



# Get started

## AFX

NetApp  
February 06, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-afx/get-started/ontap-afx-storage.html> on February 06, 2026. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

Get started .....	1
Learn about your AFX system .....	1
Learn about AFX storage systems .....	1
Details of the AFX storage system architecture .....	3
Compare AFX storage system to AFF and FAS systems .....	5
Quick start for setting up an AFX storage system .....	7
Install your AFX system .....	7
Installation and setup workflow for AFX 1K storage systems .....	7
Installation requirements for AFX 1K storage systems .....	8

# 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 as well as the other Unified ONTAP 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 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

### Data replication (SnapMirror)



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

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

### Manageability

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

### Changes to the command line interface

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

- New AFX commands related to:

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

## Related information

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

# 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 your cluster and SVMs.

1

### Install and set up your hardware

[Install](#) your AFX storage system and prepare to set up the cluster.

2

### Set up your cluster

Follow the quick and easy process to [set up](#) your ONTAP cluster using System Manager.

3

### 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.

## 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/d20260206-1392446-ka57ya/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

- Bezel
- Storage system
- Rail kits with instructions
- Storage shelf
- Cisco Nexus 9332D-GX2B or 9364D-GX2A switch

#### Cables

- 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

SSN: XXYYYYYYYYYY



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"> <li>Sign in with your username and password.</li> <li>Select <b>Systems &gt; My Systems</b>.</li> <li>Confirm that the new serial number is listed.</li> <li>If the serial number is not listed, follow the instructions for new NetApp customers.</li> </ol>
New NetApp customer	<ol style="list-style-type: none"> <li>Click <b>Register Now</b>, and create an account.</li> <li>Select <b>Systems &gt; Register Systems</b>.</li> <li>Enter the storage system's serial number and requested details.</li> </ol> <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/d20260206-1392446-ka57ya/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](#).

## 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.

- 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/d20260206-1392446-ka57ya/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.

## 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/d20260206-1392446-ka57ya/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"> <li>Cisco Nexus 9332D-GX2B (400GbE)</li> <li>Cisco Nexus 9364D-GX2A (400GbE)</li> </ul>	<ul style="list-style-type: none"> <li>400GbE QSFP-DD breakout to 4x100GbE QSFP breakout cable cables</li> </ul> <p> Breakout cables are used for 100GbE connections between the switches, controllers, and disk shelves.</p> <ul style="list-style-type: none"> <li>100GbE cables to controller cluster and HA ports</li> <li>100GbE cables to disk shelves</li> <li>2 x 400GbE cables for ISL connections between switch A and switch B</li> <li>RJ-45 cables for management connections</li> </ul>

### 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/d20260206-1392446-ka57ya/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 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/d20260206-1392446-ka57ya/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 port on each switch to either of the management switches (if ordered) or connect them directly to your management network.

The management port is the upper-right port located on the PSU side of the switch. The CAT6 cable for each switch needs to be routed through the pass-through panel after the switches are installed to connect to the management switches or 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 A (Example)
    - e2a
    - e3a
  - Controller B (Example)
    - e2a
    - e3a

#### 100GbE cables






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

- Controller A (Example)
  - e2b
  - e3b
- Controller B (Example)
  - e2b
  - e3b

### 100GbE cables




== Step 3: Cable the cluster and HA connections Use the Cluster and HA interconnect cable to connect ports e1a and e7a to switch A and e1b and e7b to switch B. The e1a/e1b ports are used for the HA connections, and the e7a/e7b ports are used for the cluster connections.

### Steps

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

- Controller A
  - e1a (HA)
  - e7a (Cluster)
- Controller B
  - e1a (HA)
  - e7a (Cluster)

### 100GbE cables




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

- Controller A
  - e1b (HA)
  - e7b (Cluster)
- Controller B
  - e1b (HA)
  - e7b (Cluster)

### 100GbE cables






== Step 4: Cable the controller-to-switch storage connections Connect the controller storage ports to the switches. Ensure you have the correct cables and connectors for your switches. See [Hardware Universe](#) for more information.

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

- Controller A
  - e10a
  - e11a
- Controller B
  - e10a
  - e11a

#### 100GbE cables




2. Connect the following storage ports to any non-ISL port on switch B.

- Controller A

- e10b
- e11b
- Controller B
  - e10b
  - e11b

### 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/d20260206-1392446-ka57ya/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/d20260206-1392446-ka57ya/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:

- Insert the straightened end of a paperclip or narrow tipped ball point pen into the small hole to gently press the shelf ID button.
- 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.

- 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:

- 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 you've turned on your AFX 1K storage system, you [set up an AFX cluster](#).

= Set up your AFX storage system ONTAP cluster :icons: font :relative\_path: ./install-setup/ :imagesdir: /tmp/d20260206-1392446-ka57ya/source./install-setup/..media/

After your AFX hardware has been installed, you can complete the ONTAP cluster setup. This process involves two related configuration phases that you need to perform sequentially.

== 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.

### Before you begin

You need to have the following information:

- 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.

You'll also need the following:

- Network gateway IP address
- DNS domain names

- DNS name server IP addresses
- NTP server IP addresses
- Subnet mask

## 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 cluster and controller nodes.

Provide IP addresses and the subnet masks.

4. Configure the network services. Select **Continue**.

Define the details for your DNS and NTP servers.

5. 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

You should [Prepare to administer AFX](#) before using your AFX cluster in a production environment.

## Related information

- [Configure an AFX SVM](#)
- [Prepare to administer AFX](#)

= Prepare to administer your AFX storage system :icons: font :relative\_path: ./get-started/ :imagesdir: /tmp/d20260206-1392446-ka57ya/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 within 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.

### 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 including:

### Controller nodes per cluster

Each AFX cluster must have at least four nodes. The maximum number of nodes varies based on the ONTAP release.

## 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.

== 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](#)

## Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—with prior written permission of the copyright owner.

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

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

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

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

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

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

## Trademark information

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