



ASA r2 documentation

ASA r2

NetApp
February 11, 2026

Table of Contents

ASA r2 documentation	1
Release notes	2
What's new in ONTAP 9.18.1 for ASA r2 systems	2
Data protection	2
Networking	2
SAN data migration	2
Security	2
Storage efficiency	2
What's new in ONTAP 9.17.1 for ASA r2 systems	3
SAN data migration	3
Data protection	3
Storage management	4
What's new in ONTAP 9.16.1 for ASA r2 systems	4
Systems	4
Data protection	4
Protocol support	4
Storage efficiency	4
What's new in ONTAP 9.16.0 for ASA r2 systems	5
Systems	5
System Manager	5
Storage management	5
Data security	6
Changes to ONTAP limits and defaults affecting ASA r2 systems	6
Changes to ONTAP limits	6
Get started	7
Learn about ASA r2 storage systems	7
Quick start for ASA r2 storage systems	8
Install your ASA r2 system	8
Installation and setup workflow for ASA r2 storage systems	8
Installation requirements for ASA r2 storage systems	9
Prepare to install an ASA r2 storage system	11
Install your ASA r2 storage system	14
Cable the hardware for your ASA r2 storage system	15
Power on your ASA r2 storage system	47
Set up your ASA r2 system	53
Set up an ONTAP cluster on your ASA r2 storage system	53
SAN host configuration with ASA r2 systems	55
Enable data access from SAN hosts to your ASA r2 storage system	55
Use ONTAP to manage your data	58
ASA r2 storage system video demonstrations	58
Manage your storage	58
Provision ONTAP SAN storage on the ASA r2 systems	58
Clone data on ASA r2 storage systems	64

Manage host groups	67
Manage storage units	68
Migrate storage VMs	70
ASA r2 storage limits	75
Protect your data	76
Create snapshots to back up your data on ASA r2 storage systems	76
Manage snapshot reserve	81
Create an intercluster storage VM peer relationship on ASA r2 storage systems	83
Set up snapshot replication	83
Set up SnapMirror active sync	89
Manage SnapMirror active sync	93
Restore data on ASA r2 storage systems	97
Manage consistency groups	99
Manage ONTAP data protection policies and schedules on ASA r2 storage systems	106
Secure your data	108
Encrypt data at rest on ASA r2 storage systems	108
Migrate ONTAP data encryption keys between key managers on your ASA r2 system	109
Protect against ransomware attacks	111
Secure NVMe connections on your ASA r2 storage systems	116
Secure IP connections on your ASA r2 storage systems	117
Administer and monitor	119
Upgrade and revert ONTAP	119
Upgrade ONTAP on ASA r2 storage systems	119
Revert ONTAP on ASA r2 storage systems	119
Update firmware on ASA r2 storage systems	120
Manage client access to storage VMs on ASA r2 storage systems	121
Create a storage VM	122
Create IPspaces	122
Create subnets	123
Create a LIF (network interface)	123
Modify a LIF (network interfaces)	126
Manage cluster networking on ASA r2 storage systems	126
Add a broadcast domain	126
Reassign ports to a different broadcast domain	127
Create a VLAN	127
Monitor usage and increase capacity	128
Monitor cluster and storage unit performance on ASA r2 storage systems	128
Monitor cluster and storage unit utilization on ASA r2 storage systems	129
Increase storage capacity on ASA r2 storage systems	130
Optimize cluster security and performance with ASA r2 storage system insights	132
View cluster events and jobs on ASA r2 storage systems	132
Send email notifications for cluster events and audit logs	133
Manage nodes	133
Add ASA r2 nodes to an ONTAP cluster	133
Reboot a node on an ASA r2 storage system	134

Rename a node on an ASA r2 storage system	134
Manage user accounts and roles on ASA r2 storage systems	135
Configure active directory domain controller access	135
Configure LDAP	135
Configure SAML authentication	135
Create user account roles	136
Create an administrator account	136
Manage security certificates on ASA r2 storage systems	137
Generate a certificate signing request	137
Add a trusted certificate authority	137
Renew or delete a trusted certificate authority	138
Add a client/server certificate or local certificate authorities	138
Renew or delete a client/server certificate or local certificate authorities	138
Verify host connectivity on your ASA r2 storage system	139
Maintain your ASA r2 storage system	140
Learn more	141
ASA r2 for ONTAP power users	141
Compare ASA r2 systems to other ONTAP systems	141
ONTAP software support and limitations for ASA r2 storage systems	143
ONTAP CLI support for ASA r2 storage systems	144
REST API support for ASA r2	149
Common ONTAP features supported on ASA r2 systems	151
Data protection	151
Data security	151
Networking	152
SAN Protocols	152
System Manager	153
Get help	154
Manage AutoSupport on ASA r2 storage systems	154
Test AutoSupport connectivity	154
Add AutoSupport recipients	154
Send AutoSupport data	155
Suppress support case generation	155
Resume support case generation	155
Submit and view support cases for ASA r2 storage systems	155
Legal notices	157
Copyright	157
Trademarks	157
Patents	157
Privacy policy	157
Open source	157
ONTAP	157

ASA r2 documentation

Release notes

What's new in ONTAP 9.18.1 for ASA r2 systems

Learn about the new capabilities available in ONTAP 9.18.1 for ASA r2 systems.

Data protection

Update	Description
Increased support for SnapMirror active sync configurations	Support for SnapMirror active sync is increased from two-node clusters to four-node clusters.

Networking

Update	Description
IPsec hardware offload IPv6 support	IPsec hardware offload support is extended to IPv6.
OpenSSL PQC algorithms	ONTAP supports postquantum computing cryptographic algorithms for SSL. These algorithms provide additional protection against potential future quantum computing attacks, and are available when SSL FIPS mode is disabled.

SAN data migration

Update	Description
Support for storage VM migration	You can non-disruptively migrate a storage virtual machine (VM) from an ASA cluster to an ASA r2 cluster. This enables block workloads to move to ASA r2 systems while preserving data integrity and ensuring no application impact. The migration process is designed to maintain existing host mappings and LUN configurations, reducing operational effort and risk during migration.

Security

Update	Description
Support for automatic ARP/AI enablement	When you initialize a new 9.18.1 ASA r2 cluster or upgrade your cluster to 9.18.1, ARP/AI is automatically enabled by default on all newly created storage units after a 12-hour grace period. If you do not disable ARP/AI during the grace period, it is enabled cluster wide for newly created storage units when the grace period ends.

Storage efficiency

Update	Description
Support for NVMe copy offload	NVMe copy offload enables an NVMe host to offload copy operations from its CPU to the CPU of the ONTAP storage controller. The host can copy data from one NVMe namespace to another while reserving its CPU resources for application workloads.
Support for modification of snapshot reserve and automatic snapshot deletion	You can modify snapshot reserve and enable automatic snapshot deletion to limit the amount of space used for snapshots in your ASA r2 storage units. When snapshot reserve is set with automatic snapshot deletion, older snapshots are automatically deleted when space used by snapshots exceeds the snapshot reserve. This prevents application disruptions by not allowing snapshots to consume space in your storage unit intended for user data.

What's new in ONTAP 9.17.1 for ASA r2 systems

Learn about the new capabilities available in ONTAP 9.17.1 for ASA r2 systems.

SAN data migration

Update	Description
Support for data migration from a third-party storage system	SAN data migration using Foreign LUN Import (FLI) is supported for ASA r2 systems. FLI allows you to migrate data from a LUN on a third-party storage system to an ASA r2 system.

Data protection

Update	Description
Support for Autonomous Ransomware Protection with Artificial Intelligence (ARP/AI)	ARP/AI can be enabled on ASA r2 storage units. ARP/AI offers additional data protection by detecting and reporting potential ransomware attacks with no learning period.
SnapMirror Active Sync support for NVMe protocols	SnapMirror Active Sync adds support for VMware workloads with NVMe/TCP and NVMe/FC host access for two-node ONTAP clusters. VMware workload support for NVMe/TCP is contingent on the resolution of VMware Bug ID: TR1049746
Support for geometry changes to consistency groups in replication relationships	ASA r2 systems support geometry changes to consistency groups in a SnapMirror active sync or a asynchronous replication relationship without deleting the SnapMirror active sync relationship or breaking the asynchronous relationship. When geometry a change occurs on the primary consistency group, the change is replicated to the secondary consistency group.
Support for asynchronous replication of child consistency groups	Asynchronous replication policies can be applied to consistency groups in hierarchical relationships.

Storage management

Update	Description
Support for automatic workload balancing	Workloads are automatically balanced between the nodes of an HA pair to optimize performance and resource utilization.

What's new in ONTAP 9.16.1 for ASA r2 systems

Learn about the new capabilities available in ONTAP 9.16.1 for ASA r2 systems.

Systems

Update	Description
Systems	<p>The following NetApp ASA r2 systems are supported beginning with ONTAP 9.16.1. These systems deliver a unified hardware and software solution that creates a simplified experience specific to the needs of SAN-only customers.</p> <ul style="list-style-type: none">ASA A50ASA A30ASA A20ASA C30

Data protection

Update	Description
Support for encryption key migration between key managers	When you switch from the ONTAP onboard key manager to an external key manager at the cluster level, you can use the ONTAP command line interface (CLI) to easily migrate the encryption keys from one key manager to the other.
Support for hierarchical consistency groups	Hierarchical consistency groups allow you to create a parent consistency group that contains multiple child consistency groups. This simplifies data protection and management for complex data structures.

Protocol support

Update	Description
NVMe support for symmetric Active/Active multipathing	NVMe/FC and NVMe/TCP now support symmetric active-active architecture for multipathing so that all paths between the hosts and storage are active/optimized.

Storage efficiency

Update	Description
Support for automatic rebalancing of storage units	ONTAP will automatically rebalance storage units across your storage availability zones for optimal performance and capacity utilization.
NVMe space deallocation enabled by default	<p>Space deallocation (also called “hole punching” and “unmap”) is enabled for NVMe namespaces by default. Space deallocation allows a host to deallocate unused blocks from namespaces to reclaim space.</p> <p>This greatly improves overall storage efficiency, especially with file systems that have data high turnover.</p>

What's new in ONTAP 9.16.0 for ASA r2 systems

Learn about the new capabilities available in ONTAP 9.16.0 for ASA r2 systems.

Systems

Update	Description
Systems	<p>The following NetApp ASA r2 systems are available. These systems deliver a unified hardware and software solution that creates a simplified experience specific to the needs of SAN-only customers.</p> <ul style="list-style-type: none"> ASA A1K ASA A70 ASA A90

System Manager

Update	Description
Streamlined support for SAN-only customers	System Manager is streamlined to provide support for essential SAN functionality while removing visibility of features and functions not supported in SAN environments.

Storage management

Update	Description
Simplified storage management	<p>ASA r2 systems introduce the use of storage units with consistency groups for simplified storage management.</p> <ul style="list-style-type: none"> A <i>storage unit</i> makes storage space available to your SAN hosts for data operations. A storage unit refers to a LUN for SCSI hosts or an NVMe namespace for NVMe hosts. A <i>consistency group</i> is a collection of storage units that are managed as a single unit.

Data security

Update	Description
Onboard key manager and dual-layer encryption	ASA r2 systems support an onboard key manager and dual-layer (hardware and software) encryption.

Changes to ONTAP limits and defaults affecting ASA r2 systems

Learn about the changes to limits and defaults affecting ASA r2 systems. NetApp strives to help its customers understand the most important default and limit changes in each ONTAP release.

Changes to ONTAP limits

Feature	Limit change	Changed in release...
Storage VMs per cluster	The maximum number of supported storage virtual machines (VMs) per HA pair is increased from 32 to 256.	ONTAP 9.18.1
SnapMirror active sync	Support for SnapMirror active sync is increased from two-node clusters to four-node clusters.	ONTAP 9.18.1
Nodes per cluster	The maximum number of nodes per cluster is increased from 2 to 12.  If you are running ONTAP 9.16.1 with more than 2 nodes in a cluster, you cannot revert to ONTAP 9.16.0.	ONTAP 9.16.1
Storage units	The maximum number of storage units is increased from 2500 per HA pair to 10,000 per HA pair.	ONTAP 9.16.1

Get started

Learn about ASA r2 storage systems

The NetApp ASA r2 systems deliver a unified hardware and software solution that creates a simplified experience specific to the needs of SAN-only customers.

The following are classified as ASA r2 systems:

- ASA A1K
- ASA A90
- ASA A70
- ASA A50
- ASA A30
- ASA A20
- ASA C30

ASA r2 systems support all SAN protocols (iSCSI, FC, NVMe/FC, NVMe/TCP). The iSCSI, FC, NVMe/FC and NVMe/TCP protocols support symmetric active-active architecture for multipathing so that all paths between the hosts and storage are active/optimized. The iSCSI and NVMe/TCP protocols support direct attach between the hosts and storage. For Fibre Channel and NVMe/FC protocols direct attach is not supported.

On an ASA r2 system, ONTAP software and System Manager are streamlined to provide support for essential SAN functionality while removing features and functions not supported in SAN environments.

ASA r2 systems introduce the use of storage units with consistency groups:

- A *storage unit* makes storage space available to your SAN hosts for data operations. A storage unit refers to a LUN for SCSI hosts or an NVMe namespace for NVMe hosts.
- A *consistency group* is a collection of storage units that are managed as a single unit.

ASA r2 systems use storage units with consistency groups to simplify storage management and data protection. For example, suppose you have a database consisting of 10 storage units in a consistency group, and you need to back up the entire database. Instead of backing up each storage unit individually, you can protect the entire database by backing up the consistency group.

To help secure your data against malicious attacks such as theft or ransomware, ASA r2 systems support an on-board key manager, dual-layer encryption, multi-factor authentication and multi-admin verification. Tamper-proof snapshots are also supported on secondary ASA r2 systems.

ASA r2 systems do not support cluster mixing with ASA, AFF, or FAS systems.

For more information

- Learn more about ASA r2 systems support and limitations in the [NetApp Hardware Universe](#).
- Learn more about [the ASA r2 systems in comparison to the ASA systems](#).
- Learn more about the [NetApp ASA](#).

Quick start for ASA r2 storage systems

To get up and running with your ASA r2 system, you install your hardware components, set up your cluster, set up data access from your hosts to the storage system, and provision your storage.

1

Install and set up your hardware

[Install and set up your ASA r2 system and deploy it in your ONTAP environment.](#)

2

Set up your cluster

Use System Manager to guide you through a quick and easy process to [set up your ONTAP cluster](#).

3

Set up data access

[Connect your ASA r2 system to your SAN clients.](#)

4

Provision your storage

[Provision storage](#) to begin serving data to your SAN clients.

What's next?

You can now use System Manager to protect your data by [creating snapshots](#).

Install your ASA r2 system

Installation and setup workflow for ASA r2 storage systems

To install and configure your ASA r2 system, you review the hardware requirements, prepare your site, install and cable the hardware components, power on the system, and set up your ONTAP cluster.

1

[Review the hardware installation requirements](#)

Review the hardware requirements to install your ASA r2 storage system.

2

[Prepare to install the ASA r2 storage system](#)

To prepare to install your ASA r2 system, you need to get the site ready, check the environmental and electrical requirements, and ensure there's enough rack space. Then, unpack the equipment, compare its contents to the packing slip, and register the hardware to access support benefits.

3

[Install the hardware for the ASA r2 storage system](#)

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system in the cabinet or telco rack. Next, slide the shelves onto the rails. Finally, attach cable management devices to the rear of the storage system for organized cable routing.

4

Cable the controllers and storage shelves for the ASA r2 storage system

To cable the hardware, first connect the storage controllers to your network and then connect the controllers to your storage shelves.

5

Power on the ASA r2 storage system

Before you power on the controllers, power on each NS224 shelf and assign a unique shelf ID to ensure each shelf is uniquely identified within the setup.

Installation requirements for ASA r2 storage systems

Review the equipment needed and the lifting precautions for your ASA r2 storage system and storage shelves.

Equipment needed for install

To install your ASA r2 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 ball point pen for setting storage shelf IDs
- Phillips #2 screwdriver

Lifting precautions


ASA r2 storage systems and storage shelves are heavy. Exercise caution when lifting and moving these items.

Storage system weights

Take the necessary precautions when moving or lifting your ASA r2 storage system.


A1K

An ASA A1K storage system can weigh up to 62.83 lbs (28.5 kg). To lift the storage system, use two people or a hydraulic lift.




A70 and A90

An ASA A70 or ASA A90 storage system can weigh up to 151.68 lbs (68.8 kg). To lift the storage system, use four people or a hydraulic lift.




A20, A30, and A50

An ASA A20, ASA A30, or ASA A50 storage system can weigh up to 61.5 lbs (27.9 kg). To lift the storage system, use two people or a hydraulic lift.



C30

An ASA C30 storage system can weigh up to 61.5 lbs (27.9 kg). To lift the storage system, use two people or a hydraulic lift.




Storage shelf weights

Take the necessary precautions when moving or lifting your shelf.


NS224 shelf

An NS224 shelf can weigh up to 66.78 lbs (30.29 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.



NS224 shelf with NSM100B modules

An NS224 shelf with NSM100B modules can weigh up to 56.8 lbs (25.8 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 ASA r2 storage system](#).

Prepare to install an ASA r2 storage system

Prepare to install your ASA r2 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 ASA r2 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 any switches:

A1K

- 4U in an HA configuration
- 2U for each NS224 storage shelf
- 1U for most switches

A70, and A90

- 4U in an HA configuration
- 2U for each NS224 storage shelf
- 1U for most switches

A20, A30, and A50

- 2U for a storage system
- 2U for each NS224 storage shelf
- 1U for most switches

C30

- 2U for a storage system
- 2U for each NS224 storage shelf
- 1U for most switches

3. Install any required network switches.

See the [Switch documentation](#) for installation instructions and [NetApp Hardware Universe](#) for compatibility information.

Step 2: Unpack the boxes

After you've ensured that the site and the cabinet or rack that you plan to use for your ASA r2 storage system meet the required specifications, unpack all boxes and compare the contents to the items on 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 there are any discrepancies, note them down 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 (optional)
- Storage shelf (if you ordered additional storage)
- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables (if you ordered additional storage)
- USB-C serial port cable

Step 3: Register your storage system

After you've ensured that your site meets the requirements for your ASA r2 storage system specifications, and you've verified that you have all the parts you ordered, you should register your 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. Determine whether you need to register your storage system:

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

What's next?

After you've prepared to install your ASA r2 hardware, you [install the hardware for your ASA r2 storage system](#).

Install your ASA r2 storage system


After you prepare to install your ASA r2 storage system, install the hardware for the 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.
- Be aware of the safety concerns associated with 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 storage system 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. Make sure that the guiding pins on the cabinet or telco rack fit securely into the storage system guide slots.
 - c. Secure the storage system to the cabinet or telco rack using the included mounting screws.
3. Attach the bezel to the front of the storage system.
4. If your ASA r2 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.

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 ASA r2 system, you [cable the controllers and storage shelves for your ASA r2 system](#).

Cable the hardware for your ASA r2 storage system

After you install the rack hardware for your ASA r2 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 and slot priority details, see [NetApp Hardware Universe](#).
- The cluster/HA and host network cabling procedures show common configurations.


If you do not see your configuration in the cabling procedures, go to [NetApp Hardware Universe](#) for comprehensive configuration and slot priority information to properly cable your storage system.

- If you have an ASA A1K, ASA A70, or ASA A90 storage system, the I/O slots are numbered 1 through 11.



- The cabling graphics have arrow icons showing 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.



- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Cable the cluster/HA connections

Cable the controllers to your ONTAP cluster. This procedure differs depending on your storage system model and I/O module configuration.



The cluster interconnect traffic and the HA traffic share the same physical ports.

A1K

Create the ONTAP cluster connections. For switchless clusters, connect the controllers to each other. For switched clusters, connect the controllers to the cluster network switches.


Switchless cluster cabling

Use the Cluster/HA interconnect cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A to port e1a on Controller B.
2. Connect port e7a on Controller A to port e7a on Controller B.

Cluster/HA interconnect cables



Switched cluster cabling

Use the 100 GbE cable to connect ports e1a to e1a and ports e7a to e7a.




Switched cluster configurations are supported in 9.16.1 and later.

Steps

1. Connect port e1a on Controller A and port e1a on Controller B to cluster network switch A.
2. Connect port e7a on Controller A and port e7a on Controller B to cluster network switch B.

100 GbE cable



A70 and A90

Create the ONTAP cluster connections. For switchless clusters, connect the controllers to each other. For switched clusters, connect the controllers to the cluster network switches.


Switchless cluster cabling

Use the the Cluster/HA interconnect cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A to port e1a on Controller B.
2. Connect port e7a on Controller A to port e7a on Controller B.

Cluster/HA interconnect cables



Switched cluster cabling

Use the 100 GbE cable to connect ports e1a to e1a and ports e7a to e7a.




Switched cluster configurations are supported in 9.16.1 and later.

Steps

1. Connect port e1a on Controller A and port e1a on Controller B to cluster network switch A.
2. Connect port e7a on Controller A and port e7a on Controller B to cluster network switch B.

100 GbE cable



A20, A30, and A50

Create the ONTAP cluster connections. For switchless clusters, connect the controllers to each other. For switched clusters, connect the controllers to the cluster network switches.

The cluster/HA cabling examples show common configurations.



If you do not see your configuration here, go to [NetApp Hardware Universe](#) for comprehensive configuration and slot priority information to cable your storage system.

Switchless cluster cabling

Connect the controllers to each other to create the ONTAP cluster connections.

ASA A30 and ASA A50 with two 2-port 40/100 GbE I/O modules

Steps

1. Connect the Cluster/HA interconnect connections:



The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

- Connect controller A port e2a to controller B port e2a.
- Connect controller A port e4a to controller B port e4a.




I/O module ports e2b and e4b are unused and available for host network connectivity.

100 GbE Cluster/HA interconnect cables



Controller A



Controller B

ASA A30 and ASA A50 with one 2-port 40/100 GbE I/O module

Steps

1. Connect the Cluster/HA interconnect connections:




The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

- a. Connect controller A port e4a to controller B port e4a.
- b. Connect controller A port e4b to controller B port e4b.

100 GbE Cluster/HA interconnect cables



Controller A



Controller B

ASA A20 with one 2-port 10/25 GbE I/O module

Steps

1. Connect the Cluster/HA interconnect connections:




The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 10/25 GbE.

- a. Connect controller A port e4a to controller B port e4a.
- b. Connect controller A port e4b to controller B port e4b.

25 GbE Cluster/HA interconnect cables



Controller A



Controller B

Switched cluster cabling

Connect the controllers to the cluster network switches to create the ONTAP cluster connections.

ASA A30 or ASA A50 with two 2-port 40/100 GbE I/O modules

Steps

1. Cable the Cluster/HA interconnect connections:




The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

- a. Connect controller A port e4a to cluster network switch A.
- b. Connect controller A port e2a to cluster network switch B.
- c. Connect controller B port e4a to cluster network switch A.
- d. Connect controller B port e2a to cluster network switch B.



I/O module ports e2b and e4b are unused and available for host network connectivity.

40/100 GbE Cluster/HA interconnect cables



ASA A30 or ASA A50 with one 2-port 40/100 GbE I/O module

Steps


1. Cable the controllers to the cluster network switches:



The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

- a. Connect controller A port e4a to cluster network switch A.
- b. Connect controller A port e4b to cluster network switch B.
- c. Connect controller B port e4a to cluster network switch A.
- d. Connect controller B port e4b to cluster network switch B.

40/100 GbE Cluster/HA interconnect cables



ASA A20 with one 2-port 10/25 GbE I/O module


1. Cable the controllers to the cluster network switches:



The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 10/25 GbE.

- a. Connect controller A port e4a to cluster network switch A.
- b. Connect controller A port e4b to cluster network switch B.
- c. Connect controller B port e4a to cluster network switch A.
- d. Connect controller B port e4b to cluster network switch B.

10/25 GbE Cluster/HA interconnect cables



C30

Create the ONTAP cluster connections. For switchless clusters, connect the controllers to each other. For switched clusters, connect the controllers to the cluster network switches.

The cluster/HA cabling examples show common configurations.



If you do not see your configuration here, go to [NetApp Hardware Universe](#) for comprehensive configuration and slot priority information to cable your storage system.

Switchless cluster cabling

Connect the controllers to each other to create the ONTAP cluster connections.

ASA C30 with two 2-port 40/100 GbE I/O modules

Steps

1. Cable the Cluster/HA interconnect connections:



The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

- a. Connect controller A port e2a to controller B port e2a.
- b. Connect controller A port e4a to controller B port e4a.




I/O module ports e2b and e4b are unused and available for host network connectivity.

100 GbE Cluster/HA interconnect cables



Controller A



Controller B

ASA C30 with one 2-port 40/100 GbE I/O module

Steps

1. Cable the Cluster/HA interconnect connections:




The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

- a. Connect controller A port e4a to controller B port e4a.
- b. Connect controller A port e4b to controller B port e4b.

100 GbE Cluster/HA interconnect cables



Controller A



Controller B

Switched cluster cabling

Connect the controllers to the cluster network switches to create the ONTAP cluster connections.

ASA C30 with two 2-port 40/100 GbE I/O modules

Steps

1. Cable the Cluster/HA interconnect connections:




The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O modules in slots 2 and 4). The ports are 40/100 GbE.

- a. Connect controller A port e4a to cluster network switch A.
- b. Connect controller A port e2a to cluster network switch B.
- c. Connect controller B port e4a to cluster network switch A.
- d. Connect controller B port e2a to cluster network switch B.



I/O module ports e2b and e4b are unused and available for host network connectivity.

40/100 GbE Cluster/HA interconnect cables



ASA C30 with one 2-port 40/100 GbE I/O module

Steps


1. Connect the controllers to the cluster network switches:



The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

- a. Connect controller A port e4a to cluster network switch A.
- b. Connect controller A port e4b to cluster network switch B.
- c. Connect controller B port e4a to cluster network switch A.
- d. Connect controller B port e4b to cluster network switch B.

40/100 GbE Cluster/HA interconnect cables



Step 2: Cable the host network connections

Connect the controllers to your host network.

This procedure differs depending on your storage system model and I/O module configuration.

A1K

Connect the Ethernet module ports to your host network.

The following are some typical host network cabling examples. See [NetApp Hardware Universe](#) for your specific system configuration.


Steps

1. Connect ports e9a and e9b to your Ethernet data network switch.



For maximum system performance for cluster and HA traffic, do not use ports e1b and e7b ports for host network connections. Use a separate host card to maximize performance.


100 GbE cable



2. Connect your 10/25 GbE host network switches.

10/25 GbE Host





A70 and A90

Connect the Ethernet module ports to your host network.

The following are some typical host network cabling examples. See [NetApp Hardware Universe](#) for your specific system configuration.


Steps

1. Connect ports e9a and e9b to your Ethernet data network switch.




For maximum system performance for cluster and HA traffic, do not use ports e1b and e7b ports for host network connections. Use a separate host card to maximize performance.

100 GbE cable



2. Connect your 10/25 GbE host network switches.

4-ports, 10/25 GbE Host



A20, A30, and A50

Connect the Ethernet module ports or the Fibre Channel (FC) module ports to your host network.

The host network cabling examples show common configurations.



If you do not see your configuration here, go to [NetApp Hardware Universe](#) for comprehensive configuration and slot priority information to cable your storage system.

Ethernet host cabling


ASA A30 and ASA A50 with two 2-port 40/100 GbE I/O modules

On each controller, connect ports e2b and e4b to the Ethernet host network switches.



The ports on I/O modules in slot 2 and 4 are 40/100 GbE (host connectivity is 40/100 GbE).


40/100 GbE cables



ASA A20, A30, and A50 with one 4-port 10/25 GbE I/O module

On each controller, connect ports e2a, e2b, e2c and e2d to the Ethernet host network switches.

10/25 GbE cables




FC host cabling

ASA A20, A30, and A50 with One 4-port 64 Gb/s FC I/O module

On each controller, connect ports 1a, 1b, 1c and 1d to the FC host network switches.

64 Gb/s FC cables



C30

Connect the Ethernet module ports or the Fibre Channel (FC) module ports to your host network.

The host network cabling examples show common configurations.



If you do not see your configuration here, go to [NetApp Hardware Universe](#) for comprehensive configuration and slot priority information to cable your storage system.

Ethernet host cabling

ASA C30 with two 2-port 40/100 GbE I/O modules


Steps

1. On each controller, cable ports e2b and e4b to the Ethernet host network switches.



The ports on I/O modules in slot 2 and 4 are 40/100 GbE (host connectivity is 40/100 GbE).

40/100 GbE cables




ASA C30 with one 4-port 10/25 GbE I/O module

Steps

1. On each controller, cable ports e2a, e2b, e2c and e2d to the Ethernet host network switches.

10/25 GbE cables




ASA C30 with one 4-port 64 Gb/s FC I/O module

Steps

1. On each controller, cable ports 1a, 1b, 1c and 1d to the FC host network switches.

64 Gb/s FC cables



Step 3: Cable the management network connections

Connect the controllers to your management network.


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

A1K

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.

A70 and A90

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.

A20, A30, and A50

Connect the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables



Controller B




Do not plug in the power cords yet.

C30

Connect the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables



Controller B



Do not plug in the power cords yet.

Step 4: Cable the shelf connections

The following cabling procedures show how to connect your controllers to a storage shelf.

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

A1K

The AFF A1K storage systems support NS224 shelves with either the NSM100 or NSM100B module. The major differences between the modules are:

- NSM100 shelf modules use built-in port e0a and e0b.
- NSM100B shelf modules use ports e1a and e1b in slot 1.

The following cabling example shows NSM100 modules in the NS224 shelves when referring to shelf module ports.


Choose one of the following cabling options that matches your setup.

Option 1: One NS224 storage shelf


Connect each controller to the NSM modules on the NS224 shelf. The graphics show cabling from each of the controllers: Controller A cabling is shown in blue and Controller B cabling is shown in yellow.

Steps

1. On controller A, connect the following ports:
 - a. Connect port e11a to NSM A port e0a.
 - b. Connect port e11b to port NSM B port e0b.



2. On controller B, connect the following ports:
 - a. Connect port e11a to NSM B port e0a.
 - b. Connect port e11b to NSM A port e0b.




Option 2: Two NS224 storage shelves


Connect each controller to the NSM modules on both NS224 shelves. The graphics show cabling from each of the controllers: Controller A cabling is shown in blue and Controller B cabling is shown in yellow.

Steps

1. On controller A, connect the following ports:
 - a. Connect port e11a to shelf 1 NSM A port e0a.
 - b. Connect port e11b to shelf 2 NSM B port e0b.
 - c. Connect port e10a to shelf 2 NSM A port e0a.
 - d. Connect port e10b to shelf 1 NSM A port e0b.



2. On controller B, connect the following ports:
 - a. Connect port e11a to shelf 1 NSM B port e0a.
 - b. Connect port e11b to shelf 2 NSM A port e0b.
 - c. Connect port e10a to shelf 2 NSM B port e0a.
 - d. Connect port e10b to shelf 1 NSM A port e0b.



A70 and A90

The AFF A70 and 90 storage systems support NS224 shelves with either the NSM100 or NSM100B module. The major differences between the modules are:

- NSM100 shelf modules use built-in ports e0a and e0b.
- NSM100B shelf modules use ports e1a and e1b in slot 1.

The following cabling example shows NSM100 modules in the NS224 shelves when referring to shelf module ports.

Choose one of the following cabling options that matches your setup.

Option 1: One NS224 storage shelf


Connect each controller to the NSM modules on the NS224 shelf. The graphics show cabling from each of the controllers: Controller A cabling is shown in blue and Controller B cabling is shown in yellow.

100 GbE QSFP28 copper cables




Steps

1. Connect controller A port e11a to NSM A port e0a.
2. Connect controller A port e11b to port NSM B port e0b.



3. Connect controller B port e11a to NSM B port e0a.
4. Connect controller B port e11b to NSM A port e0b.



Option 2: Two NS224 storage shelves


Connect each controller to the NSM modules on both NS224 shelves. The graphics show cabling from each of the controllers: Controller A cabling is shown in blue and Controller B cabling is shown in yellow.

100 GbE QSFP28 copper cables




Steps

1. On controller A, connect the following ports:
 - a. Connect port e11a to shelf 1, NSM A port e0a.
 - b. Connect port e11b to shelf 2, NSM B port e0b.
 - c. Connect port e8a to shelf 2, NSM A port e0a.
 - d. Connect port e8b to shelf 1, NSM B port e0b.



2. On controller B, connect the following ports:
 - a. Connect port e11a to shelf 1, NSM B port e0a.
 - b. Connect port e11b to shelf 2, NSM A port e0b.
 - c. Connect port e8a to shelf 2, NSM B port e0a.
 - d. Connect port e8b to shelf 1, NSM A port e0b.



A20, A30, and A50

The NS224 shelf cabling procedure shows NSM100B modules instead of NSM100 modules. The cabling is the same regardless of the type of NSM modules used, only the port names are different:

- NSM100B modules use ports e1a and e1b on an I/O module in slot 1.
- NSM100 modules use built-in (onboard) ports e0a and e0b.

You cable each controller to each NSM module on the NS224 shelf using the storage cables that came with your storage system, which could be the following cable type:


100 GbE QSFP28 copper cables



The graphics show controller A cabling in blue and controller B cabling in yellow.


Steps

1. Connect controller A to the shelf:
 - a. Connect controller A port e3a to NSM A port e1a.
 - b. Connect controller A port e3b to NSM B port e1b.



2. Connect controller B to the shelf:

- Connect controller B port e3a to NSM B port e1a.
- Connect controller B port e3b to NSM A port e1b.



C30

The NS224 shelf cabling procedure shows NSM100B modules instead of NSM100 modules. The cabling is the same regardless of the type of NSM modules used, only the port names are different:

- NSM100B modules use ports e1a and e1b on an I/O module in slot 1.
- NSM100 modules use built-in (onboard) ports e0a and e0b.

You cable each controller to each NSM module on the NS224 shelf using the storage cables that came with your storage system, which could be the following cable type:

100 GbE QSFP28 copper cables




The graphics show controller A cabling in blue and controller B cabling in yellow.

Steps


- Connect controller A to the shelf:
 - Connect controller A port e3a to NSM A port e1a.

b. Connect controller A port e3b to NSM B port e1b.



2. Connect controller B to the shelf:

- a. Connect controller B port e3a to NSM B port e1a.
- b. Connect controller B port e3b to NSM A port e1b.



What's next?

After you've connected the storage controllers to your network and then connected the controllers to your storage shelves, you [power on the ASA r2 storage system](#).

Power on your ASA r2 storage system

After you install the rack hardware for your ASA r2 storage system and install the cables for the controllers and storage shelves, you should power on your storage shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

About this task

- A valid shelf ID is 01 through 99.

If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed shelf ID of 00.


- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, 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 powers on and boots automatically when plugged into the power source.

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 press the shelf ID button.
- b. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

It can take up to 15 seconds for the number to blink. This activates the 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.

A power supply is powered on as soon as the power cord is plugged in. Its bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

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

Steps


1. Connect your laptop to the serial console port. This will allow 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.


A1K




A70 and A90



A20, A30, and A50




C30




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.


A1K




A70 and A90



A20, A30, and A50



C30



- The system initiates the boot process. The initial boot sequence can take up to eight minutes.
- During the boot process, you will observe the LEDs flashing and the fans activating, signaling that the controllers are powering up.
- Be aware that the fans may emit a high level of noise when they first start up. The fan noise during start-up is normal.
- For ASA A20, A30, A50 and ASA C30 storage systems, the shelf ID display on the front of the system chassis does not illuminate.

4. Secure the power cords using the securing device on each power supply.

What's next?

After you've turned on your ASA r2 storage system, you [set up an ONTAP ASA r2 cluster](#).

Set up your ASA r2 system

Set up an ONTAP cluster on your ASA r2 storage system

ONTAP System Manager guides you through a quick and easy workflow to set up an ONTAP ASA r2 cluster.

During cluster setup, your default data storage virtual machine (VM) is created. Optionally, you can enable the Domain Name System (DNS) to resolve host names, set your cluster to use the Network Time Protocol (NTP) for time synchronization and enable encryption of data at rest.

In certain cases, you might need to [use the ONTAP command line interface \(CLI\) to set up your cluster](#). You should use the CLI, for example, if your security protocols do not allow you to connect a laptop to your management switches, or if you are using a non-windows operating system.

Before you begin

Gather 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 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
- Partner node IP address
- DNS domain names
- DNS name server IP addresses
- NTP server IP addresses
- Data 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.

System Manager opens.

2. Under **Password**, create a strong password for the admin account.

The password must be at least eight characters long and must contain at least one letter and one number.

3. Reenter the password to confirm and then select **Continue**.

4. Under **Network addresses**, enter a storage system name or accept the default name.

If you change the default storage system name, the new name must begin with a letter and must be fewer than 44 characters. You can use a period (.), hyphen (-) or underscore (_) in the name.

5. Enter the cluster management IP address, subnet mask, gateway IP address and the IP address of the partner node; then select **Continue**.

6. Under **Network services**, select the desired options to **Use the Domain Name System (DNS) for resolving host names** and to **Use the Network Time Protocol (NTP) to keep times synchronized**.

If you choose to use the DNS, enter the DNS domain and name servers. If you choose to use NTP, enter the NTP servers; then select **Continue**.

7. Under **Encryption**, enter a passphrase for the Onboard Key Manager (OKM).

Encryption of data at rest using an Onboard Key Manager (OKM) is selected by default. If you want to use an external key manager, update the selections.

Optionally, you can configure your cluster for encryption after cluster setup is complete.

8. Select **Initialize**.

When setup is complete, you are redirected to the cluster's management IP address.

9. Under **Network**, select **Configure protocols**.

To configure IP (iSCSI and NVMe/TCP), do this...	To configure FC and NVMe/FC, do this...
<ul style="list-style-type: none">a. Select IP; then select Configure IP interfaces.b. Select Add a subnet.c. Enter a name for the subnet, then enter the subnet IP addresses.d. Enter the subnet mask, and optionally enter a gateway; then select Add.e. Select the subnet you just created; then select Save.f. Select Save.	<ul style="list-style-type: none">a. Select FC; then select Configure FC interfaces and/or Configure NVMe/FC interfaces.b. Select the FC and/or NVMe/FC ports; then select Save.

10. Optionally, download and run [ActiveIQ Config Advisor](#) to confirm your configuration.

ActiveIQ Config Advisor is a tool for NetApp systems that checks for common configuration errors.

What's next?

You are ready to [set up data access](#) from your SAN clients to your ASA r2 system.


SAN host configuration with ASA r2 systems

ASA r2 systems follow the same recommendations and guidelines for SAN host configuration as all other ONTAP systems.

It is recommended that you use two or more switches to connect your storage system to one or more SAN hosts. For iSCSI configurations, the network topology connecting your hosts, switches and storage system is referred to as a *network*. For FC and FC-NVMe configurations, this same network topology is referred to as a *fabric*.

Multi-network or multi-fabric configurations (those using two or more switches) are recommended because they provide redundancy at both the switch and storage layer. This redundancy makes your storage system more fault tolerant and provides support for nondisruptive operations.

The following illustration is an example of an FC configuration with multiple hosts using two fabrics to access a single HA pair. The FC target port numbers (0c, 0d, 1a, 1b) are also examples. The actual port numbers vary depending on your system model and whether you are using expansion adapters.



Learn more about [SAN configuration for iSCSI hosts](#).

Learn more about [SAN configuration for FC and FC/NVMe hosts](#).

Zoning recommendation for FC hosts

You should configure your FC hosts to use zoning. ASA r2 systems follow the same FC host zoning recommendations and guidelines as all other ONTAP systems.

A zone is a logical grouping of one or more ports within a fabric. For devices to be able to discover each other, establish sessions with one another, and communicate, both ports need to have a common zone membership.

Learn more about [FC/FC-NVMe zoning](#).

Enable data access from SAN hosts to your ASA r2 storage system

To set up data access, you should ensure that the critical parameters and settings on

your SAN client for proper operation with ONTAP are configured correctly. If you are providing storage for your VMware environment, you should install OTV 10.3 to simplify the management of your ASA r2 storage.

Set up data access from SAN hosts

The configuration necessary to set up data access to your ASA r2 system from your SAN hosts varies depending on the host operating system and the protocol. Correct configuration is important for best performance and successful failover.

See the ONTAP SAN host documentation for [VMware vSphere SCSI clients](#), [VMware vSphere NVMe clients](#) and [other SAN clients](#) to properly configure your hosts to connect to your ASA r2 system.

Migrate VMware virtual machines

If you need to migrate your VM workload from an ASA storage system to an ASA r2 storage system, NetApp recommends that you use [VMware vSphere vMotion](#) to perform a live, non-disruptive migration of your data.

ASA r2 storage units are thin provisioned by default. When migrating your VM workload, virtual disks (VMDKs) should also be thin provisioned.

Related information

- Learn more about [the advantages of using ONTAP for vSphere](#).
- Learn about [VMware Live Site Recovery with ONTAP](#).
- Learn about [continuous availability solutions for vSphere environments](#).
- Learn more about [how to set up Broadcom VMware ESXi iSCSI MPIO with ONTAP SAN ASA storage systems](#).

Migrate data from a third-party storage system

Beginning with ONTAP 9.17.1, you can use Foreign LUN Import (FLI) to migrate data from a LUN on a third-party storage system to an ASA r2 system. Using FLI for your data migration can help you mitigate the risk of data loss and downtime during the migration process.

FLI supports both online and offline migrations. In an online migration, the client system stays online while data is copied from the third-party storage system to the ONTAP storage system. Online migrations are supported by Windows, Linux, and ESXi host operating systems. In an offline migration, the client system is taken offline, the LUN data is copied from the third-party storage system to the ONTAP storage system, and then the client system is brought back online.

- Learn how to perform an [FLI offline migration](#).
- Learn how to perform an [FLI online migrations](#).

Configure your ASA r2 system as a storage provider in your VMware environment

You can use ONTAP tools for VMware to easily enable your ASA r2 system as a storage provider in your VMware environment.

ONTAP tools for VMware vSphere is a set of tools that work in conjunction with VMware vCenter Server Virtual Appliance (vCSA) for easy management of virtual machines on your VMware ESXi hosts.

ASA r2 systems are supported by [ONTAP tools for VMware vSphere 10.3](#) and later.

Learn how to [Deploy ONTAP tools for VMware](#) and then use it to do the following:

- [Add vCenter Server instances](#)
- [Configure your ESXi host settings](#)
- [Discover your ASA r2 storage system and hosts](#)

What's next?

You are ready to provision storage to enable your SAN hosts to read and write data to storage units.

Use ONTAP to manage your data

ASA r2 storage system video demonstrations

View short videos that demonstrate how to use ONTAP System Manager to quickly and easily perform common task on your ASA r2 storage systems.

[Configure SAN protocols on your ASA r2 system](#)

[Video transcript](#)

[Provision SAN storage on your ASA r2 system](#)

[Video transcript](#)

[Replicate data to a remote cluster from an ASA r2 system](#)

[Video transcript](#)

Manage your storage

Provision ONTAP SAN storage on the ASA r2 systems

When you provision storage, you enable your SAN hosts to read from and write data to ASA r2 storage systems. To provision storage, you use ONTAP System Manager to create storage units, add host initiators, and map the host to a storage unit. You also need to perform steps on the host to enable read/write operations.

Create storage units

On an ASA r2 system, a storage unit makes storage space available to your SAN hosts for data operations. A storage unit refers to a LUN for SCSI hosts or an NVMe namespace for NVMe hosts. If your cluster is configured to support SCSI hosts, you are prompted to create a LUN. If your cluster is configured to support NVMe hosts, you are prompted to create an NVMe namespace.

An ASA r2 storage unit has a maximum capacity of 128 TB. See the [NetApp Hardware Universe](#) for the most current storage limits for ASA r2 systems.

You add and map host initiators to the storage unit as part of the storage unit creation process. You can also [add](#) and [map](#) host initiators after you create the storage units.

Beginning with ONTAP 9.18.1, you can modify the snapshot reserve and enable automatic snapshot deletion when you create a storage unit. The snapshot reserve is the amount of space in the storage unit reserved specifically for snapshots. When snapshot reserve is set with automatic snapshot deletion, older snapshots are automatically deleted when space used by snapshots exceeds the snapshot reserve.

[Learn more about snapshot reserve on ASA r2 systems.](#)

Storage units are thin provisioned by default. Thin provisioning allows the storage unit to grow up to the size allocated but doesn't reserve the space in advance. Space is allocated dynamically from available free space as needed. This allows you to realize greater storage efficiency by [over provisioning](#) your available space. For

example, suppose you have 1 TB of free space and you need to create four 1 TB storage units. Instead of immediately adding 3 TB of additional storage capacity to your system, you can create the storage units, monitor space utilization, and increase your storage capacity as the the storage units consume actual space. Learn more about [thin provisioning](#).

Steps

1. In System Manager, select **Storage**; then select  .
2. Enter a name for the new storage unit.
3. Enter the number of units you want to create.

If you create more than one storage unit, each unit is created with the same capacity, host operating system, and host mapping.

To optimize workload balancing across the storage availability zone, create an even number of storage units.

4. Enter the storage unit capacity; then select the host operating system.



If you are creating more than one storage unit, each unit is created with the same capacity. Multiply the number of storage units you are creating by the desired capacity to ensure you have enough usable space. If you don't have enough free space and you chose to over provision, monitor utilization closely to avoid running out of space and losing data.

5. Accept the auto-selected **host mapping** or select a different host group for the storage unit to be mapped to.

Host mapping refers to the host group that the new storage unit will be mapped to. If there is a pre-existing host group for the type of host you selected for your new storage unit, the pre-existing host group is auto-selected for your host mapping. You can accept the host group that is auto-selected or you can select a different host group.

If there is no pre-existing host group for hosts running on the operating system you specified, ONTAP creates a new host group automatically .

6. If you want to do any of the following, select **More Options** and complete the required steps.

Option	Steps
Change the default Quality of Service (QoS) policy If the default QoS policy has not previously been set on the storage virtual machine (VM) on which the storage unit is being created, this option is not available.	<ol style="list-style-type: none">a. Under Storage and optimization, next to Quality of service (QoS), select .b. Select an existing QoS policy.

Option	Steps
Create a new QoS policy	<p>a. Under Storage and optimization, next to Quality of service (QoS), select .</p> <p>b. Select Define new policy.</p> <p>c. Enter a name for the new QoS policy.</p> <p>d. Set a QoS limit, a QoS guarantee, or both.</p> <p>i. Optionally, under Limit, enter a maximum throughput limit, a maximum IOPS limit, or both.</p> <p>Setting a maximum throughput and IOPS for a storage unit restricts its impact on system resources so that it does not degrade the performance of critical workloads.</p> <p>ii. Optionally, under Guarantee, enter a minimum throughput, a minimum IOPS, or both.</p> <p>Setting a minimum throughput and IOPS for a storage unit guarantees that it meets minimum performance targets regardless of demand by competing workloads.</p> <p>e. Select Add.</p>
Change the default performance service level.	<p>a. Under Storage and optimization, next to the Performance service level, select .</p> <p>b. Select Performance.</p> <p>ASA r2 systems offer two performance levels. The default performance level is Extreme, which is the highest available level. You can lower the level to Performance.</p>
Modify the default snapshot reserve and enable automatic snapshot deletion.	<p>a. Under Snapshot reserve %, enter the numeric value for the percentage of the storage unit space you want to allocate to snapshots.</p> <p>b. Select Automatically delete older snapshots.</p>
Add a new SCSI host	<p>a. Under Host information, select SCSI for the connection protocol.</p> <p>b. Select the host operating system.</p> <p>c. Under Host Mapping, select New hosts.</p> <p>d. Select FC or iSCSI.</p> <p>e. Select existing host initiators or select Add initiator to add a new host initiator.</p> <p>An example of a valid FC WWPN is "01:02:03:04:0a:0b:0c:0d". Examples of valid iSCSI initiator names are "iqn.1995-08.com.example:string" and "eui.0123456789abcdef".</p>

Option	Steps
Create a new SCSI host group	<ol style="list-style-type: none"> Under Host information, select SCSI for the connection protocol. Select the host operating system. Under Host Mapping, select New host group. Enter a name for the host group; then select the hosts to add to the group.
Add a new NVMe subsystem	<ol style="list-style-type: none"> Under Host information, select NVMe for the connection protocol. Select the host operating system. Under Host Mapping, select New NVMe subsystem. Enter a name for the subsystem or accept the default name. Enter a name for the initiator. If you want to enable in-band authentication or Transport Layer Security (TLS), select <input checked="" type="checkbox"/>; then select your options. <p>In-band authentication allows secure bidirectional and unidirectional authentication between your NVMe hosts and your ASA r2 system.</p> <p>TLS encrypts all data sent over the network between your NVMe/TCP hosts and your ASA r2 system.</p> Select Add initiator to add more initiators. <p>Format the host NQN as <nqn.yyyy-mm> followed by a fully qualified domain name. The year should be equal to or later than 1970. The total maximum length should be 223. An example of a valid NVMe initiator is nqn.2014-08.com.example:string</p>

7. Select **Add**.

What's next?

Your storage units are created and mapped to your hosts. You can now [create snapshots](#) to protect the data on your ASA r2 system.

For more information

Learn more about [how ASA r2 systems use storage virtual machines](#).

Add host initiators

You can add new host initiators to your ASA r2 system at any time. Initiators make the hosts eligible to access storage units and perform data operations.

Before you begin

If you want to replicate the host configuration to a destination cluster during the process of adding your host initiators, your cluster must be in a replication relationship. Optionally, you can [create a replication relationship](#) after your host is added.

Add host initiators for SCSI or NVMe hosts.

SCSI hosts

Steps

1. Select **Host**.
2. Select **SCSI**; then select  .
3. Enter the host name, select the host operating system and enter a host description.
4. If you want to replicate the host configuration to a destination cluster, select **Replicate host configuration**; then select the destination cluster.

Your cluster must be in a replication relationship to replicate the host configuration.

5. Add new or existing hosts.

Add new hosts	Add existing hosts
<ol style="list-style-type: none">Select New hosts.Select FC or iSCSI; then select the host initiators. Optionally, select Configure host proximity. Configuring host proximity enables ONTAP to identify the controller nearest to the host for data path optimization and latency reduction. This is only applicable if you have replicated data to a remote location. If you have not set up snapshot replication, you do not need to select this option.If you need to add new initiators, select Add initiators.	<ol style="list-style-type: none">Select Existing hosts.Select the host you want to add.Select Add.

6. Select **Add**.

What's next?

Your SCSI hosts are added to your ASA r2 system and you are ready to map your hosts to your storage units.

NVMe hosts

Steps

1. Select **Host**.
2. Select **NVMe**; then select  .
3. Enter a name for the NVMe subsystem, select the host operating system and enter a description.
4. Select **Add initiator**.

What's next?

Your NVMe hosts are added to your ASA r2 system and you are ready to map your hosts to your storage units.

Map the storage unit to a host

After creating ASA r2 storage units and adding host initiators, map hosts to storage units to begin serving data. Storage units are mapped to hosts as part of the storage unit creation process. You can also map existing storage units to new or existing hosts at any time.

Steps

1. Select **Storage**.
2. Hover over the name of the storage unit you want to map.
3. Select ; then select **Map to hosts**.
4. Select the hosts you want to map to the storage unit; then select **Map**.

What's next?

Your storage unit is mapped to your hosts and you are ready to complete the provisioning process on your hosts.

Complete host-side provisioning

After you have created your storage units, added your host initiators and mapped your storage units, there are steps you must perform on your hosts before they can read and write data on your ASA r2 system.

Steps

1. For FC and FC/NVMe, zone your FC switches by WWPN.
Use one zone per initiator and include all target ports in each zone.
2. Discover the new storage unit.
3. Initialize the storage unit and create a file system.
4. Verify that your host can read and write data on the storage unit.

What's next?

You have completed the provisioning process and are ready to begin serving data. You can now [create snapshots](#) to protect the data on your ASA r2 system.

For more information

For more details about host-side configuration, see the [ONTAP SAN host documentation](#) for your specific host.

Clone data on ASA r2 storage systems

Data cloning creates copies of storage units and consistency groups on your ASA r2 system using ONTAP System Manager that can be used for application development, testing, backups, data migration or other administrative functions.

Clone storage units

When you clone a storage unit, you create a new storage unit on your ASA r2 system that is a point-in-time, writable copy of the storage unit you cloned.

Steps

1. In System Manager, select **Storage**.

2. Hover over the name of the storage unit you want to clone.
3. Select ; then select **Clone**.
4. Accept the default name for the new storage unit that will be created as a clone or enter a new one.
5. Select the host operating system.

A new snapshot is created for the clone by default.

6. If you want to use an existing snapshot, create a new host group, or add a new host, select **More Options**.

Option	Steps
Use an existing snapshot	<ol style="list-style-type: none"> a. Under Snapshot to clone, select Use an existing snapshot. b. Select the snapshot you want to use for the clone.
Create a new host group	<ol style="list-style-type: none"> a. Under Host mapping, select New host group. b. Enter a name for the new host group; then select the host initiators to include in the group.
Add a new host	<ol style="list-style-type: none"> a. Under Host mapping, select New hosts. b. Enter the a name for the new host; then select FC or iSCSI. c. Select the host initiators from the list of existing initiators or select Add to add new initiators for the host.

7. Select **Clone**.

What's next?

You have created a new storage unit that is identical to the storage unit you cloned. You are now ready to use the new storage unit as needed.

Clone consistency groups

When you clone a consistency group, you create a new consistency group that's identical in structure, storage units, and data to the consistency group you cloned. Use a consistency group clone to perform application testing or to migrate data. Suppose, for example, you need to migrate a production workload out of a consistency group. You can clone the consistency group to create a copy of your production workload to maintain as a backup until the migration is complete.

The clone is created from a snapshot of the consistency group being cloned. The snapshot used for the clone is taken at the point in time that the cloning process is initiated by default. You can modify the default behavior to use a pre-existing snapshot.

Storage unit mappings are copied as part of the cloning process. Snapshot polices are not copied as part of the cloning process.

You can create clones from consistency groups stored locally on your ASA r2 system or from consistency groups that have been replicated to remote locations.

Clone using local snapshot

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the consistency group you want to clone.
3. Select , then select **Clone**.
4. Enter a name for consistency group clone or accept the default name.
5. Select the host operating system.
6. If you want to dissociate the clone from the source consistency group and allocate disk space, select **Split clone**.
7. If you want to use an existing snapshot, create a new host group or add a new host for the clone, select **More Options**.

Option	Steps
Use an existing snapshot	<ol style="list-style-type: none">a. Under Snapshot to clone, select Use an existing snapshot.b. Select the snapshot you want to use for the clone.
Create a new host group	<ol style="list-style-type: none">a. Under Host mapping, select New host group.b. Enter a name for the new host group; then select the host initiators to include in the group.
Add a new host	<ol style="list-style-type: none">a. Under Host mapping, select New hosts.b. Enter the name new host name; then select FC or iSCSI.c. Select the host initiators from the list of existing initiators or select Add initiator to add new initiators for the host.

8. Select **Clone**.

Clone using remote snapshot

Steps

1. In System Manager, select **Protection > Replication**.
2. Hover over the **Source** you want to clone.
3. Select , then select **Clone**.
4. Select the source cluster and storage VM; then enter a name for the new consistency group or accept the default name.
5. Select the snapshot to clone; then select **Clone**.

What's next?

You have cloned a consistency group from your remote location. The new consistency group is locally available on your ASA r2 system to use as needed.

What's next?

To protect your data, you should [create snapshots](#) of the cloned consistency group.

Split consistency group clone

When you split a consistency group clone, you dissociate the clone from the source consistency group and allocate disk space for the clone. The clone becomes a standalone consistency group that can be used independently of the source consistency group.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the consistency group clone that you want you want to split.
3. Select **Split clone**.
4. Select **Split**.

Result

The clone is dissociated from the source consistency group and disk space is allocated for the clone.

Manage host groups

Create host groups on your ASA r2 system

On an ASA r2 system, a *host group* is the mechanism used to give hosts access to storage units. A host group refers to an igroup for SCSI hosts or to an NVMe subsystem for NVMe hosts. A host can only see the storage units that are mapped to the host groups to which it belongs. When a host group is mapped to a storage unit, the hosts that are members of the group, are then able to mount (create directories and file structures on) the storage unit.

Host groups are automatically or manually created when you create your storage units. You can optionally use the following steps to create host groups before or after storage unit creation.

Steps

1. From System Manager, select **Host**.
2. Select the hosts you want to add to the host group.

After you select the first host, the option to add to a host group appears above the list of hosts.

3. Select **Add to host group**.
4. Search for and select the host group to which you want to add the host.

What's next?

You have created a host group and you can now [map it to a storage unit](#).

Delete a host group on your ASA r2 system

On an ASA r2 system, a host group is the mechanism used to give hosts access to storage units. A host group refers to an igroup for SCSI hosts or to an NVMe subsystem for NVMe hosts. A host can only see the storage units that are mapped to the host groups to which it belongs. You might want to delete a host group if you no longer want the hosts in the group to have access to the storage units that are mapped to the group.

Steps

1. In System Manager, select **Storage**.
2. Under **Host mapping** select the host group you want to delete.
3. Select **Mapped storage**.
4. Select **More**; then select **Delete**.
5. Select to verify that you want to continue; then select **Delete**.

What's next?

The host group is deleted. The hosts that were in the group no longer have access to the storage units that were mapped to the host group.

Manage storage units

Modify storage units on ASA r2 storage systems

To optimize performance on your ASA r2 system, you might need to modify your storage units to increase their capacity, update QoS policies or to change the hosts that are mapped to the units. For example, if a new, critical application workload is added to an existing storage unit, you might need to change the Quality of Service (QoS) policy applied to the storage unit to support the performance level needed for the new application.

Increase capacity

Increase the size of a storage unit before it reaches full capacity to prevent a loss of data access that can occur if the storage unit runs out of writeable space. The capacity of a storage unit can be increased to 128 TB which is the maximum size allowed by ONTAP.

Modify host mappings

Modify the hosts that are mapped to a storage unit to assist in balancing workloads or reconfiguring system resources.

Modify QoS policy

Quality of service (QoS) policies guarantee that the performance of critical workloads is not degraded by competing workloads. You can use QoS policies to set a QoS throughput *limit* and a QoS throughput *guarantee*.

- QoS throughput limit

The QoS throughput *limit* restricts the impact of a workload on system resources by limiting the throughput for the workload to a maximum number of IOPS or MBps, or IOPS and MBps.

- QoS throughput guarantee

The QoS throughput *guarantee* ensures that critical workloads meet minimum throughput targets, regardless of demand by competing workloads, by guaranteeing that the throughput for the critical workload does not fall below a minimum number of IOPS or MBps, or IOPS and MBps.

Steps

1. In System Manager, select **Storage**.
2. Hover over the name of the storage unit you want to edit.
3. Select ; then select **Edit**.
4. Update the storage unit parameters as needed to increase capacity, change the QoS policy, and update the host mapping.

What's next?

If you have increased the size of your storage unit, you must rescan the storage unit on the host for the host to recognize the change in size.

Move storage units on ASA r2 storage systems

If a storage availability zone is running low on space, you can move storage units to another storage availability zone to balance the storage utilization across the cluster.

You can move a storage unit while the storage unit is online and serving data. The move operation is non-disruptive.

Before you begin

- You must be running ONTAP 9.16.1 or later.
- Your cluster must consist of four or more nodes.

Steps

1. In System Manager, select **Storage**; then select the storage unit you want to move.
2. Select ; then select **Move**.
3. Select the storage availability zone you want to move the storage unit to; then select **Move**.

Delete storage units on ASA r2 storage systems

Delete a storage unit if you no longer need to maintain the data contained in the unit. Deleting storage units that are no longer needed can help you free space needed for other host applications.

Before you begin

If the storage unit you want to delete is in a consistency group that is in replication relationship, you must [remove the storage unit from the consistency group](#) before you delete it.

Steps

1. In System Manager, select **Storage**.
2. Hover over the name of the storage unit you want to delete.
3. Select ; then select **Delete**.

4. Acknowledge that the deletion cannot be undone.

5. Select **Delete**.

What's next?

You can use the space freed from the deleted storage unit to [increase the size](#) of storage units that need additional capacity.

Migrate storage VMs

Migrate a storage VM from an ASA cluster to an ASA r2 cluster

Beginning with ONTAP 9.18.1, you can non-disruptively migrate a storage virtual machine (VM) from any ASA cluster to any ASA r2 cluster. Migrating from an ASA cluster to an ASA r2 cluster allows you to adopt the simplified and streamlined architecture of ASA r2 systems for SAN-only environments.

Storage VM migration between ASA and ASA r2 storage systems is supported as follows:

From any of the following ASA systems:	To any of the following ASA r2 systems:
<ul style="list-style-type: none">• ASA C800• ASA C400• ASA C250• ASA A900• ASA A800• ASA A400• ASA A250• ASA A150• ASA AFF A800• ASA AFF A700• ASA AFF A400• ASA AFF A250• ASA AFF A220	<ul style="list-style-type: none">• ASA A1K• ASA C30• ASA A90• ASA A70• ASA A50• ASA A30• ASA A20



For the most current list of ASA and ASA r2 systems, see [NetApp Hardware Universe](#). ASA r2 systems are listed in NetApp Hardware Universe as "ASA A-Series/C-Series (New)".

You can migrate a storage VM to an ASA r2 cluster from an ASA cluster only. Migration from any other type of ONTAP system is not supported.

Before you begin

All nodes in the ASA r2 cluster and the ASA cluster must be running ONTAP 9.18.1 or later. The ONTAP 9.18.1 patch versions on the cluster nodes can vary.

Step 1: Verify the status of the ASA storage VM

Before you migrate a storage VM from an ASA system, there should be no NVMe namespaces or vVols present and each volume in the storage VM should contain only one LUN. Migration of NVMe namespaces and vVols is not supported. The architecture of ASA r2 systems requires that volumes contain a single LUN.

Steps

1. Verify that no NVMe namespaces are present in the storage VM:

```
vserver nvme namespace show -vserver <storage_VM>
```

If entries are displayed, the NVMe objects must be [converted](#) to LUNs or removed. See the `vserver nvme namespace delete` and the `vserver nvme subsystem delete` commands in the [ONTAP command reference](#) for more information.

2. Verify that there are no vVols present in the storage VM:

```
lun show -vserver <storage_VM> -class protocol-endpoint,vvol
```

If any vVols are present, they should be copied to another storage VM and then deleted from the storage VM to be migrated. See the `lun copy` and `lun delete` commands in the [ONTAP command reference](#) for more information.

3. Verify that each volume in the storage VM contains a single LUN:

```
lun show -vserver <storage_VM>
```

If a volume contains more than one LUN, use the `volume create` and `lun move` commands to create a 1:1 volume-to-LUN ratio. See the [ONTAP command reference](#) for more information.

What's next?

You are ready to create a cluster peer relationship between your ASA and ASA r2 clusters.

Step 2: Create a cluster peer relationship between your ASA and ASA r2 clusters

Before you can migrate a storage VM from an ASA cluster to an ASA r2 cluster, you need to create a peer relationship. A peer relationship defines network connections that enable ONTAP clusters and storage VMs to exchange data securely.

Before you begin

You must have created intercluster LIFs on every node in the clusters being peered using one of the following methods.

- [Configure intercluster LIFs on shared data ports](#)
- [Configure intercluster LIFs on dedicated data ports](#)
- [Configure intercluster LIFs in custom IPspaces](#)

Steps

1. On the ASA r2 cluster, create a peer relationship with the ASA cluster and generate a passphrase:

```
cluster peer create -peer-addrs <ASA_cluster_LIF_IPs> -generate  
-passphrase
```

The following example creates a cluster peer relationship between cluster 1 and cluster 2 and creates a system-generated passphrase:

```
cluster1::> cluster peer create -peer-addrs 10.98.191.193 -generate  
-passphrase  
Passphrase: UCa+6lRVICXeL/gqlWrK7ShR  
Peer Cluster Name: cluster2  
Initial Allowed Vserver Peers: -  
Expiration Time: 6/7/2017 09:16:10 +5:30  
Intercluster LIF IP: 10.140.106.185  
Warning: make a note of the passphrase - it cannot be displayed again.
```

2. Copy the generated passphrase.

3. On the ASA cluster, create a peer relationship with the ASA r2 cluster:

```
cluster peer create -peer-addrs <ASA_r2_LIF_IPs>
```

4. Enter the passphrase generated on the ASA r2 cluster.

5. Verify that the cluster peer relationship is created:

```
cluster peer show
```

The following example displays the expected output for successfully peered clusters.

```
cluster1::> cluster peer show  
  
Peer Cluster Name      Cluster Serial Number      Availability  
Authentication  
-----  -----  -----  
-----  
cluster2                1-80-123456          Available      ok
```

Result

The ASA and ASA r2 clusters are peered and storage VM data can be securely transferred.

What's next?

You are ready to prepare your ASA storage VM for migration.

Step 3: Prepare for storage VM migration from an ASA to an ASA r2 cluster

Before you migrate a storage virtual machine (VM) from an ASA cluster to an ASA r2 cluster, you must run a migration pre-check and fix any required issues. You cannot perform the migration until the pre-check passes successfully.

Step

1. From your ASA r2 cluster, execute the migration pre-check:

```
vserver migrate start -vserver <storage_VM> -source-cluster  
<asa_cluster> -check-only true
```

If you need to fix any issues to prepare your ASA cluster for migration, the issue and the corrective action is displayed. Fix the issue and repeat the pre-check until it completes successfully.

What's next?

You are ready to migrate your storage VM from your ASA cluster to an ASA r2 cluster.

Step 4: Migrate an ASA storage VM to an ASA r2 cluster

After you have prepared your ASA cluster and created the necessary cluster peer relationship with the ASA r2 cluster, you can begin the storage VM migration.

When performing a storage VM migration, it is a best practice to leave 30% CPU headroom on both the ASA cluster and the ASA r2 cluster to enable the CPU workload to execute.

About this task

After the storage VM migration, clients are automatically cut over to the ASA r2 cluster and the storage VM on the ASA cluster is automatically removed. Automatic cutover and automatic storage VM removal are enabled by default. You can optionally disable them both and perform the cutover and the storage VM removal manually.

Before you begin

- The ASA r2 cluster must have enough free space to hold the migrated storage VM.
- If the ASA storage VM contains encrypted volumes, the onboard key manager or the external key manager on the ASA r2 system must be configured at the cluster level.
- The following operations cannot be running on the source ASA cluster:
 - Failover operations
 - WAFLIRON
 - Fingerprint
 - Volume move, rehost, clone, create, convert or analytics

Steps

1. From the ASA r2 cluster, start the storage VM migration:

```
vserver migrate start -vserver <storage_VM_name> -source-cluster  
<ASA_cluster>
```

To disable automatic cutover, use the `-auto-cutover false` parameter. To disable the automatic removal of the ASA storage VM, use the `-auto-source-cleanup false` parameter.

2. Monitor the status of the migration

```
vserver migrate show -vserver <storage_VM_name>
```

When the migration is complete, the **status** displays as **migration-complete**.

 If you need to pause or cancel the migration before automatic cutover begins, use the `vserver migrate pause` and the `vserver migrate abort` commands. You must pause the migration before cancelling it. You cannot cancel the migration after cutover starts.

Result

The storage VM is migrated from the ASA cluster to the ASA r2 cluster. The storage VM's name and UUID, the data LIF name, IP address, and object names, such as the volume name, remain unchanged. The UUID of the migrated objects in the storage VM are updated.

What's next?

If you disabled automatic cutover and automatic storage VM removal, [manually cut over your ASA clients to your ASA r2 cluster and remove the storage VM from the ASA cluster](#).

Cutover clients and clean up source storage VM after migration to an ASA r2 system

After a storage virtual machine (VM) is migrated from an ASA cluster to an ASA r2 cluster, by default, clients are automatically cut over to the ASA r2 cluster and the storage VM on the ASA cluster is automatically removed. If you chose to disable automatic cutover and removal of the ASA storage VM during the migration, you need to perform these steps manually after the migration is complete.

Manually cut over clients to an ASA r2 system after a storage VM migration

If you disable automatic client cutover during the migration of a storage VM from an ASA cluster to an ASA r2 cluster, after the migration is successfully complete, perform the cutover manually so the ASA r2 storage VM can serve data to clients.

Steps

1. On the ASA r2 cluster, manually execute client cutover:

```
vserver migrate cutover -vserver <storage_VM_name>
```

2. Verify that the cutover operation is complete:

```
vserver migrate show
```

Result

Data is being served to your clients from the storage VM on your ASA r2 cluster.

What's next?

You are now ready to remove the storage VM from the source ASA cluster.

Manually remove an ASA storage VM after migration to an ASA r2 cluster

If you disable automatic source cleanup during the migration of a storage VM from an ASA cluster to an ASA r2 cluster, after the migration is complete, remove the storage VM from the ASA cluster to free the storage space.

Before you begin

Your clients should be serving data from the ASA r2 cluster.

Steps

1. From the ASA cluster, verify that the status of the ASA storage VM is **Ready for source cleanup**:

```
vserver migrate show
```

2. Remove the ASA storage VM:

```
vserver migrate source-cleanup -vserver <storage_VM_name>
```

Result

The storage VM on your ASA cluster is removed.

ASA r2 storage limits

For optimal performance, configuration and support, you should be aware of ASA r2 storage limits.

For a complete list of the most current ASA r2 storage limits, see [NetApp Hardware Universe](#).

ASA r2 systems support the following storage limits:

	Maximum per HA pair	Maximum per cluster
Consistency groups	256	256
Enterprise applications	100	350
Nodes	2	12
Replication groups	50	50
Storage availability zone size	2 PB	2 PB

	Maximum per HA pair	Maximum per cluster
Storage units	10,000	30,000
Storage unit size	128 TB	128 TB
Storage units per consistency group	256	256
Child consistency groups per parent consistency group	64	64
Storage virtual machines	<ul style="list-style-type: none"> • 256 (ONTAP 9.18.1 and later) • 32 (ONTAP 9.17.1 and earlier) 	<ul style="list-style-type: none"> • 256 (ONTAP 9.18.1 and later) • 32 (ONTAP 9.17.1 and earlier)
Virtual machines	800	1200

Limits for SnapMirror asynchronous relationships

The following limits apply to storage units and consistency groups in a SnapMirror asynchronous replication relationship. For a complete list of the most current ASA r2 storage limits, [NetApp Hardware Universe](#).

Limit maximum	Per HA pair	Per cluster
Consistency groups	250	750
Storage units	4,000	6,000

Limits for SnapMirror active sync relationship

The following limits apply to storage units and consistency groups in a SnapMirror active sync replication relationship. SnapMirror active sync is supported beginning with ONTAP 9.17.1 on two-node clusters only. Beginning with ONTAP 9.18.1, SnapMirror active sync is supported on four-node clusters.

For a complete list of the most current ASA r2 storage limits, [NetApp Hardware Universe](#).

Limit maximum	Per HA pair
Consistency groups	50
Storage units	400

Protect your data

Create snapshots to back up your data on ASA r2 storage systems

Create a snapshot to back up data on your ASA r2 system. Use ONTAP System Manager to create a manual snapshot of a single storage unit, or to create a consistency group and schedule automatic snapshots of multiple storage units at the same time.

Step 1: Optionally, create a consistency group

A consistency group is a collection of storage units that are managed as a single unit. Create consistency groups to simplify storage management and data protection for application workloads spanning multiple

storage units. For example, suppose you have a database consisting of 10 storage units in a consistency group, and you need to back up the entire database. Instead of backing up each storage unit, you can back up the entire database by simply adding snapshot data protection to the consistency group.

Create a consistency group using new storage units or create a consistency group using existing storage units.

Beginning with ONTAP 9.18.1, you can set the snapshot reserve percentage and enable automatic snapshot deletion when creating a consistency group with new storage units. The snapshot reserve is the amount of space in the storage unit reserved specifically for snapshots. When snapshot reserve is set with automatic snapshot deletion, older snapshots are automatically deleted when space used by snapshots exceeds the snapshot reserve. If snapshot reserve and automatic snapshot deletion is enabled on a parent consistency group, it is enabled on all existing child consistency groups. If new child consistency groups are added they do not inherit the snapshot reserve and snapshot deletion settings of the parent.

[Learn more about snapshot reserve on ASA r2 storage systems.](#)

Beginning with ONTAP 9.16.1, you when you create consistency groups using new storage units, you can configure up to five child consistency group.

[Learn more about child consistency groups on ASA r2 systems.](#)

Use new storage units

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Select **+ Add**; then select **Using new storage units**.
3. Enter a name for the new storage unit, the number of units, and the capacity per unit.

If you create more than one unit, each unit is created with the same capacity and the same host operating system by default. You can optionally assign a different capacity to each unit.

4. If you want to do any of the following, select **More Options** and complete the required steps.

Option	Steps
Assign a different capacity to each storage unit	Select Add a different capacity .
Change the default performance service level	Under Performance service level , select a different service level. ASA r2 systems offer two performance levels. The default performance level is Extreme , which is the highest available level. You can lower the performance level to Performance .
Modify the default snapshot reserve and enable automatic snapshot deletion	<ol style="list-style-type: none">a. Under Snapshot reserve %, enter the numeric value for the percentage of the storage unit's space that you want allocated to snapshots.b. Select Automatically delete older snapshots.
Create a child consistency group	Select Add child consistency group .

5. Select the host operating system and host mapping.

6. Select **Add**.

What's next?

You have created a consistency group containing the storage units you want to protect. Now you can create a snapshot.

Use existing storage units

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Select **+ Add**; then select **Using existing storage units**.
3. Enter a name for the consistency group; then search for and select the storage units you want to include in the consistency group.
4. Select **Add**.

What's next?

You have created a consistency group containing the storage units you want to protect. Now you can

create a snapshot.

Step 2: Create a snapshot

A snapshot is a local, read-only copy of your data that you can use to restore storage units to specific points in time.

Snapshots can be created on demand, or they can be created automatically in regular intervals based on a [snapshot policy and schedule](#). The snapshot policy and schedule specifies 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 storage unit or of a consistency group. You can create automated snapshots of a consistency group containing multiple storage units. You cannot create automated snapshots of a single storage unit.

- On-demand snapshots

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

- Automated snapshots

Automated snapshots are created using snapshot policies. To apply a snapshot policy to a storage unit for automated snapshot creation, the storage unit must be a member of a consistency group. If you apply a snapshot policy to a consistency group, all the storage units in the consistency group are protected with automated snapshots.

Create a snapshot of a consistency group or a storage unit.

Snapshot of a consistency group

Steps

1. In System Manager, select **Protection > 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 storage unit.

- 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  ; 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**.

Snapshot of storage unit

Steps

1. In System Manager, select **Storage**.
2. Hover over the name of the storage unit 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 storage unit.

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 <input checked="" type="checkbox"/> 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**.

What's next?

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

Manage snapshot reserve

Learn about ONTAP snapshot reserve on ASA r2 storage

The snapshot reserve is the amount of space in the storage unit reserved specifically for snapshots. When snapshot reserve is set with automatic snapshot deletion, older snapshots are automatically deleted when space used by snapshots exceeds the snapshot reserve. This prevents snapshots from consuming space in your storage unit intended for user data.

Snapshot reserve is set as a percentage of the total storage unit size. For example, if the storage unit is 50 GB and you set the snapshot reserve to 10%, the amount of space reserved for snapshots is 5 GB. When the amount of space used by snapshots grows to 5 GB, older snapshots are automatically deleted to make room for new snapshots. If the storage unit size increases to 100 GB, then the snapshot reserve increases to 10 GB. The maximum snapshot reserve you can set is 200%. If your storage unit grows to the maximum size of 128 TB, a 200% snapshot reserve allows you to take 2 complete snapshots.

By default, snapshot reserve is set to 0% and snapshot auto-delete is not enabled.

Beginning with ONTAP 9.18.1, you can modify the default snapshot reserve during or after the creation of storage units and during the creation of consistency groups. You can also modify the default snapshot reserve on existing storage virtual machines (VMs). In ONTAP 9.17.1 and earlier, you cannot modify these settings.

Snapshot reserve is set to the same percentage for all storage units in a consistency group at the time the consistency group is created. Snapshot reserve must be individually set on any storage units added later.

Modify snapshot reserve on an ASA r2 storage system

The snapshot reserve is the amount of space in the storage unit reserved specifically for snapshots. By default, snapshot reserve is set to 0%. Beginning with ONTAP 9.18.1, you can modify the storage unit's default snapshot reserve and enable automatic snapshot deletion. Automatic deletion of snapshots is disabled by default. When a snapshot reserve value is set and automatic snapshot deletion is enabled, older snapshots are automatically deleted when space used by snapshots exceeds the snapshot reserve. This prevents snapshots from consuming space in your storage unit intended for user data.

[Learn more about snapshot reserve on ASA r2 storage systems.](#)

Modify snapshot reserve on storage units

To set different snapshot reserve values, configure each storage unit individually. To use the same value for all storage units, modify the snapshot reserve on the storage VM.

Steps

1. In System Manager, select **Storage**.
2. Hover over the name of the storage unit for which you want to set the snapshot reserve.
3. Select , then select **Edit**.
4. Under **Snapshot reserve %**, enter the numeric value for the percentage of the storage unit's space that you want allocated to snapshots.
5. Verify that **Automatically delete older snapshots** is selected.
6. Select **Save**.

Result

The snapshot reserve is set to the percentage you specified. If the amount of space consumed by snapshots reaches the reserve, older snapshots are automatically deleted.

Modify snapshot reserve on a storage VM

To set the same snapshot reserve for all storage units in a storage VM, apply the desired percentage to the storage VM. When snapshot reserve is applied to the storage VM, it is applied to all newly created storage units within the storage VM. It is not applied to storage units created before you modified the setting.

Steps

1. In System Manager, select **Cluster > Storage VMs**; then select **Settings**.
2. Under **Policies**, next to **Snapshots**, select ; then select **Set/edit snapshot reserve default**.
3. Under **Snapshot reserve %**, enter the numeric value for the percentage of the storage unit's space that you want allocated to snapshots.
4. Verify that **Automatically delete older snapshots** is selected.
5. Select **Save**.

Result

The snapshot reserve for newly created storage units is set to the percentage you specified. If the amount of space consumed by snapshots in those storage units reaches the reserve, older snapshots are automatically deleted.

Create an intercluster storage VM peer relationship on ASA r2 storage systems

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

[Learn more about peer relationships.](#)

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**.

Set up snapshot replication

Replicate snapshots to a remote cluster from ASA r2 storage systems

Snapshot replication is a process in which consistency groups on your ASA r2 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.



Snapshot replication for an ASA r2 storage system is only supported to and from another ASA r2 storage system. You cannot replicate snapshots from an ASA r2 system to an ASA, AFF or FAS system or from an ASA, AFF or FAS system to an ASA r2 system.

To set up Snapshot replication, you need to establish a replication relationship between your ASA r2 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.

Beginning with ONTAP 9.17.1, you can apply asynchronous replication policies to consistency groups in a hierarchical relationship. Asynchronous replication is not supported for consistency groups in hierarchical relationships in ONTAP 9.16.1.

[Learn more about hierarchical \(parent/child\) consistency groups.](#)

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 ASA r2 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 ASA r2 cluster with a remote cluster. You can now create a replication relationship.

Step 2: Optionally, create a custom replication policy

The replication policy defines when updates performed on the ASA r2 cluster are replicated to the remote site. ONTAP include various pre-defined data protection policies that you can use for your replication relationships. If the pre-defined policies do not meet your needs, you can create a custom replication policy.

Learn about [pre-defined ONTAP data protection policies](#).

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 storage units in a specific storage VM, select **Storage VM**.

5. For the **Policy type**, select **Asynchronous**.



With the asynchronous policy, data is copied to the remote site after it is written to the source. Synchronous replication is not supported for ASA r2 systems.

6. Under **Transfer snapshots from source**, accept the default transfer schedule or select a different one.
7. Select to transfer all snapshots or to create rules to determine which snapshots to transfer.
8. Optionally, enable network compression.
9. Select **Save**.

What's next?

You have created a replication policy and are now ready to create a replication relationship between your ASA r2 system and your remote location.

For more information

Learn more about [storage VMs for client access](#).

Step 3: Create a replication relationship

A snapshot replication relationship establishes a connection between your ASA r2 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 storage unit is compromised by a ransomware attack.

Before you begin

- [Learn about replication policies](#).

When you create a replication relationship, you must select the appropriate replication policy for your replication relationship. You can use a pre-defined policy or create a custom policy.

- If you want to lock your destination snapshots, you must [initialize the Snapshot compliance clock](#) before you create the replication relationship.

Create a replication relationship with or without locked destination snapshots.

With locked 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.

10. 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.

11. Select **Save**.

Without locked snapshots

Steps

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

Option	Steps
Local destinations	<ol style="list-style-type: none">1. Select Local destinations, then select  Replicate.2. Search for and select the source consistency group. <p>The <i>source</i> consistency group refers to the consistency group on your local cluster that you want to replicate.</p>

Option	Steps
Local sources	<ol style="list-style-type: none"> 1. Select Local sources, then select  . 2. Search for and select the source consistency group. 3. Under Replication destination, select the cluster to replicate to; then select the storage VM. <p>3. Select a replication policy.</p> <p>4. To delay the start of the data transfer, select Destination settings; then deselect Start transfer immediately.</p> <p>The initial data transfer begins immediately by default.</p> <p>5. Optionally, to override the default transfer schedule, select Destination settings, then select Override transfer schedule.</p> <p>Your transfer schedule must be a minimum of 30 minutes to be supported.</p> <p>6. Select Save.</p>

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 ASA r2 system goes offline.

Step 4: Test replication failover

Optionally, validate that you can successfully serve data from replicated storage units 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 ASA r2 system is repurposed, returned, misplaced or stolen.

Learn about pre-defined ONTAP data protection policies

The replication policy defines when updates performed on the ASA r2 cluster are replicated to the remote site. ONTAP includes various pre-defined data protection policies that you can use for your replication relationships.

If the pre-defined policies do not meet your needs, you can [create a custom replication policy](#).



ASA r2 systems do not support synchronous replication.

ASA r2 systems support the following pre-defined protection policies.

Policy	Description	Policy type
Asynchronous	A unified SnapMirror asynchronous and vault policy for mirroring the latest active file system and daily and weekly snapshots with an hourly transfer schedule.	Asynchronous
AutomatedFailOverDuplex	Policy for SnapMirror synchronous with zero RTO guarantee and bi-directional sync replication.	SnapMirror active sync
CloudBackupDefault	Vault policy with daily rule.	Asynchronous
DailyBackup	Vault policy with a daily rule and a daily transfer schedule.	Asynchronous
DPDefault	SnapMirror asynchronous policy for mirroring all snapshots and the latest active file system.	Asynchronous
MirrorAllSnapshots	SnapMirror asynchronous policy for mirroring all snapshots and the latest active file system.	Asynchronous
MirrorAllSnapshotsDiscardNetwork	SnapMirror asynchronous policy for mirroring all snapshots and the latest active file system excluding the network configurations.	Asynchronous
MirrorAndVault	A unified SnapMirror asynchronous and vault policy for mirroring the latest active file system and daily and weekly snapshots.	Asynchronous
MirrorAndVaultDiscardNetwork	A unified SnapMirror asynchronous and vault policy for mirroring the latest active file system and daily and weekly snapshots excluding the network configurations.	Asynchronous
MirrorLatest	SnapMirror asynchronous policy for mirroring the latest active file system.	Asynchronous
Unified7year	Unified SnapMirror policy with 7-year retention.	Asynchronous
XDPDefault	Vault policy with daily and weekly rules.	Asynchronous

Break an asynchronous replication relationship on your ASA r2 system

In certain situations, you might need to break an asynchronous replication relationship. For example, if you are running ONTAP 9.16.1 and you want to increase the size of a consistency group that is in an asynchronous replication relationship, you must break the relationship before you can modify the consistency group's size.

Steps

1. In System Manager, select **Protection > Replication**.
2. Select **Local destinations** or **Local sources**.
3. Next to the relationship that you want to break, select ; then select **Break**.
4. Select **Break**.

Result

The asynchronous relationship between the primary and secondary consistency group is broken.

Set up SnapMirror active sync

SnapMirror active sync setup workflow

ONTAP SnapMirror active sync data protection enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. There is no manual intervention or custom scripting required to trigger a failover with SnapMirror active sync.

While the System Manager procedures for configuring SnapMirror active sync are different on ASA r2 systems than NetApp FAS, AFF, and ASA systems running the unified ONTAP personality, the requirements, architecture and operation of SnapMirror active sync is the same.

[Learn more about ONTAP personalities.](#)



Beginning with ONTAP 9.18.1, SnapMirror active sync is supported on four-node configurations. In ONTAP 9.17.1, SnapMirror active sync is supported on two-node configurations only.

[Learn more about SnapMirror active sync.](#)

[Learn more about disaster recovery with SnapMirror active sync on your ASA r2 system](#)

On ASA r2 systems, SnapMirror active sync supports symmetric active/active configurations. In a symmetric active/active configuration, both sites can access local storage for active I/O.

[Learn more about symmetric active/active configurations.](#)

1

Prepare to configure SnapMirror active sync.

To [prepare to configure SnapMirror active sync](#) on your ASA r2 system you should review the configuration prerequisites, confirm support for your host operating systems, and be aware of object limits that might impact specific configuration.

2

Confirm your cluster configuration.

Before you configure SnapMirror active sync, you should [confirm that your ASA r2 clusters are in the proper peering relationships and meet other configuration requirements](#).

3

Install ONTAP Mediator.

You can use ONTAP Mediator or ONTAP Cloud Mediator to monitor the health of your cluster and enable

business continuity. If you are using ONTAP Mediator, you must [install it](#) on your host. If you are using ONTAP Cloud Mediator, you can skip this step.

4

Configure ONTAP Mediator or ONTAP Cloud Mediator using self-signed certificates.

You must [configure ONTAP mediator or ONTAP cloud mediator](#) before you can begin using it with SnapMirror active sync for cluster monitoring.

5

Configure SnapMirror active sync.

[Configure SnapMirror active sync](#) to create a copy of your data at a secondary site and enable your host applications to automatically and transparently fail over in the event of a disaster.

Prepare to configure SnapMirror active sync on ASA r2 systems

To prepare to configure SnapMirror active sync on your ASA r2 system you should review the configuration prerequisites, confirm support for your hosts operating systems, and be aware of object limits that might impact specific configuration.

Steps

1. Review the SnapMirror active sync [prerequisites](#).
2. [Confirm that your host operating systems are supported](#) for SnapMirror active sync.
3. Review the [object limits](#) that might impact your configuration.
4. Verify host protocol support for SnapMirror active sync on your ASA r2 system.

Support for SnapMirror active sync on ASA r2 systems varies based upon ONTAP version and host protocol.

Beginning with ONTAP...	SnapMirror active sync supports...
9.17.1	<ul style="list-style-type: none">• iSCSI• FC• NVMe/FC• NVMe/TCP
9.16.0	<ul style="list-style-type: none">• iSCSI• FC

NVMe protocol limitations with SnapMirror active sync on ASA r2 systems

Before you configure SnapMirror active sync on an ASA r2 system with NVMe hosts, you should be aware of certain NVMe protocol limitations.

All NVMe storage units in the NVMe subsystem must be members of the same consistency group and must all be part of the same SnapMirror active sync relationship.

The NVMe/FC and NVMe/TCP protocols are supported with SnapMirror active sync as follows:

- Only on 2-node clusters
- Only on ESXi hosts
- Only with symmetric active/active configurations

Asymmetric active/active configurations are not supported with NVMe hosts.

SnapMirror active sync with NVMe does not support the following:

- Subsystems mapped to more than one consistency group

A consistency group can be mapped with multiple subsystems, but each subsystem can be mapped to only one consistency group.

- Expansion of consistency groups in a SnapMirror active sync relationship
- Mapping NVMe storage units that are not in a SnapMirror active sync relationship to replicated subsystems
- Removing a storage unit from a consistency group
- Consistency group geometry change
- [Microsoft Offloaded Data Transfer \(ODX\)](#)

What's next?

After you have completed the preparation necessary to enable SnapMirror active sync, you should [confirm your cluster configuration](#).

Confirm your ASA r2 cluster configuration before configuring SnapMirror active sync

SnapMirror active sync relies on peered clusters to protect your data in the event of a failover. Before you configure SnapMirror active sync, you should confirm that your ASA r2 clusters are in a supported peering relationship and meet other configuration requirements.

Steps

1. Confirm that a cluster peering relationship exists between the clusters.



The default IPspace is required by SnapMirror active sync for cluster peer relationships. A custom IPspace isn't supported.

[Create a cluster peer relationship.](#)

2. Confirm that a peer relationship exists between the storage virtual machines (VMs) on each cluster.

[Create an intercluster storage VM peer relationship.](#)

3. Confirm that at least one LIF is created on each node in the cluster.

[Create a LIF.](#)

4. Confirm that the necessary storage units are created and mapped to host groups.

[Create a storage unit](#) and [map the storage unit to a host group](#).

5. Rescan the application host to discover any new storage units.

What's next?

After you have confirmed your cluster configuration, you are ready to [install ONTAP Mediator](#).

Install ONTAP Mediator on ASA r2 systems

To install ONTAP Mediator for your ASA r2 system, you should follow the same procedure used to install ONTAP Mediator for all other ONTAP systems.

Installing ONTAP Mediator includes preparing for installation, enabling access to repositories, downloading the ONTAP Mediator package, verifying the code signature, installing the package on the host and performing post-installation tasks.

To install ONTAP Mediator, follow [this workflow](#)

What's next

After ONTAP Mediator is installed you should [configure ONTAP Mediator using self-signed certificates](#).

Configure ONTAP Mediator or ONTAP Cloud Mediator on ASA r2 systems

You must configure ONTAP Mediator or ONTAP Cloud Mediator before you can begin using SnapMirror active sync for cluster monitoring. ONTAP Mediator and ONTAP Cloud Mediator both provide a persistent and fenced store for high availability (HA) metadata used by the ONTAP clusters in a SnapMirror active sync relationship. Additionally, both mediators provide a synchronous node health query functionality to aid in quorum determination and serves as a ping proxy for controller liveness detection.

Before you begin

If you are using ONTAP Cloud Mediator, verify that your ASA r2 system meets the necessary [prerequisites](#).

Steps

1. In System Manager, select **Protection > Overview**.
2. In the right pane under **Mediators**, select **Add a mediator**.
3. Select the **Mediator type**.
4. For a **Cloud** mediator enter the organization ID, client ID and client secret. For an **On-premises** mediator enter the IP address, port, mediator user name and mediator password.
5. Select the cluster peer from the list of eligible cluster peers or select **Add a cluster peer** to add a new one.
6. Add the certificate information
 - If you are using a self signed certificate, copy the content of the `intermediate.crt` file and paste it into the **Certificate** field, or select **Import** to navigate to the `intermediate.crt` file and import the certificate information.
 - If you are using a third-party certificate, enter the certificate information into the **Certificate** field.
7. Select **Add**.

What's next?

After you have initialized the mediator, you can [configure SnapMirror active sync](#) to create a copy of your data at a secondary site and enable your host applications to automatically and transparently failover in the event of

a disaster.

Configure SnapMirror active sync on ASA r2 systems

Configure SnapMirror active sync to create a copy of your data at a secondary site and enable your host applications to automatically and transparently failover in the event of a disaster.

On ASA r2 systems, SnapMirror active sync supports symmetric active/active configurations. In a symmetric active/active configuration, both sites can access local storage for active I/O.



If you are using the iSCSI or FC protocol and use ONTAP tools for VMware Sphere, you can optionally [use ONTAP Tools for VM ware to configure SnapMirror active sync](#).

Before you begin

[Create a consistency group](#) on the primary site with new storage units. If you want to create a non-uniform symmetric active/active configuration, also create a consistency group on the secondary site with new storage units.

Learn more about [non-uniform](#) symmetric active/active configurations.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the name of the consistency group you want to protect with SnapMirror active sync.
3. Select and then select **Protect**.
4. Under **Remote protection**, select **Replicate to a remote cluster**.
5. Select an existing cluster peer or choose to **Add a new one**.
6. Select the storage VM.
7. For the replication policy, select **AutomatedFailOverDuplex**.
8. If you are creating a non-uniform symmetric active/active configuration, select **Destination settings**; then input the name of the new destination consistency group you create before beginning this procedure.
9. Select **Save**.

Result

SnapMirror active sync is configured to protect your data so that you can continue operations with near zero recovery point objective (RPO) and near zero recovery time objective (RTO) in the event of a disaster.

Manage SnapMirror active sync

Reconfigure ONTAP Mediator or ONTAP Cloud Mediator to use a third-party certificate on ASA r2 systems

If you configure ONTAP Mediator or ONTAP Cloud Mediator with a self-signed certificate, you can reconfigure the mediator to use a third-party certificate. Third party certificates might be preferred or required by your organization for security reasons.

Step 1: Remove the mediator configuration

To reconfigure the mediator, you must first remove its current configuration from the cluster.

Steps

1. In System Manager, select **Protection > Overview**.
2. In the right pane, under **Mediators**, select  next to the cluster peer with the mediator configuration that you want to remove; then select **Remove**.
If you have multiple mediators installed, and you want to remove all configurations, select  next to **Mediators**; then select **Remove**.
3. Select **Remove** to confirm that you want to remove the mediator configuration.

Step 2: Remove the self-signed certificate

After the mediator configuration is removed, you should remove the associated self-signed certificate from the cluster.

Steps

1. Select **Cluster > Settings**.
2. Under **Security**, select **Certificates**.
3. Select the certificate that you want to remove.
4. Select ; then select **Delete**.

Step 3: Reinstall the mediator with a third-party certificate

After you have removed the associated self-signed certificate, you can reconfigure the mediator with the third-party certificate.

Steps

1. Select **Protection > Overview**.
2. In the right pane, under **Mediators**, select **Add a mediator**.
3. Select the **Mediator type**.
4. For a **Cloud** mediator enter the organization ID, client ID and client secret. For an **On-premises** mediator enter the IP address, port, mediator user name, and mediator password.
5. Select a cluster peer from the list of eligible cluster peers or select **Add a cluster peer** to add a new one.
6. Under **Certificate**, enter the third-party certificate information.
7. Select **Add**.

Result

The ONTAP Mediator or ONTAP Cloud Mediator is reconfigured to use the third-party certificate. You can now use the mediator to manage SnapMirror active sync relationships.

Perform a planned failover of ASA r2 clusters in a SnapMirror active sync relationship

SnapMirror active sync offers continuous availability for business-critical applications by creating a copy of your data at a secondary site and enabling your host applications to automatically and transparently fail over in the event of a disaster. You might need to

perform a planned failover of your SnapMirror active sync relationship to test the failover process or to perform maintenance on the primary site.

Before you begin

- The SnapMirror active sync relationship must be in sync.
- You cannot initiate a planned failover when a nondisruptive operation, such as a storage unit move, is in process.
- ONTAP Mediator or ONTAP Cloud Mediator must be configured, connected, and in quorum.

Steps

1. Select **Protection > Replication**.
2. Select the SnapMirror active sync relationship you want to fail over.
3. Select ; then select **Failover**.

What's next

Use the `snapmirror failover show` command in the ONTAP command line interface (CLI) to monitor the status of the failover.

Reestablish the SnapMirror active sync relationship after an unplanned failover of your ASA r2 clusters

On ASA r2 systems, SnapMirror active sync supports symmetric active/active configurations. In a symmetric active/active configuration both sites can access local storage for active I/O. If the source cluster fails or is isolated, the mediator triggers an automatic unplanned failover (AUFO) and serves all I/O from the destination cluster until the source cluster recovers.

If you experience an AUFO of your SnapMirror active sync relationship, you should reestablish the relationship and resume operations on the original source cluster after it comes back online.

Before you begin

- The SnapMirror active sync relationship must be in sync.
- You cannot initiate a planned failover when a nondisruptive operation, such as a storage unit move, is in process.
- The ONTAP Mediator must be configured, connected, and in quorum.
- To recover lost I/O paths or update I/O path states on your hosts, you need to perform a storage/adapter rescan on the hosts after the primary storage cluster resumes operation.

Steps

1. Select **Protection > Replication**.
2. Select the SnapMirror active sync relationship you need to reestablish.
3. Wait for the relationship status to display **InSync**.
4. Select ; then select **Failover** to resume operations on the original primary cluster.

Delete a SnapMirror active sync relationship on your ASA r2 system

If you no longer require near zero RPO and RTO for a business application, you should remove SnapMirror active sync protection by deleting the associated SnapMirror active

sync relationship. If you are running ONTAP 9.16.1 on an ASA r2 system, you might also need to delete the SnapMirror active sync relationship before you can make certain geometry changes to consistency groups in a SnapMirror active sync relationship.

Step 1: Terminate host replication

If the host group from the source cluster is replicated to the destination cluster and destination consistency groups are mapped to the replicated host group, you must terminate host replication on the source cluster before you can delete the SnapMirror active sync relationship.

Steps

1. In System Manager, select **Host**.
2. Next to a host containing the host group you want to stop replicating, select ; and then select **Edit**.
3. Deselect **Replicate host configuration**, and then select **Update**.

Step 2: Delete the SnapMirror active sync relationship

To remove SnapMirror active sync protection from a consistency group, you must delete the SnapMirror active sync relationship.

Steps

1. In System Manager, select **Protection > Replication**.
2. Select **Local destinations** or **Local sources**.
3. Next to the SnapMirror active sync relationship that you want to remove, select ; then select **Delete**.
4. Select **Release the source consistency group base snapshots**.
5. Select **Delete**.

Result

The SnapMirror active sync relationship is removed and the source consistency group base snapshots are released. The storage units in the consistency group are no longer protected by SnapMirror active sync.

What's next?

[Set up snapshot replication](#) to copy the consistency group to a geographically remote location for backup and disaster recovery.

Remove ONTAP Mediator or ONTAP Cloud Mediator from your ASA r2 system

You can use only one type of mediator at a time for SnapMirror active sync on your ASA r2 system. If you choose to change your mediator type, you must remove your current instance before you install another instance.

Steps

You must use the ONTAP command line interface (CLI) to remove ONTAP Mediator or ONTAP Cloud Mediator.

ONTAP Mediator

1. Remove ONTAP Mediator:

```
snapmirror mediator remove -mediator-address <address> -peer-cluster  
<peerClusterName>
```

Example:

```
snapmirror mediator remove -mediator-address 12.345.678.90 -peer  
-cluster cluster_xyz
```

ONTAP Cloud Mediator

1. Remove ONTAP Cloud Mediator:

```
snapmirror mediator remove -peer-cluster <peerClusterName> -type cloud
```

Example:

```
snapmirror mediator remove -peer-cluster cluster_xyz -type cloud
```

Related information

- [snapmirror mediator remove](#)

Restore data on ASA r2 storage systems

Data in a consistency group or storage unit that is protected by snapshots can be restored if it is lost or corrupted.

Restore a consistency group

Restoring a consistency group replaces the data in all the storage units in the consistency group with the data from a snapshot. Changes made to the storage units after the snapshot was created are not restored..

You can restore a consistency group from a local or remote snapshot.

Restore from a local snapshot

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Double-click the consistency group containing the data you need to restore.

The consistency group details page opens.

3. Select **Snapshots**.
4. Select the snapshot you want to restore; then select .
5. Select **Restore consistency group from this snapshot**; then select **Restore**.

Restore from a remote snapshot

Steps

1. In System Manager, select **Protection > Replication**.
2. Select **Local destinations**.
3. Select the **Source** you want to restore, then select .
4. Select **Restore**.
5. Select the cluster, storage VM, and consistency group to which you want to restore data.
6. Select the snapshot you want to restore from.
7. When prompted, enter "restore"; then select **Restore**.

Result

Your consistency group is restored to the point in time of the snapshot used for restoration.

Restore a storage unit

Restoring a storage unit replaces all the data in the storage unit with the data from a snapshot. Changes made to the storage unit after the snapshot was created are not restored.

Steps

1. In System Manager, select **Storage**.
2. Double-click the storage unit containing the data you need to restore.

The storage unit details page opens.

3. Select **Snapshots**.
4. Select the snapshot you want to restore.
5. Select ; then select **Restore**.
6. Select **Use this snapshot to restore the storage unit**; then select **Restore**.

Result

Your storage unit is restored to the point in time of the snapshot used for restoration.

Manage consistency groups

Learn about ONTAP consistency groups on ASA r2 storage systems

A consistency group is a collection of storage units that are managed as a single unit. Use consistency groups for simplified storage management.

For example, suppose you have a database consisting of 10 storage units in a consistency group, and you need to back up the entire database. Instead of backing up each storage unit, you can back up the entire database by simply adding snapshot data protection to the consistency group. Backing up the storage units as a consistency group instead of individually also provides a consistent backup of all the units, while backing up units individually could potentially create inconsistencies.

Beginning with ONTAP 9.16.1, you can use System Manager to create hierarchical consistency groups on your ASA r2 system. In an hierarchical structure, one or more consistency groups are configured as children under a parent consistency group.

Hierarchical consistency groups allow you to apply individual snapshot policies to each child consistency group and to replicate the snapshots of all the child consistency groups to a remote cluster as a single unit by replicating the parent. This simplifies data protection and management for complex data structures. For example, suppose you create a parent consistency group called SVM1_app which contains two child consistency groups: SVM1app_data for application data and SVM1app_logs for application logs. Snapshots of SVM1app_data are taken every 15 minutes and snapshots of SVM1app_logs are taken every hour. The parent consistency group, SVM1_app, has a SnapMirror policy that replicates the snapshots of both SVM1app_data and SVM1app_logs to a remote cluster every 24 hours. The parent consistency group SVM1_app is managed as a single unit and the child consistency groups are managed as separate units.

Consistency groups in replication relationships

Beginning with ONTAP 9.17.1, you can make the following geometry changes to consistency groups in an asynchronous replication relationship or in a SnapMirror active sync relationship without breaking or deleting the relationship. When a geometry change occurs on the primary consistency group, the change is replicated to the secondary consistency group.

- [Modify the size of a storage unit](#) by adding or removing storage units.
- [Promote a single consistency group](#) to a parent consistency group.
- [Demote a parent consistency group](#) to a single consistency group.
- [Detach a child consistency group](#) from a parent consistency group.
- [Create a child consistency group](#) using an existing consistency group.

In ONTAP 9.16.1, you must [break the asynchronous replication relationship](#) and [delete the SnapMirror active sync relationship](#) before making geometry changes to the consistency group.

Protect consistency groups on your ASA r2 system with snapshots

Create snapshots of the consistency groups in your ASA r2 storage system to protect the data in the storage units that are part of the consistency group. If you no longer need to protect the data in any of the storage units in the consistency group, you can remove snapshot protection from the consistency group.

If you no longer need to protect the data from specific storage units in the consistency group, you can remove

those storage units from the consistency group.

Add snapshot data protection to a consistency group

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

You can use snapshots to [restore data](#) that is lost or corrupted.

Steps

1. In System Manager, select **Protection > 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.
Create a new snapshot policy	<ol style="list-style-type: none">a. Select  Add; then enter the new policy name.b. Select the policy scope.c. Under Schedules select  Add.d. Select the name that appears under Schedule name; then select .e. Select the policy schedule.f. Under Maximum snapshots, enter the maximum number of snapshots that you want to retain of the consistency group.g. Optionally, under SnapMirror label enter a SnapMirror label.h. Select Save.

6. Select **Save**.

What's next

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

Remove snapshot data protection from a consistency group

When you remove snapshot data protection from a consistency group, snapshots are disabled for all the storage units in the consistency group.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the consistency group you want to stop protecting.
3. Select ; then select **Edit**.
4. Under **Local protection**, deselect Schedule snapshots.
5. Select **Edit**.

Result

Snapshots will not be taken for any of the storage units in the consistency group.

Modify the size of consistency groups on your ASA r2 system

Increase or decrease the size of a consistency group by modifying the number of storage units in the consistency group.

Add storage units to a consistency group

Expand the amount of storage managed by a consistency group by adding new or existing storage units to the consistency group.

Beginning with ONTAP 9.18.1, you can set snapshot reserve and automatic snapshot deletion to limit the amount of space used by snapshots in your storage units. When you add a storage unit to an existing consistency group, snapshot reserve and automatic snapshot deletion are set as follows by default.

If you add...	The snapshot reserve percentage is set to...	Automatic snapshot deletion is...
New storage units	0	Disabled
Existing storage units	Unchanged	Unchanged

You can modify the default settings for new storage units when you create the storage units. You can also [modify existing storage units](#) to update their current settings.

[Learn more about snapshot reserve on ASA r2 storage systems.](#)

Before you begin

If you are running ONTAP 9.16.1 and the consistency group you want to expand is in an SnapMirror active sync relationship, you must [delete the SnapMirror active sync relationship](#) before you can add storage units. If you are running ONTAP 9.16.1 and the consistency group is in an asynchronous replication relationship, you must [break the relationship](#) before you can expand the consistency group. Deleting the SnapMirror active sync relationship or breaking the asynchronous relationship before expanding a consistency group is not required in ONTAP 9.17.1 and later releases.

Add existing storage units

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the consistency group you want to expand.
3. Select ; then select **Expand**.
4. Select **Using existing storage units**.
5. Select the storage units to add to the consistency group; then select **Expand**.

Add new storage units

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the consistency group you want to expand.
3. Select ; then select **Expand**.
4. Select **Using new storage units**.
5. Enter the number of units you want to create and the capacity per unit.

If you create more than one unit, each unit is created with the same capacity and the same host operating system. To assign a different capacity to each unit, select **Add a different capacity** to assign a different capacity to each unit.

6. Select **Expand**.

What's next

After you create a new storage unit, you should [add host initiators](#) and [map the newly created storage unit to a host](#). Adding host initiators makes hosts eligible to access the storage units and perform data operations. Mapping a storage unit to a host allows the storage unit to begin serving data to the host it is mapped to.

What's next?

Existing snapshots of the consistency group won't include your newly added storage units. You should [create an immediate snapshot](#) of your consistency group to protect your newly added storage units until the next scheduled snapshot is automatically created.

Remove a storage unit from a consistency group

Remove a storage unit from a consistency group to delete it, manage it as part of a different consistency group, or stop protecting its data. Removing a storage unit from a consistency group breaks the relationship between the storage unit and the consistency group, but does not delete the storage unit.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Double-click the consistency group from which you want to remove a storage unit.
3. In the **Overview** section, under **Storage units**, select the storage unit you want to remove; then select **Remove from consistency group**.

Result

The storage unit is no longer a member of the consistency group.

What's next

If you need to continue data protection for the storage unit, add the storage unit to another consistency group.

Delete consistency groups on your ASA r2 system

If you no longer need to manage the members of a consistency group as a single unit, you can delete the consistency group. After a consistency group is deleted, the storage units previously in the group remain active on the cluster. If the consistency group was in a replication relationship, the replicated copies remain on the remote cluster.

Before you begin

If you are running ONTAP 9.16.1, and the consistency group you want to delete is in a SnapMirror active sync relationship, you must [delete the SnapMirror active sync relationship](#) before you delete the consistency group. Deleting this relationship before modifying a consistency group is not required in ONTAP 9.17.1 and later releases.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the consistency group you want to delete.
3. Select ; then select **Delete**.
4. Accept the warning, then select **Delete**.

What's next?

After you delete a consistency group, the storage units previously in the consistency group are no longer protected by snapshots. Consider adding these storage units to another consistency group to protect them against data loss.

Manage hierarchical consistency groups on your ASA r2 system

Beginning with ONTAP 9.16.1, you can use System Manager to create hierarchical consistency groups on your ASA r2 system. In an hierarchical structure, one or more consistency groups are configured as children under a parent consistency group. You can apply individual snapshot policies to each child consistency group and replicate the snapshots of all the child consistency groups to a remote cluster as a single unit by replicating the parent. This simplifies data protection and management for complex data structures.

Promote an existing consistency group into a parent consistency group

If you promote an existing consistency group to a parent, a new child consistency group is created and the storage units belonging to the promoted consistency group are moved to the new child consistency group. Storage units cannot be directly associated with a parent consistency group.

Before you begin

If you are running ONTAP 9.16.1 and the consistency group you want to promote is in a SnapMirror active sync relationship, you must [delete the SnapMirror active sync relationship](#) before the consistency group can be promoted. If you are running ONTAP 9.16.1 and the consistency group is in an asynchronous replication relationship, you must [break the relationship](#) before you can promote the consistency group. Deleting the

SnapMirror active sync relationship or breaking the asynchronous relationship before promoting a consistency group is not required in ONTAP 9.17.1 and later releases.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the consistency group you want convert into a parent consistency group.
3. Select ; then select **Promote to parent consistency group**.
4. Enter a name for the new child consistency group or accept the default name; then select the consistency group component type.
5. Select **Promote**.

What's next?

You can create additional child consistency groups under the parent consistency group. You can also [set up snapshot replication](#) to copy the parent and child consistency groups to a geographically remote location for backup and disaster recovery.

Demote a parent consistency group to a single consistency group

When you demote a parent consistency group to a single consistency group, the storage units of the associated child consistency groups are added to the parent consistency group. The child consistency groups are deleted and the parent is then managed as a single consistency group.

Before you begin

If you are running ONTAP 9.16.1 and the consistency group you want to demote is in a SnapMirror active sync relationship, you must [delete the SnapMirror active sync relationship](#) before the consistency group can be demoted. If you are running ONTAP 9.16.1 and the consistency group is in an asynchronous replication relationship, you must [break the relationship](#) before you can demote the consistency group. Deleting the SnapMirror active sync relationship or breaking the asynchronous relationship before expanding a consistency group is not required in ONTAP 9.17.1 and later releases.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the parent consistency group you want to demote.
3. Select ; then select **Demote to a single consistency group**.
4. Select **Demote**.

What's next?

[Add a snapshot policy](#) to the demoted consistency group to protect the storage units that were previously managed by the child consistency groups.

Create a child consistency group

Creating child consistency groups allows you to apply individual snapshot policies to each child. Beginning with ONTAP 9.17.1, you can also apply individual replication policies directly to each child. In ONTAP 9.16.1, replication policies can be applied only at the parent level.

You can create a child consistency group from a new or existing consistency group.

From a new consistency group

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Hover over the parent consistency group you want to add a child consistency group to.
3. Select ; then select **Add a new child consistency group**.
4. Enter a name for the child consistency group or accept the default name; then select the consistency group component type.
5. Select to add existing storage units to the child consistency group or to create new storage units.

If you create new storage units, enter the number of units you want to create and the capacity per unit; then enter the host information.

If you create more than one storage unit, each unit is created with the same capacity and the same host operating system. To assign a different capacity to each unit, select **Add a different capacity**.

6. Select **Add**.

From an existing consistency group

Before you begin

If the consistency group you would like to use is already the child of another consistency group, you must [detach it from the existing parent consistency group](#) before you can move it to a new parent consistency group.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Select the existing consistency group that you would like to make a child consistency group.
3. Select ; then select **Move under different consistency group**.
4. Enter a new name for the child consistency group or accept the default name; then select the consistency group component type.
5. Select the existing consistency group that you would like to make the parent consistency group or select to create a new parent consistency group.

If you select to create a new parent consistency group, enter a name for the parent consistency group or accept the default name; then select the consistency application component type.

6. Select **Move**.

What's next

After you create a child consistency group, you can [apply individual snapshot protection policies](#) to each child consistency group. You can also [set up replication policies](#) on the parent and child consistency groups to replicate the consistency groups to a remote location.

Detach a child consistency group from a parent consistency group

When you detach a child consistency group from a parent consistency group, the child consistency group is removed from the parent consistency group and is managed as a single consistency group. The replication policy applied to the parent is no longer applied to the detached child consistency group.

Before you begin

If you are running ONTAP 9.16.1 and the consistency group you want to detach is in a SnapMirror active sync relationship, you must [delete the SnapMirror active sync relationship](#) before the consistency group can be detached. If you are running ONTAP 9.16.1 and the consistency group is in an asynchronous replication relationship, you must [break the relationship](#) before you can detach the consistency group. Deleting the SnapMirror active sync relationship or breaking the asynchronous relationship before expanding a consistency group is not required in ONTAP 9.17.1 and later releases.

Steps

1. In System Manager, select **Protection > Consistency groups**.
2. Select the parent consistency group.
3. Select over the child consistency group you want to detach.
4. Select ; then select **Detach from parent**.
5. Enter a new name for the consistency group you are detaching or accept the default name; then select the consistency group application type.
6. Select **Detach**.

What's next?

[Set up a replication policy](#) to replicate the snapshots of the detached child consistency group to a remote cluster.

Manage ONTAP data protection policies and schedules on ASA r2 storage systems

Use snapshot policies to protect data in your consistency groups on an automated schedule. Use policy schedules within snapshot policies to 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 > 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 > 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 > 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 > 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 > 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**.

Result

You have modified the replication policy.

Secure your data

Encrypt data at rest on ASA r2 storage systems

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

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.



In ASA r2 systems, SEDs are supported only for NVMe based SSD's.

Use the ASA r2 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 > 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.

Option	Steps
Configure an external key manager	<ol style="list-style-type: none"> Select External key manager to add the key servers. Select + Add to add the key servers. Add the KMIP server CA certificates. Add the KMIP client certificates.

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

What's next?

Now that you have encrypted your data at rest, if you are using the NVMe/TCP protocol, you can [encrypt all the data sent over the network](#) between your NVMe/TCP host and your ASA r2 system.

Migrate ONTAP data encryption keys between key managers on your ASA r2 system

You can manage your data encryption keys using either the ONTAP onboard key manager on your ASA r2 system or an external key manager (or both). External key managers can only be enabled at the storage VM level. At the ONTAP cluster level, you can enable either the onboard key manager or an external key manager.

If you enable your key manager at the...	You can use...
Cluster level only	Either the onboard key manager or an external key manager
Storage VM level only	An external key manager only
Both the cluster and storage VM level	<p>One of the following key manager combinations:</p> <ul style="list-style-type: none"> • Option 1 <ul style="list-style-type: none"> Cluster level: Onboard key manager Storage VM level: External key manager • Option 2 <ul style="list-style-type: none"> Cluster level: External key manager Storage VM level: External key manager

Migrate keys between key managers at the ONTAP cluster level

Beginning with ONTAP 9.16.1 you can use the ONTAP command line interface (CLI) to migrate keys between key managers at the cluster level.

From onboard to external

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Create an inactive external key manager configuration:

```
security key-manager external create-config
```

3. Switch to the external key manager:

```
security key-manager keystore enable -vserver <storage_vm_name>  
-type KMIP
```

4. Delete the onboard key manager configuration:

```
security key-manager keystore delete-config -vserver  
<storage_vm_name> -type OKM
```

5. Set the privilege level to admin:

```
set -privilege admin
```

From external to onboard

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Create an inactive onboard key manager configuration:

```
security key-manager onboard create-config
```

3. Enable the onboard key manager configuration:

```
security key-manager keystore enable -vserver <storage_vm_name>
-type OKM
```

4. Delete the external key manager configuration

```
security key-manager keystore delete-config -vserver
<storage_vm_name> -type KMIP
```

5. Set the privilege level to admin:

```
set -privilege admin
```

Migrate keys between key managers across ONTAP cluster and storage VM levels

You can use the ONTAP command line interface (CLI) to migrate keys between the key manager at the cluster level and a key manager at the storage VM level.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Migrate the keys:

```
security key-manager key migrate -from-vserver <storage_vm_name> -to
-vserver <storage_vm_name>
```

3. Set the privilege level to admin:

```
set -privilege admin
```

Protect against ransomware attacks

Create tamper-proof snapshots to protect against ransomware attacks on ASA r2 storage systems

For enhanced protection against ransomware attacks, replicate snapshots to a remote cluster, then lock the destination snapshots to make them tamper-proof. Locked snapshots cannot be deleted accidentally or maliciously. You can use locked snapshots to recover data if a storage unit is ever compromised by a ransomware attack.

Initialize the SnapLock compliance clock

Before you can create tamper-proof snapshots, you must initialize the SnapLock compliance clock on your local and destination clusters.

Steps

1. Select **Cluster > Overview**.
2. In the **Nodes** section, select **Initialize SnapLock Compliance Clock**.
3. Select **Initialize**.
4. Verify that the compliance clock is initialized.
 - a. Select **Cluster > Overview**.
 - b. In the **Nodes** section, select ; then select **SnapLock Compliance Clock**.

What's next?

After you have initialized the SnapLock compliance clock on your local and destination clusters, you are ready to [create a replication relationship with locked snapshots](#).

Enable autonomous ransomware protection with AI on your ASA r2 storage systems

Beginning with ONTAP 9.17.1, you can use Autonomous Ransomware Protection with Artificial Intelligence (ARP/AI) to protect the data on your ASA r2 system. ARP/AI quickly detects potential ransomware threats, automatically creates an ARP snapshot to protect your data, and displays a warning message in System Manager to alert you of suspicious activity.

ARP improves cyber resiliency by adopting a machine-learning model for anti-ransomware analytics that detects constantly evolving forms of ransomware with 98% accuracy for SAN environments. ARP's machine-learning model is pre-trained on a large dataset of files both before and after a simulated ransomware attack. This resource-intensive training is done outside ONTAP, and the pre-trained model that results from this training is included on-box with ONTAP. This model is not accessible or modifiable. ARP/AI is active immediately after enablement; there is no [learning period](#).



No ransomware detection or prevention system can completely guarantee safety from a ransomware attack. Although an attack might go undetected, ARP/AI acts as an important additional layer of defense if anti-virus software fails to detect an intrusion.

About this task

- ARP/AI support is included with the [ONTAP One license](#).
- ARP/AI is not supported on storage units protected by SnapMirror active sync, SnapMirror synchronous or SnapLock.
- Beginning with ONTAP 9.18.1, ARP/AI is enabled by default on all newly created storage units 12 hours after upgrading to ONTAP 9.18.1 or initializing a new ONTAP 9.18.1 ASA r2 cluster.
- After you have enabled ARP/AI, you should [enable automatic updates for your security files](#) to automatically receive new security updates.

Enable ARP/AI on all storage units in the cluster

If you are running ONTAP 9.17.1, you can enable ARP/AI on all storage units created in the cluster by default.

In ONTAP 9.18.1 and later, ARP/AI is enabled by default on all new storage units. If you have storage units created in ONTAP 9.17.1 for which ARP/AI is not enabled, you can enable it manually.

Steps

1. In System Manager, select **Cluster > Settings**.
2. Next to **Anti-ransomware**, select  and then select **Enable on all existing storage units**.
3. Select **Enable**.

Enable ARP/AI on all storage units in a storage VM

If you are running ONTAP 9.17.1, you can enable ARP/AI on all storage units created in a storage virtual machine (VM) by default. This means that any new storage units created in the storage VM will have ARP/AI enabled automatically. You can also apply ARP/AI to existing storage units in the storage VM.

In ONTAP 9.18.1 and later, ARP/AI is enabled by default on all new storage units. If you have storage units created in ONTAP 9.17.1 for which ARP/AI is not enabled, you can enable it manually.

Steps

1. In System Manager, select **Cluster > Storage VMs**.
2. Select the storage VM on which you want to enable ARP/AI.
3. In the **Security** section, next to **Anti-ransomware**, select ; then select **Edit anti-ransomware settings**.
4. Select **Enable anti-ransomware**.

This enables ARP/AI on all future storage units created on the selected storage VM by default.

5. To apply ARP to existing storage units on the selected storage VM, select **Apply this change to all applicable existing storage units on this storage VM**.
6. Select **Save**.

Result

All new storage units you create on the storage VM are protected against ransomware attacks by default, and suspicious activity is reported to you in System Manager.

Enable ARP/AI on specific storage units in a storage VM

If you are running ONTAP 9.17.1, and you do not want ARP/AI enabled on all the storage units in an storage VM, you can select the specific units you want enabled.

In ONTAP 9.18.1 and later, ARP/AI is enabled by default on all new storage units. If you have storage units created in ONTAP 9.17.1 for which ARP/AI is not enabled, you can enable it manually.

Steps

1. In System Manager, select **Storage**.
2. Select the storage units for which you want to enable ARP/AI.
3. Select ; then select **Enable anti-ransomware**.
4. Select **Enable**.

Result

The storage units you selected are protected against ransomware attacks, and suspicious activity is reported to you in System Manager.

Disable default autonomous ransomware protection on your ASA r2 storage systems

When you initialize a new ONTAP 9.18.1 ASA r2 cluster or upgrade your cluster to ONTAP 9.18.1, ARP/AI is automatically enabled by default on all new storage units after a 12-hour grace period. If you don't disable ARP/AI during the grace period, it is enabled cluster-wide for new storage units when the grace period ends.

Storage units created in ONTAP 9.17.1 must be [manually enabled](#) for ARP/AI.

Steps

You can disable the default enablement during or after the initial 12-hour grace period.

System Manager

1. Select **Cluster > Settings**.
2. Disable ARP:
 - To disable during the 12-hour grace period:
 - a. Under **Anti-ransomware**, select **Don't enable** and then select **Disable**.
 - To disable after the 12-hour grace period:
 - a. Under **Anti-ransomware**, select  and then deselect **Enable for new storage units**.
 - b. Select **Save**

CLI

1. Check the default enablement status:

```
security anti-ransomware auto-enable show
```

2. Disable default enablement for existing and new volumes:

```
security anti-ransomware auto-enable modify -default-existing-volume
-state false -default-new-volume-state false
```

Modify ARP/AI snapshot retention periods on ASA r2 storage systems

If Autonomous Ransomware Protection with Artificial Intelligence (ARP/AI) detects abnormal activity on one or more of your ASA r2 system storage units, it automatically creates an ARP snapshot to protect the storage unit's data. Depending upon your storage capacity and the business requirements for your data, you might want to increase or decrease the default ARP snapshot retention period. For example, you might want to increase the retention period for business critical applications so that, if needed, you have longer retention periods for data recovery, or you might want to decrease the retention period for non-critical applications to save storage space.

The default retention period for the ARP snapshot varies depending on the action you take in response to the

abnormal activity.

If you take this action...	ARP snapshots are retained by default for...
Mark as false positive	12 hours
Mark as potential ransomware attack	7 days
Do not take immediate action	10 days

The default retention periods can be modified using the ONTAP command line interface (CLI). See [Modify options for ONTAP automatic snapshots](#) for steps to change the default retention period.

Respond to autonomous ransomware protection with AI alerts on ASA r2 storage systems

If Autonomous Ransomware Protection with Artificial Intelligence (ARP/AI) detects abnormal activity on one or more of your ASA r2 system storage units, a warning is generated on the System Manager dashboard. You should view the warning, verify the activity and, if necessary, take action to stop any potential threat to your data.

If an ARP/AI warning message is displayed, before you take action, you should use the appropriate application integrity checker to verify the integrity of the data on the storage unit. Verifying the storage unit's data integrity helps you determine if the activity is acceptable or if it is a potential ransomware attack.

If the abnormal activity is ...	Then do this...
Acceptable	Mark the activity as a false positive.
A potential ransomware attack	Mark the activity as a potential ransomware attack.
Indeterminate	Do not take immediate action. Monitor the storage unit for up to 7 days. If the storage unit continues to operate normally, mark the activity as a false positive. If the storage unit continues to exhibit abnormal activity, mark the activity as a potential ransomware attack.

Steps

1. In System Manager, select **Dashboard**.

If ARP has detected abnormal activity on one or more storage units, a message appears under **Warnings**.

2. Select the warning message.
3. Under **Events overview**, select the **Warnings** message that indicates the number of storage units with abnormal activity.
4. Under **Storage units with abnormal activity**, select the storage unit.
5. Select **Security**.

If there is abnormal activity on the storage unit, a message is displayed under **Anti-ransomware**.

6. Select **Choose an action**.
7. Select **Mark as false positive** or select **Mark as potential ransomware attack**.

What's next?

If you know of surges in your storage unit activity, either one-time surges or a surge that is characteristic of a new normal, you should report them as safe. Manually reporting these surges as safe helps to improve the accuracy of ARP's threat assessments. Learn how to [report known ARP/AI surges](#).

Pause or resume autonomous ransomware protection with AI on your ASA r2 storage systems

Beginning with ONTAP 9.17.1, you can use Autonomous Ransomware Protection with Artificial Intelligence (ARP/AI) to protect the data on your ASA r2 system. If you are planning an unusual workload event, you can temporarily suspend ARP/AI analysis to prevent false positive detections of ransomware attacks. After your workload event is complete, you can resume ARP/AI analysis.

Pause ARP/AI

Before you begin an unusual workload event, you might need to temporarily suspend the ARP/AI analysis to prevent false positive detections of ransomware attacks.

Steps

1. In System Manager, select **Storage**.
2. Select the storage units for which you want to pause ARP/AI.
3. Select **Pause anti-ransomware**.

Result

ARP/AI analysis is paused for the selected storage units, and no suspicious activity is reported to you in System Manager until you resume ARP/AI.

Resume ARP/AI

If you pause ARP/AI during an unusual workload, after your workload is complete, you should resume it to protect your data against ransomware attacks.

Steps

1. In System Manager, select **Storage**.
2. Select the storage units for which you want to resume ARP/AI.
3. Select **Resume anti-ransomware**.

Result

Analysis of potential ransomware attacks is resumed, and suspicious activity is reported to you in System Manager.

Secure NVMe connections on your ASA r2 storage systems

If you are using the NVMe protocol, you can configure in-band authentication to enhance your data security. In-band authentication allows secure bidirectional and unidirectional authentication between your NVMe hosts and your ASA r2 system. In-band authentication is available for all NVMe hosts. If you are using the NVMe/TCP protocol, you can further enhance your data security by configuring transport layer security (TLS) to encrypt all data sent over the network between your NVMe/TCP hosts and your ASA r2

system.

Steps

1. Select **Hosts**; then select **NVMe**.
2. Select  .
3. Enter the host name; then select the host operating system.
4. Enter a host description; then select the storage VM to connect to the host.
5. Select  next to the host name.
6. Select **In-band authentication**.
7. If you are using the NVMe/TCP protocol, select **Require Transport Layer Security (TLS)**.
8. Select **Add**.

Result

The security of your data is enhanced with in-band authentication and/or TLS.

Secure IP connections on your ASA r2 storage systems

If you are using the IP protocol on your ASA r2 system, you can configure IP security (IPsec) to enhance your data security. IPsec is an internet standard that provides data-in-flight encryption, authentication for the traffic flowing between the network endpoints at an IP level, and protection against replay and malicious man-in-the-middle attacks on your data.

For ASA r2 systems, IPsec is available for iSCSI and NVMe/TCP hosts.

On certain ASA r2 systems, several of the cryptographic operations, such as encryption and integrity checks, can be offloaded to a supported network interface controller (NIC) card. The throughput for operations offloaded to the NIC card is approximately 5% or less. This can significantly improve the performance and throughput of the network traffic protected by IPsec.

Beginning with ONTAP 9.18.1, IPsec hardware offload supported is extended to IPv6 traffic.

The following NIC cards are supported for hardware offload on the following ASA r2 systems and ONTAP versions:

Supported NIC card	ASA r2 systems	ONTAP Version
X50135A (2p, 40G/100G Ethernet Controller)	<ul style="list-style-type: none">• ASA A1K• ASA A90• ASA A70	ONTAP 9.17.1 and later
X60135A (2p, 40G/100G Ethernet Controller)	<ul style="list-style-type: none">• ASA A50• ASA A30• ASA A20	ONTAP 9.17.1 and later

Supported NIC card	ASA r2 systems	ONTAP Version
X50131A - (2p, 40G/100G/200G/400G Ethernet Controller)	<ul style="list-style-type: none"> • ASA A1K • ASA A90 • ASA A70 	ONTAP 9.16.1 and later
X60132A - (4p, 10G/25G Ethernet Controller)	<ul style="list-style-type: none"> • ASA A50 • ASA A30 • ASA A20 	ONTAP 9.16.1 and later

See the [NetApp Hardware Universe](#) for more information about the supported systems and cards.

What's next?

IPsec is configured on your ASA r2 system the same way as on other ONTAP systems. For more information, see [Prepare to configure IP security for the ONTAP network](#).

Administer and monitor

Upgrade and revert ONTAP

Upgrade ONTAP on ASA r2 storage systems

When you upgrade your ONTAP software on your ASA r2 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.

ONTAP software upgrades for ASA r2 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.

Learn more about [ONTAP software upgrades](#).

Revert ONTAP on ASA r2 storage systems

ONTAP software reverts for ASA r2 systems follow the same process as reverts for other ONTAP systems.

Reverting an ONTAP cluster is disruptive. You must take the cluster offline for the duration of the reversion. You should not revert a production cluster without assistance from technical support. You can revert a new or test cluster without assistance. If the revert of a new or test system fails or if it finishes successfully, but you are not satisfied with the cluster performance in your production environment, you should contact technical support for assistance.

[Revert an ONTAP cluster](#).

Revert requirements for ASA r2 systems

Certain ASA r2 cluster configurations require you to take specific actions before you begin an ONTAP software revert.

Reverting from ONTAP 9.17.1

If you are reverting from ONTAP 9.17.1 on an ASA r2 system, you should perform the following actions before you begin the revert:



[dynamic space balancing](#) is enabled by default 14 days after either upgrading to ONTAP 9.17.1 or initializing a new ONTAP 9.17.1 ASA r2 cluster. You cannot revert from ONTAP 9.17.1 on your ASA r2 system after dynamic space balancing is enabled.

If you have...	Before you revert you should...
Hierarchical consistency groups in a SnapMirror active sync relationship	Delete the SnapMirror active sync relationship.
Active import relationships	Delete the active import relationships. Learn about import relationships.
Anti-ransomware protection enabled	Pause anti-ransomware protection.

Update firmware on ASA r2 storage systems

ONTAP automatically downloads and updates firmware and system files on your ASA r2 system by default. If you want the flexibility of viewing recommended updates before they are downloaded and installed, you can use ONTAP System Manager to disable automated updates or to edit update parameters to show you notifications of available updates before any action is performed.

Enable automatic updates

Recommended updates for storage firmware, SP/BMC firmware and system files are automatically downloaded and installed on your ASA r2 system by default. If automatic updates have been disabled, you can enable them to reinstate the default behavior.

Steps

1. In System Manager, select **Cluster > 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 ASA r2 system based upon your update selections.

Disable automatic updates

Disable automatic updates only if you want to manage updates entirely yourself. With automatic updates turned off, the system will not notify, download, or install updates. You are responsible for monitoring, downloading, scheduling, and installing all 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 , 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.

Manage client access to storage VMs on ASA r2 storage systems


Storage units on an ASA r2 system are contained inside storage virtual machines (VMs). Storage VMs are used to serve data to your SAN clients. Use ONTAP System Manager to create a LIF (network interface) for your SAN clients to connect to a storage VM and

access data in the storage units. You can optionally use subnets to simplify LIF creation and IPspaces to provide your storage VMs with their own secure storage, administration, and routing.

Create a storage VM

During cluster setup, your default data storage virtual machine (VM) is created. All new storage units are created inside your default data storage VM unless you create and select a different storage VM. You might want to create an additional storage VM to segregate your storage units for different applications, departments or clients. For example, you might want to create a storage VM for your development environment and another storage VM for your production environment, or you might want to create a storage VM for your finance department and another storage VM for your marketing department.

Steps

1. Select **Cluster > Storage VMs**.
2. Select .
3. Enter a name for the storage VM or accept the default name.
4. Under **Configure protocols**, select the protocols for the storage VM.

Select **IP** for iSCSI and NVMe/TCP. Select **FC** for Fibre Channel or NVMe/FC.

5. Under **Storage VM administration**, select **Manage administrator account**; then enter the username and password for the administrator account.
6. Add a network interface for the storage VM.
7. Select **Save**.

What's next?

You have created a storage VM. You can now use the storage VM to [provision storage](#).

Create IPspaces

An IPspace is a distinct IP address space in which storage VMs reside. When you create IPspaces, you enable your storage VMs to have their own secure storage, administration, and routing. You also enable clients in administratively separate network domains to use overlapping IP addresses from the same IP address subnet range.

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

Steps

1. Select **Network > Overview**.
2. Under **IPspaces**, select .
3. Enter a name for the IPspace or accept the default name.

An IPspace name cannot be “all” because “all” is a system-reserved name.

4. Select **Save**.

What's next?

Now that you have created an IPspace, you can use it to create a subnet.

Create subnets

A subnet allows you to allocate specific blocks of IPv4 or IPv6 addresses to use when you create a LIF (network interface). A subnet simplifies LIF creation by allowing you to specify the subnet name instead of a specific IP address and network mask for each LIF.

Before you begin

- You must be a cluster administrator to perform this task.
- The [broadcast domain](#) and IPspace where you plan to add the subnet must already exist.

Steps

1. Select **Network > Overview**.
2. Select **Subnets**; then select  .
3. Enter the subnet name.

All subnet names must be unique within an IPspace.

4. Enter the subnet IP address and subnet mask.
5. Specify the IP address range for the subnet.

When you specify the IP address range for the subnet, do not overlap IP addresses with other subnets. Network issues can occur when subnet IP addresses overlap and different subnets or hosts attempt to use the same IP address.

6. Select the broadcast domain for the subnet.
7. Select **Add**.

What's next?

You have created a subnet which you can now use to simplify the creation of your LIFs.

Create a LIF (network interface)

A LIF (network interface) is an IP address associated with a physical or logical port. Create LIFs on the ports you want to use to access data. Storage VMs serve data to clients through one or more LIFs. If there is a component failure, a LIF can fail over or be migrated to a different physical port, so that network communication is not interrupted.

On an ASA r2 system, you can create IP, FC, and NVMe/FC LIFs. An IP data LIF can service both iSCSI and NVMe/TCP traffic by default. Separate data LIFs must be created for FC and NVMe/FC traffic.

If you want to enable automatic iSCSI LIF failover, you must create an IP LIF for iSCSI only traffic. When automatic iSCSI LIF failover is enabled, if a storage failover occurs, the IP iSCSI LIF is automatically migrated from its home node or port to its HA partner node or port and then back once the failover is complete. Or, if the port for an IP iSCSI LIF becomes unhealthy, the LIF is automatically migrated to a healthy port in its current home node and then back to its original port once the port is healthy again.

Before you begin

- You must be a cluster administrator to perform this task.
- The underlying physical or logical network port must have been configured to the administrative **up** status.
- If you are planning to use a subnet name to allocate the IP address and network mask value for a LIF, the

subnet must already exist.

- A LIF handling intracluster traffic between nodes should not be on the same subnet as a LIF handling management traffic or a LIF handling data traffic.

Steps

1. Select **Network > Overview**.
2. Select **Network interfaces**; then select **+ Add**.
3. Select the interface type and protocol; then select the storage VM.
4. Enter a name for the LIF or accept the default name.
5. Select the home node for the network interface; then enter the IP address and subnet mask.
6. Select **Save**.

Result

You have created a LIF for data access.

What's next?

You can use the ONTAP command line interface (CLI) to create an iSCSI-only LIF with automatic failover.

Create a custom iSCSI-only LIF service policy

If you would like to create iSCSI-only LIFs with automatic LIF failover, you must first create a custom iSCSI-only LIF service policy.

You must use the ONTAP command line interface (CLI) to create the custom service policy.

Step

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Create a custom iSCSI-only LIF service policy:

```
network interface service-policy create -vserver <storage_VM_name>
-priority <service_policy_name> -services data-core,data-iscsi
```

3. Verify that the service policy was created:

```
network interface service-policy show -policy <service_policy_name>
```

4. Return the privilege level to admin:

```
set -privilege admin
```

Create iSCSI-only LIFs with automatic LIF failover

If there are iSCSI LIFs on the storage VM that are not enabled for automatic LIF failover, your newly created LIFs will not be enabled for automatic LIF failover either. If automatic LIF failover is not enabled and there is a failover event your iSCSI LIFs will not migrate.

Before you begin

You must have created a custom iSCSI-only LIF service policy.

Steps

1. Create iSCSI-only LIFs with automatic LIF failover:

```
network interface create -vserver <storage_VM_name> -lif
<iscsi_lif_name> -service-policy <service_policy_name> -home-node
<home_node> -home-port <port_name> -address <ip_address> -netmask
<netmask> -failover-policy sfo-partner-only -status-admin up
```

- It is recommended that you create two iSCSI LIFs on each node, one for fabric A and the other for fabric B. This provides redundancy and load balancing for your iSCSI traffic. In the following example, a total four iSCSI LIFs are created, two on each node, one for each fabric.

```
network interface create -vserver svml -lif iscsi-lif-01a -service
-policy custom-data-iscsi -home-node node1 -home-port e2b -address
<node01-iscsi-a-ip> -netmask 255.255.255.0 -failover-policy sfo-
partner-only -status-admin up
```

```
network interface create -vserver svml -lif iscsi-lif-01b -service
-policy custom-data-iscsi -home-node node1 -home-port e4b -address
<node01-iscsi-b-ip> -netmask 255.255.255.0 -failover-policy sfo-
partner-only -status-admin up
```

```
network interface create -vserver svml -lif iscsi-lif-02a -service
-policy custom-data-iscsi -home-node node2 -home-port e2b -address
<node02-iscsi-a-ip> -netmask 255.255.255.0 -failover-policy sfo-
partner-only -status-admin up
```

```
network interface create -vserver svml -lif iscsi-lif-02b -service
-policy custom-data-iscsi -home-node node2 -home-port e4b -address
<node02-iscsi-b-ip> -netmask 255.255.255.0 -failover-policy sfo-
partner-only -status-admin up
```

- If you are using VLANs, adjust the `home-port` parameter to include the VLAN port information for the respective iSCSI fabric, for example, `-home-port e2b-<iSCSI-A-VLAN>` for iSCSI fabric A and `-home-port e4b-<iSCSI-B-VLAN>`.
- If you are using interface groups (ifgroups) with VLANs, adjust the `home-port` parameter to include the appropriate VLAN port, for example, `-home-port a0a-<iSCSI-A-VLAN>` for iSCSI fabric A and

-home-port a0a-*<iSCSI-B-VLAN>* for iSCSI fabric B where a0a is the ifgroup and a0a-*<iSCSI-A-VLAN>* and a0a-*<iSCSI-B-VLAN>* are the respective VLAN ports for the iSCSI A fabric and the iSCSI B fabric.

2. Verify that the iSCSI LIFs were created:

```
network interface show -lif iscsi*
```

Modify a LIF (network interfaces)

LIFs can be disabled or renamed as needed. You can also change the LIF IP address and subnet mask.

About this task

ONTAP utilizes Network Time Protocol (NTP) to synchronize time across the cluster. After changing LIF IP addresses, you may need to update the NTP configuration to prevent synchronization failures. For more information, refer to the Knowledge Base article [NTP synchronization fails after LIF IP change](#).

Steps

1. Select **Network > Overview**; then select **Network interfaces**.
2. Hover over the network interface you want to edit; then select .
3. Select **Edit**.
4. You can disable the network interface, rename the network interface, change the IP address, or change the subnet mask.
5. Select **Save**.

Result

Your LIF has been modified.

Manage cluster networking on ASA r2 storage systems

You can use ONTAP System Manager to perform basic storage network administration on your ASA r2 system. For example, you can add a broadcast domain or reassign ports to a different broadcast domain.

Add a broadcast domain

Use broadcast domains to simplify management of your cluster network by grouping network ports that belong to the same layer 2 network. Storage virtual machines (VMs) can then use the ports in the group for data or management traffic.

The “Default” broadcast domain and the “Cluster” broadcast domain are created during cluster setup. The “Default” broadcast domain contains ports that are in the “Default” IPspace. These ports are used primarily to serve data. Cluster management and node management ports are also in this broadcast domain. The “Cluster” broadcast domain contains ports that are in the “Cluster” IPspace. These ports are used for cluster communication and include all 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) of the ports added to a broadcast domain are updated to the MTU value set in the broadcast domain.

Steps

1. In System Manager, select **Network > Overview**.
2. Under **Broadcast** domains, select  .
3. Enter a name for the broadcast domain or accept the default name.

All broadcast domain names must be unique within an IPspace.

4. Select the IPspace for the broadcast domain.

If you don't specify an IPspace name, the broadcast domain is created in the "Default" IPspace.

5. Enter the maximum transmission unit (MTU).

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

6. Select the desired ports; then select **Save**.

Result

You have added a new broadcast domain.

Reassign ports to a different broadcast domain

Ports can belong to only one broadcast domain. If you want to change the broadcast domain to which a port belongs, you need to reassign the port from its existing broadcast domain to a new broadcast domain.

Steps

1. In System Manager, select **Network > Overview**.
2. Under **Broadcast Domains**, select  next to the domain name; then select **Edit**.
3. Deselect the Ethernet ports that you want to reassign to another domain.
4. Select the broadcast domain to which you want to reassign the port; then select **Reassign**.
5. Select **Save**.

Result

You have reassigned ports to a different broadcast domain.

Create a VLAN

A VLAN consists of switch ports grouped together into a broadcast domain. VLANs enable you to increase security, isolate 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

- A VLAN can't be created on an interface group port that contains no 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 > Ethernet ports**; then select  **VLAN**.
2. Select the node and broadcast domain for the VLAN.
3. Select the port for the VLAN.

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

4. Enter a VLAN ID.
5. Select **Save**.

Result

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

Monitor usage and increase capacity

Monitor cluster and storage unit performance on ASA r2 storage systems

Use ONTAP System Manager to monitor the overall performance of your cluster and the performance of specific storage units to determine how latency, IOPS and throughput are impacting your critical business applications. Performance can be monitored over various spans of time ranging from one hour to one year.

For example, suppose a critical application is experiencing high latency and low throughput. When you view cluster performance for the last five business days, you notice a decrease in performance at the same time each day. You use this information to determine that the critical application is competing for cluster resources when a non-critical process begins running in the background. You are then able to modify your QoS policy to limit the impact of the non-critical workload on system resources and to ensure that your critical workload meets minimum throughput targets.

Monitor cluster performance

Use cluster performance metrics to determine whether you need to shift workloads to minimize latency and maximize IOPS and throughput for your critical applications.

Steps

1. In System Manager, select **Dashboard**.
2. Under **Performance**, view the latency, IOPS, and throughput for the cluster by hour, day, week, month, or year.
3.  Select  to download the performance data.

What's next?

Use your cluster performance metrics to analyze if you need to modify your QoS policies or make other adjustments to your application workloads to maximize your overall cluster performance.

Monitor storage unit performance

Use storage unit performance metrics to determine the impact of specific applications on latency, IOPS and throughput.

Steps

1. In System Manager, select **Storage**.
2. Select the storage unit you want to monitor; then select **Overview**.
3. Under **Performance**, view the latency, IOPS, and throughput for the storage unit by hour, day, week, month, or year.
4. Select  to download the performance data.

What's next?

Use your storage unit performance metrics to analyze if you need to modify the QoS policies assigned to your storage units to decrease latency and maximize IOPS and throughput.

Monitor cluster and storage unit utilization on ASA r2 storage systems

Use ONTAP System Manager to monitor your storage utilization to ensure you have the storage capacity you need to serve current and future workloads.

Monitor cluster utilization

Regularly monitor the amount of storage consumed by your cluster to ensure that, if needed, you are prepared to expand the cluster capacity before running out of space.

Steps

1. In System Manager, select **Dashboard**.
2. Under **Capacity**, view the amount of physical used space and the amount of available space on your cluster.

The data reduction ratio represents the amount of space saved through storage efficiency.

What's next?

If your cluster is running low on space or if it doesn't have the capacity to meet a future demand, you should plan to [add new drives](#) to your ASA r2 system to increase your storage capacity.

Monitor storage availability zone utilization

Each HA pair in an ASA r2 system uses a common pool of storage called a *storage availability zone*. The storage availability zone has access to all available disks in the storage system and is visible to both nodes in the HA pair.

If you have 4 or more nodes in your cluster, you can view the amount of space used by the storage availability zone for each HA pair. This metric is not available for 2-node clusters.

Steps

1. In System Manager, select **Cluster**; then select **Overview**.

A summary of the storage availability zone utilization is displayed for each HA pair in the cluster.

2. If you want more detailed metrics, select a specific storage availability.

Under **Overview**, the capacity of the storage availability zone, the amount of used space, and the data reduction ratio is displayed.

Under **Storage units** a list of all the storage units in the storage availability zone is displayed.

What's next?

If your storage availability zone is running low on space, you should plan to [move storage units](#) to another storage availability zone to balance the storage utilization across the cluster.

Monitor storage unit utilization

Monitor the amount of storage consumed by a storage unit so that you can proactively increase the size of the storage unit based on your business needs.

Steps

1. In System Manager, select **Storage**.
2. Select the storage unit you want to monitor; then select **Overview**.
3. Under **Storage**, view the following:
 - Size of your storage unit
 - Amount of used space
 - Data reduction ratio

The data reduction ratio represents the amount of space saved through storage efficiency

- Snapshot used

Snapshot used represents the amount of storage used by snapshots.

What's next?

If your storage unit is nearing capacity, you should [modify the storage unit](#) to increase its size.

Increase storage capacity on ASA r2 storage systems

Add drives to a node or shelf to increase the storage capacity of your ASA r2 system.

Use NetApp Hardware Universe to prepare for installation of a new drive

Before you install a new drive to a node or shelf, use the NetApp Hardware Universe to confirm that the drive you want to add is supported by your ASA r2 system and to identify the correct slot for the new drive. The correct slots for adding drives vary depending on the system model and ONTAP version. In some cases, you need to add drives to specific slots in sequence.

Steps

1. Go the [NetApp Hardware Universe](#).

2. Under **Products**, select your hardware configurations.
3. Select your ASA r2 system.
4. Select your ONTAP version; then select **Show Results**.
5. Beneath the graphic, select **Click here to see alternative views**; then choose the view that matches your configuration.
6. Use the view of your configuration to confirm that your new drive is supported and the correct slot for installation.

Result

You have confirmed that your new drive is supported and you know the appropriate slot for installation.

Install a new drive on the ASA r2

The minimum number of drives you should add in a single procedure is six. Adding a single drive might reduce performance.

About this task

You should repeat the steps in this procedure for each drive.

Steps

1. Properly ground yourself.
2. Gently remove the bezel from the front of the system.
3. Insert the new drive into the correct slot.
 - a. With the cam handle in the open position, use both hands to insert the new drive.
 - b. Push until the drive stops.
 - c. Close the cam handle so that the drive is fully seated into the mid plane and the handle clicks into place.

Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

4. Verify that the drive's activity LED (green) is illuminated.
 - If the LED is solid, the drive has power.
 - If the LED is blinking, the drive has power and I/O is in progress. The LED will also blink if the drive firmware is being updated.

Drive firmware is automatically updated (nondisruptively) on new drives that do not have current firmware versions.

5. If your node is configured for drive auto-assignment, you can wait for ONTAP to automatically assign the new drives to a node. If your node isn't configured for drive auto-assignment or if preferred, you can assign the drives manually.

The new drives are not recognized until they are assigned to a node.


What's next?

After the new drives have been recognized, verify that they have been added and their ownership is specified correctly.

Optimize cluster security and performance with ASA r2 storage system insights

View *Insights* in ONTAP System Manager to identify best practices and configuration modifications that you can implement on your ASA r2 system to optimize cluster security and performance.

For example, suppose you have Network Time Protocol (NTP) servers configured for your cluster. However, you are unaware that you have less than the recommended number of NTP servers needed for optimal cluster time management. To help you prevent problems that can occur when the cluster time is inaccurate, Insights will notify you that you have too few NTP servers configured and give you options to either learn more about this issue, fix it, or dismiss it.



The screenshot shows the 'Insights' page in ONTAP System Manager. At the top, a banner says 'Take action to address concerns and apply best practices to optimize the security and performance of your system.' Below this, the 'Apply best practices' section lists five recommendations:

- Login banner isn't configured**: You haven't configured one or more login banner messages. You can create a custom login banner for the cluster or storage VM to inform visitors about terms and conditions, acceptable use, and site permissions. [Learn more about best practices for security.](#)
- Too few NTP servers are configured**: Problems can occur when the cluster time is inaccurate. Configure Network Time Protocol (NTP) servers to synchronize the cluster time with external NTP servers. For redundancy and accuracy, you should associate at least three NTP servers with the cluster. [Learn more about best practices for security.](#)
- Cluster isn't configured for automatic updates**: You aren't receiving automatic updates for this cluster. Enable automatic updates to always get the latest disk qualification package, disk firmware, shelf firmware, and SP/BMC firmware files when available.
- Global FIPS 140-2 compliance is disabled**: Global FIPS 140-2 compliance is disabled on this cluster. For security reasons, you should ensure ONTAP communicates with external clients or server components outside of ONTAP by using SSL communication that uses FIPS 140-2 compliant cryptography. [Learn more about best practices for security.](#)
- Cluster isn't configured for notifications**: You aren't receiving notifications from ONTAP about potential problems on the cluster. You can configure ONTAP to send notifications using email, a webhook, or an SNMP traphost.

Steps

1. In System Manager, select **Insights**.
2. Review recommendations.

What's next

Perform any necessary actions to implement best practices and optimize your cluster security and performance.

View cluster events and jobs on ASA r2 storage systems

Use ONTAP System Manager to view a list of errors or alerts that have occurred in your system along with recommended corrective actions. You can also view system audit logs and a list of jobs that are active, completed, or failed.

Steps

1. In System Manager, select **Events & Jobs**.
2. View cluster events and jobs.

To view this...	Do this...
Cluster events	Select Events ; then select Event log .
Active IQ suggestions	Select Events ; then select Active IQ suggestions .
System alerts	<ol style="list-style-type: none"> Select System alerts. Select the system alert for which you want to take action. Acknowledge or suppress the alert.
Cluster jobs	Select Jobs .
Audit logs	Select Audit logs .

Send email notifications for cluster events and audit logs

Configure your system to send a notification to specific email addresses when there is a cluster event or audit log entry.

Steps

1. In System Manager, select **Cluster > Settings**.
2. Next to **Notifications management** select .
3. To configure an event destination, select **View event destinations**; then select **Event destinations**. To configure an audit log destination, select **View audit destinations**; then select **Audit log destinations**.
4. Select .
5. Enter the destination information; then select **Add**.

Result

The email address you added will now receive the specified email notifications for cluster events and audit logs.

Manage nodes

Add ASA r2 nodes to an ONTAP cluster

Beginning with ONTAP 9.16.1, ASA r2 storage systems support up to 12 nodes per cluster. After the new nodes of an HA pair have been cabled and powered on, you need to join them to the cluster.

Before you begin

Gather the following information:

- The node IP address
- The intercluster network interface IP address
- The intercluster network subnet mask

- The intercluster network gateway
- If you want to configure the onboard key manager (OKM), you need the OKM passphrase.

Steps

1. In System Manager, select **Cluster > Overview**.
2. Select  next to the node you want to join to the cluster; then select **Add node**
3. Enter the IP address for each node.
4. Enter the intercluster network interface IP address, subnet mask, and gateway.
5. If you want to configure the onboard key manager (OKM), enter the OKM passphrase.

Configure onboard key manager for encryption is selected by default.

6. Select **Add**.

Result

The new HA pair is joined to the cluster.

What's next?

After you add the new HA pair to the cluster, you can [enable data access from your SAN hosts](#) to your new nodes.

Reboot a node on an ASA r2 storage system

You might need to reboot a node for maintenance, troubleshooting, software updates or other administrative reasons. When a node is rebooted, its HA partner automatically executes a takeover. The partner node then performs an automatic giveback after the rebooted node comes back online.

Steps

1. In System Manager, select **Cluster > Overview**.
2. Select  next to the node you want to reboot; then select **Reboot**.
3. Enter the reason you are rebooting the node; then select **Reboot**.

The reason you enter for the reboot is recorded in the system audit log.

What's next?

While the node is being rebooted, its HA partner performs a takeover so that there is no interruption in data service. When the reboot is complete, the HA partner performs a giveback.

Rename a node on an ASA r2 storage system

You can use ONTAP System Manager to rename a node on your ASA r2 system. You might need to rename a node to align with the naming conventions of your organization or for other administrative reasons.

Steps

1. In System Manager, select **Cluster > Overview**.

2. Select  next to the node you want to rename; then select **Rename**.
3. Enter the new name for the node, then select **Rename**.

Result

The new name is applied to the node.

Manage user accounts and roles on ASA r2 storage systems

Use System Manager to configure active directory domain controller access, LDAP and SAML authentication for your user accounts. Create user account roles to define specific functions that users assigned to the roles can perform on your cluster.

Configure active directory domain controller access

Configure active directory (AD) domain controller access to your cluster or storage VM so that you can enable AD account access.

Steps

1. In System Manager, select **Cluster > Settings**.
2. In the **Security** section, under **Active Directory**, select **Configure**.

What's next?

You can now enable AD account access on your ASA r2 system.

Configure LDAP

Configure a Lightweight Directory Access Protocol (LDAP) server to centrally maintain user information for authentication.

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 > Settings**.
2. In the **Security** section, next to **LDAP**, select .
3. Enter the necessary LDAP server and binding information; then select **Save**.

What's next?

You can now use LDAP for user information and authentication.

Configure SAML authentication

Security Assertion Markup Language (SAML) authentication allows users to be authenticated by a secure identity provider (IdP) instead of the direct service providers such as Active Directory and LDAP.

Before you begin

- The IdP that you plan to use for remote authentication must be configured.

See the IdP documentation for configuration.

- You must have the URI of the IdP.

Steps

1. In System Manager, select **Cluster > Settings**.
2. Under **Security**, next to **SAML authentication**, select .
3. Select **Enable SAML authentication**.
4. Enter the IdP URL and the host system IP address; then select **Save**.

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

5. Go to the IdP system you specified; then copy the metadata from your clipboard to update the system metadata.
6. Return to the confirmation window in System Manager; then 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.


What's next?

You can now use SAML authentication for your user accounts.

Create user account roles

Roles for cluster administrators and storage VM administrators are automatically created when your cluster is initialized. 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 > Settings**.
2. In the **Security** section, next to **Users and roles**, select .
3. Under **Roles**, select .
4. Select the role attributes.

To add multiple attributes, select .

5. Select **Save**.


Result

A new user account is created and available for use on your ASA r2 system.

Create an administrator account

Create an administrator user account to enable the account user to perform specific actions on your cluster based on the role assigned to the account. To enhance account security, set up multi-factor authentication (MFA) when you create the account.

Steps

1. In System Manager, select **Cluster > Settings**.
2. In the **Security** section, next to **Users and roles**, select .
3. Under **Users**, select .
4. Enter a username; then select a role to assign to the user.
5. Select the user login method and the authentication method.
6. To enable MFA, select ; then select a secondary login method and authentication method
7. Enter a password for the user.
8. Select **Save**.

Result

A new administrator account is created and available for use on your ASA r2 cluster.

Manage security certificates on ASA r2 storage systems

Use digital security certificates to verify the identity of remote servers.

Online Certificate Status Protocol (OCSP) validates the status of digital certificate requests from ONTAP services using SSL and Transport Layer Security (TLS) connections.

Generate a certificate signing request

Generate a certificate signing request (CSR) to create a private key which can be used to generate a public certificate.

Steps

1. In System Manager, select **Cluster > Settings**.
2. Under **Security**, next to **Certificates**, select ; then select .
3. Enter the subject common name; then select the country.
4. If you want to change the GSR defaults, select extended key usage, or add subject alternative names, select  **More options**; then make the desired updates.
5. Select **Generate**.

Result

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

Add a trusted certificate authority

ONTAP provides a default set of trusted root certificates for applications using Transport Layer Security (TLS). You can add additional trusted certificate authorities as needed.

Steps

1. Select **Cluster > Settings**.
2. Under **Security**, next to **Certificates**, select .
3. Select **Trusted certificate authorities**.
4. Enter or import the certificate details; then select .

Result

You have added a new trusted certificate authority to your ASA r2 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 > Settings**.
2. Under **Security**, next to **Certificates**, select .
3. Select **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">Select ; then select Renew.Enter or import the certificate information; then select Renew.	<ol style="list-style-type: none">Select ; then select Delete.Confirm that you want to delete; then select Delete.

Result

You have renewed or deleted an existing trusted certificate authority on your ASA r2 system.

Add a client/server certificate or local certificate authorities

Add a client/server certificate or local certificate authorities to enable secure web services.

Steps

1. In System Manager, select **Cluster > Settings**.
2. Under **Security**, next to **Certificates**, select .
3. Select **Client/server certificates** or **Local certificate authorities**.
4. Add the certificate information; then select  **Add**.

Result

You have added a new client/server certificate or local authorities to your ASA r2 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 > Settings**.
2. Under **Security**, next to Certificates, select .
3. Select **Client/server certificates**, or **Local certificate authorities**.

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...
<p>a. Select ; then select Renew.</p> <p>b. Enter or import the certificate information; then select Renew.</p>	Select  ; then select Delete .

Result

You have renewed or deleted an existing client/server certificate or local certificate authority on your ASA r2 system.

Verify host connectivity on your ASA r2 storage system

If there is an issue with host data operations, you can use ONTAP System Manager to verify that the connection from your host to your ASA r2 storage system is active.

Steps

1. In System Manager, select **Host**.

The host connectivity status is indicated next to name of the host group as follows:

- **OK**: Indicates all initiators are connected to both nodes.
- **Partially Connected**: Indicates that some of the initiators are not connected both nodes.
- **None Connected**: Indicates that no initiators are connected.

What's next?

Make updates on your host to correct connectivity issues. ONTAP will recheck the connection status every fifteen minutes.

Maintain your ASA r2 storage system

Go to the [ASA r2 maintenance documentation](#) to learn how to perform maintenance procedures on your ASA r2 system components.

Learn more

ASA r2 for ONTAP power users

Compare ASA r2 systems to other ONTAP systems

ASA r2 systems offer a hardware and software solution for SAN-only environments built on all flash solutions. ASA r2 systems vary from other ONTAP systems (ASA, AFF, and FAS) in the implementation of its ONTAP personality, storage layer and supported protocols.

The following are classified as ASA r2 systems:

- ASA A1K
- ASA A90
- ASA A70
- ASA A50
- ASA A30
- ASA A20
- ASA C30

Personality differences

On an ASA r2 system, ONTAP software is streamlined to provide support for essential SAN functionality while limiting the visibility and availability of non-SAN related features and functions. For example, System Manager running on an ASA r2 system does not display options to create home directories for NAS clients. This streamlined version of ONTAP is identified as the *ASA r2 personality*. ONTAP running on ASA systems is identified as the *ASA ONTAP personality*. ONTAP running on AFF and FAS ONTAP systems is identified as the *unified ONTAP personality*. The differences between ONTAP personalities are referenced in the ONTAP command reference (man pages), REST API specification, and EMS messages where applicable.

You can verify the personality of your ONTAP storage from System Manager or from the ONTAP CLI.

- From the System Manager menu, select **Cluster > Overview**.
- From the CLI, enter: `system node show -personality -is-disaggregated`

For ASA r2 systems, the *personality* is *ASA r2* and the status of *is-disaggregated* is *true*.

The personality of your ONTAP storage system cannot be changed.

Storage layer differences

ASA r2 systems use a simplified storage layer that is different from the storage layer used by FAS, AFF, and ASA systems.

FAS, AFF, and ASA systems

The storage layer for FAS, AFF, and ASA systems use aggregates as the base unit of storage. An aggregate owns a specific set of the disks available in a storage system. The aggregate allocates space on the disks it

owns to volumes for LUNs and namespaces. With these systems, ONTAP users can create and modify aggregates, volumes, LUNs and namespaces.

ASA r2 systems

Instead of aggregates, the storage layer in ASA r2 systems uses storage availability zones. A storage availability zone is a common pool of storage available to both nodes of a single HA pair. Both nodes in the HA pair have access to all available disks in their shared storage availability zone. For example, in a 2-node ASA r2 system ONTAP cluster, there is one storage availability zone, accessible by both nodes in the cluster. In a 4-node ASA r2 system ONTAP cluster, there are two storage availability zones. Each HA pair in the cluster has access to one of the storage availability zones.

When a storage unit (based on either a LUN or an NVMe namespace) is created, ONTAP automatically creates a volume in the appropriate storage availability zone to house the storage unit. The newly created volume is automatically placed within the storage availability zone for optimal performance and balanced capacity utilization. Capacity utilization is balanced within the storage availability zone based on your version of ONTAP. [Learn about capacity balancing in an ASA r2 cluster](#).

Summary of ASA r2 system differences

ASA r2 systems differ from FAS, AFF, and ASA systems in the following ways:

	ASA r2	ASA	AFF	FAS
ONTAP personality	ASA r2	ASA	Unified	Unified
SAN protocol support	Yes	Yes	Yes	Yes
NAS protocol support	No	No	Yes	Yes
Storage layer support	Storage availability zone	Aggregates	Aggregates	Aggregates

Because of this automated and simplified approach to storage management, certain System Manager options, ONTAP commands, and REST API endpoints are not available or have limited usage on an ASA r2 system. For example, because volume creation and management is automated for ASA r2 systems, the **Volumes** menu does not appear in System Manager and the `volume create` command is not supported. [Learn more about unsupported ASA r2 commands](#).

The major differences between ASA r2 systems and FAS, AFF, and ASA systems relevant to the ONTAP command line interface (CLI) and REST API are described below.

Default storage VM creation with protocol services

New clusters automatically contain a default data storage virtual machine (VM) with the SAN protocols enabled. IP data LIFs support iSCSI and NVMe/TCP protocols and use the `default-data-blocks` service policy by default.

Automatic volume creation

Creating a storage unit (LUN or namespace) automatically creates a volume from the storage availability zone. This results in a simplified and common namespace. Deleting a storage unit automatically deletes the associated volume.

Changes to thin and thick provisioning

Storage units are always thinly provisioned on ASA r2 storage systems. Thick provisioning is not supported.

Changes to data compression

Temperature-sensitive storage efficiency is not applied on ASA r2 systems. On ASA r2 systems, compression is not based on *hot* (frequently accessed) data or *cold* (infrequently accessed) data. Compression begins without waiting for data to become cold.

For more information

- Learn more about [ONTAP hardware systems](#).
- See full configuration support and limitations for ASA and ASA r2 systems in [NetApp Hardware Universe](#).
- Learn more about the [NetApp ASA](#).

ONTAP software support and limitations for ASA r2 storage systems

While ASA r2 systems offers a wide range of support for SAN solutions, certain ONTAP software features are not supported.

ASA r2 systems do not support the following:

- Default automatic iSCSI LIF failover

In ASA r2 systems, the default networking LIF is shared between NVMe and SCSI hosts, so doesn't support automatic failover. To enable automatic iSCSI LIF failover, you must [create an iSCSI only LIF](#). Automatic failover is enabled on iSCSI only LIFS by default.

When automatic iSCSI LIF failover is enabled, if a storage failover occurs, the iSCSI LIF is automatically migrated from its home node or port to its HA partner node or port and then back once the failover is complete. Or, if the port for an iSCSI LIF becomes unhealthy, the LIF is automatically migrated to a healthy port in its current home node and then back to its original port once the port is healthy again.

- FabricPool
- LUN thick provisioning
- MetroCluster
- Object protocols
- ONTAP S3 SnapMirror and S3 APIs

ASA r2 systems support the following:

- Snaplock

[Learn how to lock snapshots](#) on your ASA r2 system.

- Dual-layer encryption

[Learn how to apply dual-layer encryption](#) to data on your ASA r2 system.

Support for SnapMirror replication

SnapMirror replication is supported on ASA r2 systems with the following limitations:

- SnapMirror synchronous replication is not supported.
- SnapMirror active sync is supported only between two ASA r2 systems.

Learn more about [SnapMirror active sync on ASA r2 systems](#).

- SnapMirror asynchronous replication is supported only between two ASA r2 systems. SnapMirror asynchronous replication is not supported between an ASA r2 system and an ASA, AFF or FAS system or the cloud.

Learn more about [SnapMirror replication policies supported on ASA r2 systems](#).

For more information

- See the [NetApp Hardware Universe](#) for more information on ASA r2 hardware support and limitations.

ONTAP CLI support for ASA r2 storage systems

Instead of aggregates, the storage layer in ASA r2 systems use storage availability zones. A storage availability zone is a common pool of storage available to a single HA pair. Both nodes in the HA pair have access to all available disks in their shared storage availability zone. When a storage unit (LUN or NVMe namespace) is created, ONTAP automatically creates a volume in the appropriate storage availability zone to house the storage unit.

Because of this simplified approach to storage management, storage aggregate commands are not supported on ASA r2 systems. Support for certain lun, storage and volume commands and parameters is also limited.

The following commands and command sets are not supported on ASA on r2:

Unsupported lun commands

- lun copy
- lun geometry
- lun maxsize
- lun move
- lun move-in-volume



The lun move-in-volume command is replaced with the lun rename and the vserver nvme namespace rename commands.

- lun transition

Unsupported storage commands

- storage failover show-takeover
- storage failover show-giveback
- storage aggregate relocation
- storage disk assign
- storage disk partition
- storage disk reassign

Unsupported volume command sets

- volume activity-tracking
- volume analytics
- volume conversion
- volume file
- volume flexcache
- volume flexgroup
- volume inode-upgrade
- volume object-store
- volume qtree
- volume quota
- volume reallocation
- volume rebalance
- volume recovery-queue
- volume schedule-style

Unsupported volume commands and parameters

- volume autosize
- volume create
- volume delete
- volume expand
- volume modify

The `volume modify` command is not available when used in conjunction with the following parameters:

- `-anti-ransomware-state`
- `-autosize`
- `-autosize-mode`
- `-autosize-shrik-threshold-percent`
- `-autosize-reset`
- `-group`
- `-is-cloud-write-enabled`
- `-is-space-enforcement-logical`
- `-max-autosize`
- `-min-autosize`
- `-offline`
- `-online`
- `-percent-snapshot-space`
- `-qos*`
- `-size`
- `-snapshot-policy`
- `-space-guarantee`
- `-space-mgmt-try-first`
- `-state`
- `-tiering-policy`
- `-tiering-minimum-cooling-days`
- `-user`
- `-unix-permissions`
- `-vserver-dr-protection`

- `volume make-vsroot`

- volume mount
- volume move
- volume offline
- volume rehost
- volume rename
- volume restrict
- volume transition-prepare-to-downgrade
- volume unmount

Unsupported volume clone **commands**

- volume clone create
- volume clone split

Unsupported volume snaplock **commands**

- volume snaplock modify

Unsupported volume snapshot **commands**

- volume snapshot
- volume snapshot autodelete modify
- volume snapshot policy modify

For more information

See the [ONTAP command reference](#) for a full list of supported commands

Set up an ONTAP ASA r2 cluster using the CLI

It is recommended that you [use System Manager to set up your ONTAP ASA r2 cluster](#). System Manager offers a quick and easy guided workflow to get your cluster up and running. However, if you are accustomed to working with ONTAP commands, the ONTAP command line interface (CLI) can optionally be used for cluster setup. Cluster set up using the CLI offers no additional options or advantages than cluster set up using System Manager.

During cluster setup, your default data storage virtual machine (VM) is created, an initial storage unit is created, and your data LIFs are automatically discovered. Optionally, you can enable the Domain Name System (DNS) to resolve host names, set your cluster to use the Network Time Protocol (NTS) for time synchronization, and enable encryption of data at rest.

Before you begin

Gather 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 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
- Partner node IP address
- DNS domain names
- DNS name server IP addresses
- NTP server IP addresses
- Data subnet mask

Steps

1. Power on both nodes of the HA pair.
2. Show the nodes discovered on the local network:

```
system node show-discovered -is-in-cluster false
```

3. Start the cluster setup wizard:

```
cluster setup
```

4. Acknowledge the AutoSupport statement.
5. Enter values for the node management interface port, IP address, netmask and default gateway.
6. Press **Enter** to continue setup using the command line interface; then enter **create** to create a new cluster.
7. Accept the system defaults or enter your own values.
8. After setup on the first node is complete, log into the cluster.
9. Verify that the cluster is active and the first node is healthy:

```
system node show-discovered
```

10. Add the second node to the cluster:

```
cluster add-node -cluster-ip <partner_node_ip_address>
```

11. Optionally, synchronize the system time across the cluster

Synchronize without symmetric authentication

```
cluster time-service ntp server  
create -server <server_name>
```

Synchronize with symmetric authentication

```
cluster time-service ntp server  
create -server  
<server_ip_address> -key-id  
<key_id>
```

- a. Verify that the cluster is associated with an NTP server:

```
Cluster time-service ntp show
```

12. Optionally, download and run [ActiveIQ Config Advisor](#) to confirm your configuration.

What's next?

You are ready to [set up data access](#) from your SAN clients to your system.

REST API support for ASA r2

The ASA r2 REST API is based on the REST API provided with the unified ONTAP personality, with a number of changes adapted to the unique characteristics and capabilities of the ASA r2 personality.

Types of API changes

There are several types of differences between the ASA r2 system REST API and the unified ONTAP REST API available with FAS, AFF, and ASA systems. Understanding the types of changes will help you better utilize the online API reference documentation.

New ASA r2 endpoints not supported in unified ONTAP

Several endpoints have been added to the ASA r2 REST API which are not available with unified ONTAP.

For example, a new block-volume endpoint has been added to the REST API for ASA r2 systems. The block-volume endpoint provides access to both LUN and NVMe namespace objects, enabling an aggregated view of the resources. This is only available through the REST API.

As another example, the **storage-units** endpoints provide an aggregated view of the LUNs and NVMe namespaces. There are several endpoints and they're all based on or derived from `/api/storage/storage-units`. You should also review `/api/storage/luns` and `/api/storage/namespaces`.

Restrictions on the HTTP methods used for some endpoints

Several endpoints available with ASA r2 have restrictions on which HTTP methods can be used as compared with unified ONTAP. For example, POST and DELETE are not allowed when using the endpoint `/api/protocols/nvme/services` with ASA r2 systems.

Property changes for an endpoint and HTTP method

Some ASA r2 system endpoint and method combinations do not support all the defined properties available in the unified ONTAP personality. For example, when using PATCH with the endpoint `/api/storage/volumes/{uuid}`, several properties are not supported with ASA r2, including:

- `autosize.maximum`
- `autosize.minimum`
- `autosize.mode`

Changes to internal processing

There are several changes to how ASA r2 processes certain REST API requests. For example, a DELETE request with the endpoint `/api/storage/luns/{uuid}` is processed asynchronously.

Enhanced security with OAuth 2.0

OAuth 2.0 is the industry standard authorization framework. It's used to restrict and control access to protected resources based on signed access tokens. You can configure OAuth 2.0 using System Manager to protect ASA r2 system resources.

After OAuth 2.0 is set up with System Manager, access by the REST API clients can be controlled. You need to first obtain an access token from an authorization server. The REST client then passes the token to the ASA r2 cluster as a bearer token using the HTTP authorization request header. See [Authentication and authorization using OAuth 2.0](#) for more information.

Access the ASA r2 API reference documentation through the Swagger UI

You can access the REST API reference documentation through the Swagger UI at your ASA r2 system.

About this task

You should access the ASA r2 reference documentation page for details about the REST API. As part of this, you can search for the string **Platform Specifics** to find details about ASA r2 system support for the API calls and properties.

Before you begin

You must have the following:

- IP address or host name of the ASA r2 system's cluster management LIF
- User name and password for an account with authority to access the REST API

Steps

1. Type the URL in your browser and press **Enter**:

`https://<ip_address>/docs/api`

2. Sign in using your administrator account.

The ASA r2 API documentation page is displayed with the API calls organized in major resource categories.

3. To see an example of an API call that's specifically applicable only to ASA r2 systems, scroll down to the **SAN** category and click **GET /storage/storage-units**.

Common ONTAP features supported on ASA r2 systems

Because ASA r2 systems run a streamlined version of ONTAP, many common ONTAP tasks and System Manager functions are performed the same way on ASA r2 systems as on other ONTAP systems.

For more information about common features and functions, see the following ONTAP documentation.

Data protection

Learn more about common data protection features supported on ASA r2 systems.

Clustered external key servers

You can configure connectivity to clustered external key management servers on a storage VM. With clustered key servers, you can designate primary and secondary key servers on a storage VM. When registering keys, ONTAP will first attempt to access a primary key server before sequentially attempting to access secondary servers until the operation completes successfully, preventing duplication of keys.

[Learn to configure clustered external key servers.](#)

External key management for encryption at rest

You can use one or more KMIP servers to secure the keys the cluster uses to access encrypted data.

- [Enable external key management.](#)
- [Enable external key management \(NVE\).](#)

Data security

Learn more about common data security features supported on ASA r2 systems.

Administrator access management

The role assigned to an administrator determines which functions the administrator can perform. Predefined roles for cluster administrators and storage VM administrators are provided by System Manager. You assign the role when you create the administrator's account, or you can assign a different role later.

- [Learn to manage administrator access with System Manager.](#)

Client authentication and authorization

ONTAP uses standard methods to secure client and administrator access to storage and to protect against viruses. Advanced technologies are available for encryption of data at rest and for WORM storage. ONTAP authenticates a client machine and user by verifying their identities with a trusted source. ONTAP authorizes a user to access a file or directory by comparing the user's credentials with the permissions configured on the file or directory.

[Learn about client authentication and authorization.](#)

OAuth 2.0 authentication

You can use the Open Authorization (OAuth 2.0) framework to control access to your ONTAP clusters. OAuth 2.0 restricts and controls access to protected resources using signed access tokens.

[Learn about OAuth 2.0 authentication.](#)

SAML authentication and administrator access

You can configure and enable Security Assertion Markup Language (SAML) authentication for web services. SAML authenticates users by an external Identity Provider (IdP) instead of the directory service providers such as Active Directory and LDAP.

[Learn to Configure SAML authentication.](#)

Networking

Learn more about common networking features supported on ASA r2 systems.

FIPS compliance

ONTAP is compliant in the Federal Information Processing Standards (FIPS) 140-2 for all SSL connections. You can turn SSL FIPS mode on and off, set SSL protocols globally, and turn off any weak ciphers such as RC4 within ONTAP.

Beginning with ONTAP 9.18.1 postquantum computing cryptographic algorithms are supported for SSL. These algorithms provide additional protection against potential future quantum computing attacks, and are available when SSL FIPS mode is disabled.

- [Learn to configure FIPS for all SSL connections.](#)

Link Aggregation Groups (LAGs)

An interface group, also known as a Link Aggregation Group (LAG), is created by combining two or more physical ports on the same node into a single logical port. The logical port provides increased resiliency, increased availability, and load sharing.

[Learn about Link Aggregation Groups.](#)

SAN Protocols

ASA r2 systems support all SAN protocols (iSCSI, FC, NVMe/FC, NVMe/TCP).

- [Learn more about the iSCSI protocol.](#)
- [Learn more about the Fibre Channel \(FC\) protocol.](#)
- [Learn about the NVMe protocol .](#)
 - [Learn to configure NVMe copy offload.](#)

Beginning with ONTAP 9.18.1, NVMe copy offload is supported. NVMe copy offload enables an NVMe host to offload copy operations from its CPU to the CPU of the ONTAP storage controller. The host can copy data from one NVMe namespace to another while reserving its CPU resources for application workloads.

- [Learn more about space allocation \(unmap\) for NVMe.](#)

Beginning with ONTAP 9.16.1, space deallocation (also called “hole punching” and “unmap”) is enabled for NVMe namespaces by default. Space deallocation allows a host to deallocate unused blocks from namespaces to reclaim space.

System Manager

You can search for various actions, objects, and information topics in System Manager. You can also search table data for specific entries.

[Learn to search, filter and sort information in System Manager.](#)

Get help

Manage AutoSupport on ASA r2 storage systems

AutoSupport is a mechanism that proactively monitors the health of your system and automatically sends messages to NetApp technical support, your internal support organization, and a support partner.

AutoSupport messages to technical support are enabled by default when you set up your cluster. You must set the correct options and have a valid mail host to have messages sent to your internal support organization. ONTAP begins sending AutoSupport messages 24 hours after it is enabled.

Before you begin

You must be a cluster administrator to manage AutoSupport.

Test AutoSupport connectivity

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

Steps

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

What's next?

You have confirmed that technical support can receive AutoSupport messages from your ASA r2 system, ensuring they have the necessary data to assist you in case of any issues.

Add AutoSupport recipients

Add members of your internal support organization to the list of email addresses that receive AutoSupport messages.

Steps

1. In System Manager, select **Cluster >Settings**.
2. Next to **AutoSupport** select ; then select **More options**.
3. Next to **Email**, select ; then select .
4. Enter the email address for the recipient; then the the recipient category.

For partners, select **Partner** for the recipient category. Select **General** for members of your internal support organization.

5. Select save.

What's next?

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

Send AutoSupport data

Should a problem occur on your ASA r2 system, AutoSupport data can significantly decrease the time it takes to identify and resolve issues.

Steps

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

What's next?

Your AutoSupport data is sent to technical support.

Suppress support case generation

If you are performing an upgrade or maintenance on your ASA r2 system, you might want to suppress AutoSupport generation of support cases until your upgrade or maintenance is complete.

Steps

1. In System Manager, select **Cluster >Settings**.
2. Next to **AutoSupport** select ; then select **Suppress support case generation**.
3. Specify the number of hours to suppress the generation of support cases; then select the nodes for which you do not want cases generated.
4. Select **Send**.

What's next?

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 >Settings**.
2. Next to **AutoSupport** select ; then select **Resume support case generation**.
3. Select the nodes for which you want to resume AutoSupport cases generated.
4. Select **Send**.

Result

AutoSupport cases are autogenerated for your ASA r2 system as needed.

Submit and view support cases for ASA r2 storage systems

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 closed or in progress.

You must be [registered with Active IQ](#) to view support cases for your ASA r2 system.

Steps

1. To submit a support case, in System Manager, select **Cluster >Support**; then select **Go to NetApp Support**.
2. To view a previously submitted case, in System Manager, select **Cluster >Support**; then select **View my cases**.

Legal notices

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.

ONTAP

Notice for ONTAP 9.16.1

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.