.

Can I change the ONTAP personality of my NetApp storage system?

No. The personality of your ONTAP storage system is immutable and cannot be changed. For example, you cannot convert or upgrade a FAS or AFF storage system (which run the Unified ONTAP personality) to an AFX storage system.

The System Manager interfaces for the different ONTAP personalities all look very similar. How can I determine the personality of a specific system?

In System Manager, select **Cluster** in the left navigation pane and then **Overview**. You'll see the personality displayed on the page. As an alternative, you can issue the command "system node show" at the CLI. You can also determine the personality of an ONTAP cluster using the REST API; see [Your first AFX system REST API call](#) for details.

When did the AFX storage systems become available? What's the earliest ONTAP version supported with AFX?

The AFX storage system was announced at the NetApp Insight conference in October, 2025. AFX supports ONTAP 9.17.1 and later releases. Contact your NetApp sales representative for more details.

What does "disaggregated" mean in the context of AFX storage systems?

The term "disaggregated" can have two different though related meanings with AFX depending on the context.

An important concept to begin with is the decoupling of the compute capability in the controller nodes from the storage shelves. With AFX, the cluster compute and storage components are no longer tightly coupled as they are with FAS and AFF systems, which run the Unified ONTAP personality. Instead they are connected through cluster switches. Each AFX node controller has a complete view of the entire cluster storage pool.

The second related concept with disaggregated storage is that the aggregates and RAID management are removed as administrable entities. A storage abstraction layer within AFX automatically configures and manages the low-level aspects of storage, including the physical disks and RAID groups. This allows the AFX administrator to focus on the high-level storage configuration based on volumes and buckets.

== Interoperability

Can I mix AFX system nodes with AFF, ASA, or FAS system nodes in the same ONTAP cluster?

No. You cannot mix system nodes that run different ONTAP personalities in the same cluster. For example, you cannot mix AFX nodes (which run the AFX ONTAP personality) with AFF or FAS nodes (which run the Unified ONTAP personality) in the same cluster.

Can I use FlexCache with an AFX system cluster?

Yes. AFX storage systems support FlexCache both to and from an AFF or FAS system running the Unified ONTAP personality. The only restriction is that FlexCache with write-back mode is not supported with AFX.

If I want to use an AFF or FAS system (running the Unified ONTAP personality) with an AFX system for SnapMirror or FlexCache, what ONTAP version do I need?

The SnapMirror version rules for AFX are identical to Unified ONTAP. This means that to replicate from Unified ONTAP, the source system must be within the supported version range. To replicate from AFX, the Unified ONTAP system must be at ONTAP 9.16.1 or later (the earliest release the Advanced Capacity Balancing feature is supported). For FlexCache, the same rules apply for source and destination systems as described in [TR-4743](#).

There are some differences in the support for FlexGroup volumes. A FlexGroup volume on AFX cannot be an origin volume for a FlexCache volume using the Writeback mode.

Can I make ONTAPI (ZAPI) API calls to an AFX cluster?

No. Only the ONTAP REST API is supported with AFX. Any automation code that uses ZAPI needs to be converted to the REST API for use with AFX.

== Advanced concepts

What are the data protocols supported with an AFX storage system?

The data protocols supported with AFX include the following:

- NFSv3, NFSv4.0, NFSv4.1, NFSv4.2
- SMB2.x, SMB3.x

- S3
- NDMP

Do the data protocols operate differently in AFX?

No. The data protocols in AFX operate the same way as with AFF and FAS systems.

Is Advanced Disk Partitioning (ADP) used in AFX?

No. ADP is not used with AFX. Because there are no root aggregates with AFX, the ADP feature is not needed to maximize disk space efficiency.

Can I use any type of switches in the backend cluster network for my AFX storage system?

No. Only switches specifically approved for and provided with the AFX storage platform are supported for the cluster network. Also, these backend switches are dedicated for AFX cluster operations. The client access operations (using NFS, SMB, and S3) should only be performed over the frontend client data network.

How are the cluster switches configured?

The cluster network switches are configured using a NetApp-provided configuration file. Changes to the configuration file are not supported.

How is the storage in an AFX cluster organized?

All the disks and storage shelves attached to an AFX cluster are part of a Storage Availability Zone (SAZ). Each AFX cluster supports only one SAZ which cannot be shared across AFX clusters (except for SnapMirror replication and FlexCache operations).

Every node has visibility to all the storage in the SAZ. When storage shelves are added to a cluster, ONTAP automatically adds the disks.

How is data allocated in an AFX cluster?

When data is allocated, it can be placed on any disk in the SAZ. Once the data is placed, there is no need for it to move. A volume is created based on the underlying data and assigned to a specific node. The volumes can be moved by AFX, typically as part of a balancing process. This affects which node's NVRAM commits the write operations to disk. A volume move changes which node owns the volume but the data itself can remain in place.

How does AFX manage the volumes across the SAZ?

AFX includes a feature known as Automated Topology Management (ATM) which responds to system and user object imbalances. The primary objective of ATM is to balance volumes across the AFX cluster. When an imbalance is detected, an internal job is triggered to evenly distribute the data across the active nodes. The data is reallocated using ZCVM which only needs to copy and update the object metadata.

How are volumes placed at the nodes in an AFX cluster?

NetApp AFX automatically balances the placement of volumes across all the nodes in a cluster. Beginning with ONTAP 9.18.1, the placement algorithm has been enhanced to consider the performance of the nodes when placing or moving volumes. This results in improved performance balancing across the nodes in the AFX cluster and makes it much less likely that any single node becomes overloaded. Previous AFX releases base the placement on the total number of volumes in the cluster. Every node is assigned the same number of volumes regardless of activity.

How do volume move operations work differently with AFX compared to AFF or FAS systems?

With AFF and FAS systems, which run the Unified ONTAP personality, it's possible to relocate a volume nondisruptively from one node or aggregate to another in the cluster. This is performed using a

background copy operation with SnapMirror technology. A new destination volume is created at the target destination. Depending on the size of the volume and the utilization of cluster resources, the time it takes for a volume move operation to complete can vary.

With AFX, there are no aggregates. All storage is contained within a single Storage Availability Zone (SAZ) which is accessible by every node in the cluster. As a result, volume moves never need to actually copy the data. Instead, all volume moves are performed with pointer updates among the nodes. This is referred to as a Zero Copy Volume Move (ZCVM) and happens instantaneously because no data is actually copied or moved. This is essentially the same volume move process used with Unified ONTAP without the SnapMirror copy.

In the initial 9.17.1 AFX release, volumes will move only in storage failover recovery scenarios and when nodes are added or removed from the cluster. These moves are controlled only through ONTAP.

= Legal notices for AFX storage systems

:icons: font

:relative_path: ./

:imagesdir: /tmp/d20260515-4141158-lcjb8c/source/./install-setup/./media/

Legal notices provide access to copyright statements, trademarks, patents, and more.

== Copyright

<https://www.netapp.com/company/legal/copyright/>

== Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

== Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

== Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

== Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

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—without 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.