

A blue sphere is positioned on a yellow geometric structure composed of interlocking cubes, creating a stepped, architectural effect. The lighting is soft, casting gentle shadows and highlighting the smooth surface of the sphere and the sharp edges of the cubes.

NETAPP LEARNING SERVICES

# ONTAP Cluster Administration

Student Guide  
Content Version 11





NetApp Learning Services

# ONTAP Cluster Administration

Student Guide

Course ID: STRSW-ILT-ONTAPADM

Catalog Number: STRSW-ILT-ONTAPADM-SG

## **ATTENTION**

The information contained in this course is intended only for training. This course contains information and activities that, while beneficial for the purposes of training in a closed, non-production environment, can result in downtime or other severe consequences in a production environment. This course material is not a technical reference and should not, under any circumstances, be used in production environments. To obtain reference materials, refer to the NetApp product documentation that is located at <http://mysupport.netapp.com/>.

## **COPYRIGHT**

© 2023 NetApp, Inc. All rights reserved. Printed in the U.S.A. Specifications subject to change without notice.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of NetApp, Inc.

## **U.S. GOVERNMENT RIGHTS**

Commercial Computer Software. Government users are subject to the NetApp, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

## **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.

# Table of Contents

<b>Welcome.....</b>	<b>1</b>
<b>Module 1: NetApp ONTAP 9 Clusters.....</b>	<b>1-1</b>
<b>Module 2: Cluster Setup.....</b>	<b>2-1</b>
<b>Module 3: Cluster Management.....</b>	<b>3-1</b>
<b>Module 4: Network Management.....</b>	<b>4-1</b>
<b>Module 5: Physical Storage Management.....</b>	<b>5-1</b>
<b>Module 6: Logical Storage Management.....</b>	<b>6-1</b>
<b>Module 7: Data Access.....</b>	<b>7-1</b>
<b>Module 8: Data Protection.....</b>	<b>8-1</b>
<b>Module 9: Storage Efficiency.....</b>	<b>9-1</b>
<b>Module 10: Cluster Maintenance.....</b>	<b>10-1</b>



# ONTAP Cluster Administration



1 © 2023 NetApp, Inc. All rights reserved.

# Welcome

## Getting Started

- Schedule (start, stop, breaks, breakout sessions)
- Activities and participation
- Materials
- Equipment check
- Support

## Classroom Sessions

- Sign-in sheet
- Refreshments
- Phones to vibrate
- Alarm signal
- Evacuation procedure
- Electrical safety

## Virtual Sessions


- Collaboration tools
- Ground rules
- Phones and headsets


Set your phone to vibrate to prevent disturbing your fellow students. NetApp realizes that work does not always stop when in training. If you need to take a call, feel free to step outside of the classroom.

## Instructions

Tell everyone the following:

- Your name
- The company you work for
- Your storage administration experience level

 3 © 2023 NetApp, Inc. All rights reserved.



Take time to meet and get to know one another. If you are participating in a NetApp Virtual Live class, your instructor asks you to use the chat window or a conference connection to speak. If you are using a conference connection, unmute your line to speak and be sure to mute again after you speak.

## Which tasks does a data administrator perform?

### Make storage accessible

- Manage network and data access protocols
- Provision storage and apply access permissions
- Load balance and prioritize network I/O

### Protect the data

- Manage Snapshot schedules
- Configure SnapMirror replication
- Maintain disaster recovery storage systems
- Manage administrator access
- Perform preventive maintenance

### Maximize use of the data

- Monitor and manage capacity use and future requirements
- Create test and development copies of data
- Provide reports to end users
- Optimize and manage quality-of-service (QoS) settings
- Assist end users with application integration

Depending on the size of your environment, you might perform all three tasks, or you might grow to specialize in only one.

Before you see what this course is about, you should first understand the duties of a data administrator, to put the modules of the course into context.

The first duty of a data administrator is to make storage accessible. This duty is mostly associated with networking.

The second duty is to protect data from loss or corruption.

The third duty is to help users and applications obtain the most value from the data that is stored in the cluster.

In a small organization, you might be responsible for all these duties. In a large organization, you might become a specialist in only one of these duties.



## About this course

This course focuses on enabling you to do the following:

- Define NetApp ONTAP cluster components
- Describe the role of a storage VM (storage virtual machine, also known as SVM) in the NetApp storage architecture
- Configure an ONTAP cluster
  - Configure and manage storage resources
  - Configure and manage networking resources
  - Create and configure a storage VM
  - Create, manage, and protect FlexVol volumes
  - Implement storage efficiency features
- Manage ONTAP administrator access and user accounts
- Maintain NetApp storage systems

During the next three days, you want to learn as many of the tasks and procedures for administering a NetApp ONTAP cluster as possible. Unfortunately, this condition means that you cannot go as technically deep as you might want. This course is a launching platform for continued training to master your skills.

# ONTAP 9

## Foundational

ONTAP Cluster Fundamentals

ONTAP NAS Fundamentals

ONTAP SAN Fundamentals

ONTAP Data Protection Fundamentals

## Intermediate

ONTAP Cluster Administration

ONTAP NFS Administration

ONTAP SMB Administration

ONTAP SAN Administration

ONTAP Data Protection Administration

ONTAP Compliance Solutions Administration

## What this course is not

- An ONTAP 9 fundamentals course

You are expected to be familiar with the core concepts and functionality of ONTAP 9 software.

- A certification test preparation course
- A hardware installation and configuration course
- A data protection course
- A performance troubleshooting course
- A NAS or SAN administration course

This course assumes that, although you might be new to *ONTAP Cluster Administration*, you have taken the prerequisite training to learn about NetApp ONTAP 9 software.

Although this course is recommended training for taking NetApp certification exams, it is not an exam preparation course.

Although you learn some hardware installation basics, if you want to learn how to physically install cluster hardware, you should take online courses like *Universal FAS Installation* and the model-specific installation courses.

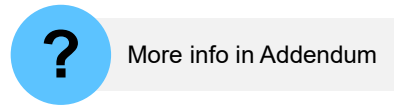
This course is often paired with *ONTAP Data Protection Administration*, so the course assumes that you will take the latter course at the end of the week.

Performance is a complicated topic with many variables, so this course covers only some recommended practices to keep cluster performance stable under general use.

This course enables you to set up basic sharing of NAS and SAN data. For advanced uses, take the protocol-specific courses.

## Student guide addendums

To maximize the time for exercises, not all the information that you will cover is included in the lecture portion. Whenever you see the following graphic, you can find more information about the topic in the module addendum in your Student Guide.



Although it is necessary to omit some content when developing a course to fit the allotted training time, you should still find this course useful. Content that was not included in the lecture can be found in addendums at the end of each module in your Student Guide. You can identify this content by the graphic with the question mark on the final slide in the lecture that covers the topic.




## **ACTION: Take the pre-class assessment**

A short quiz

- **Assessment link:**  
[ONTAP Cluster Administration Assessment](#)
- **Submit your answers.**
- **Observe your baseline score.**

This assessment requires approximately  
**7 minutes.**

 9 © 2023 NetApp, Inc. All rights reserved.

To measure your current knowledge of course topics, take the pre-class assessment by accessing the link that is provided. At the completion of the course, you can take the post-class assessment to measure how much you learned.

[ONTAP Cluster Administration pre-class assessment](#)

[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/leclassview/dowbt-0000386614](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/leclassview/dowbt-0000386614)

You will receive an acknowledgement of completion regardless of your score. You can retake the assessment as many times as you like.


## Course modules

### Day 1 Morning

- Introduction
- Module 1: NetApp ONTAP 9 clusters
- Module 2: Cluster setup

### Day 1 Afternoon

- Module 3: Cluster management
- Module 4: Network management

 10 © 2023 NetApp, Inc. All rights reserved.

This schedule is based on average completion times for modules. Each class is composed of students with different backgrounds and experience levels. This situation means that some modules might take more or less time to complete. Your instructor will adjust the schedule accordingly for breaks, meals, and the start time of each module.

## Course modules

### Day 2 Morning

- Day 1 review
- Module 5: Physical storage management
- Module 6: Logical storage management

### Day 2 Afternoon

- Module 7: Data access

## Course modules

### Day 3 Morning

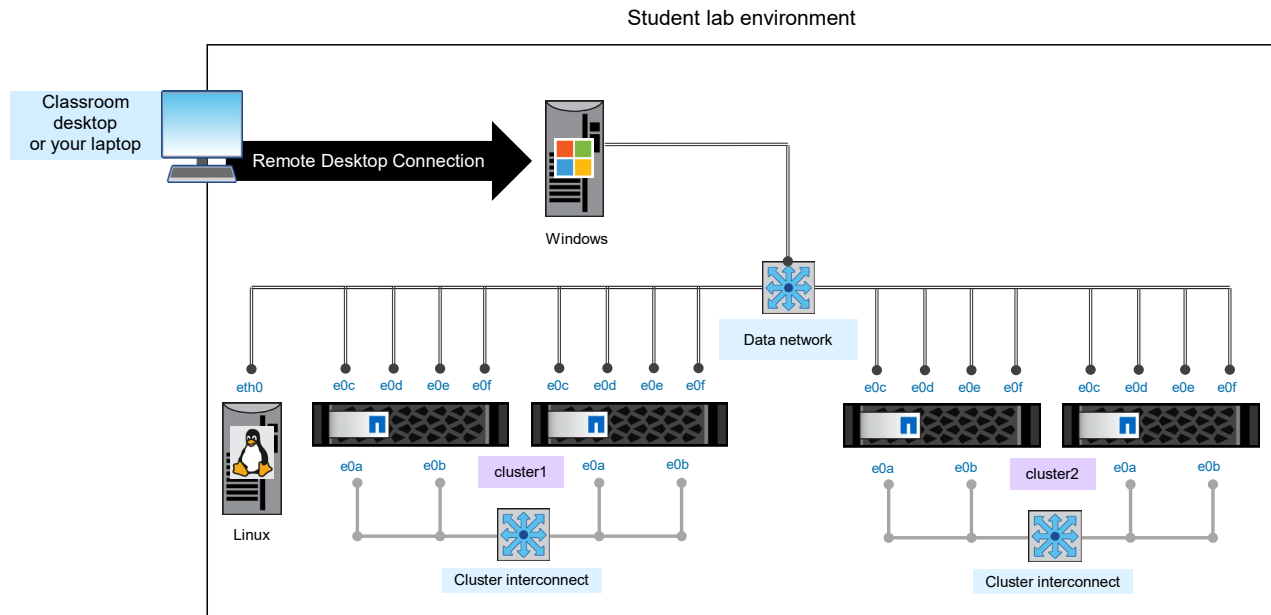
- Day 2 review
- Module 8: Data protection
- Module 9: Storage efficiency

### Day 3 Afternoon

- Module 10: Cluster maintenance
- Course review and post-class assessment



## Exercise equipment: Basic architecture



NetApp 13 © 2023 NetApp, Inc. All rights reserved.

Launch your exercise equipment kit from your laptop or from the classroom desktop. To connect to your exercise equipment, use Remote Desktop Connection or the NetApp Learning Services portal.

The Windows Server is your jump host to access the lab environment.

Your exercise equipment consists of several servers:

- A 2-node ONTAP 9.12 cluster
- A second 2-node ONTAP 9.12 cluster
- A CentOS Linux server



## Complete an exercise

Module 0: Welcome

### Checking the Exercise Equipment

- Access your lab equipment.
- Open your Exercise Guide, Module 0.
- Complete Exercise 1.
- Share your results.
- Report issues.

This exercise requires approximately  
**15 minutes.**

See the instructions in your Exercise Guide.




## Share your experiences

Roundtable questions for the  
equipment-based exercises

- Do you have questions about your lab equipment?
- Do you have an issue to report?

# Module 1

## NetApp ONTAP 9 clusters

 1 © 2023 NetApp, Inc. All rights reserved.

In this module, you learn how to configure key features of NetApp ONTAP software.

## About this module

This module focuses on enabling you to do the following:


- Identify NetApp ONTAP deployment options
- Define ONTAP cluster components
- Describe the role of a storage VM (storage virtual machine, also known as SVM) in the ONTAP storage architecture

This module covers NetApp ONTAP deployment options and cluster components for ONTAP software.

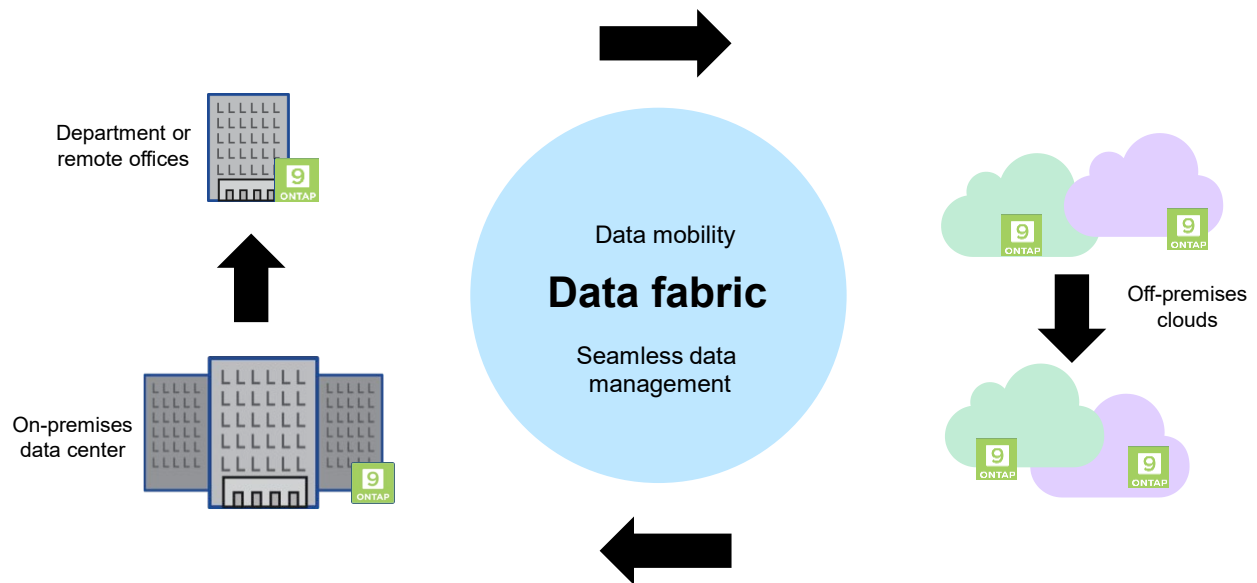


# Lesson 1

## ONTAP deployment options

 3 © 2023 NetApp, Inc. All rights reserved.

## NetApp ONTAP software: The foundation for your data fabric



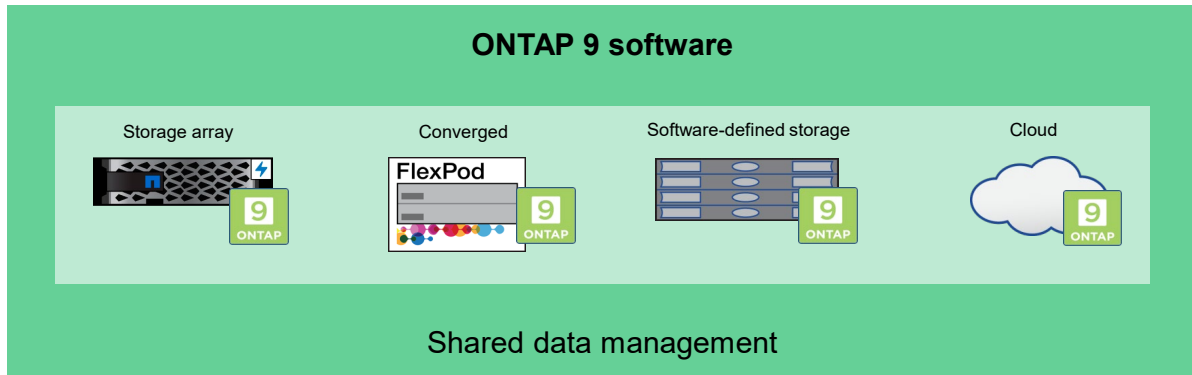
**NetApp** 4 © 2023 NetApp, Inc. All rights reserved.

The data fabric powered by NetApp weaves hybrid cloud mobility with uniform data management. NetApp works with new and existing partners to continually add to the fabric.

For more information about the data fabric, visit <https://www.netapp.com/data-fabric/what-is-data-fabric>.

## Standardize data management

ONTAP software: For any application, anywhere



NetApp 5 © 2023 NetApp, Inc. All rights reserved.

ONTAP 9 software has three major deployment options (ONTAP 9 software, NetApp ONTAP Select software, and NetApp Cloud Volumes ONTAP), which you can use in various environments. Simply put, "It is just ONTAP!"

Standardize data management:



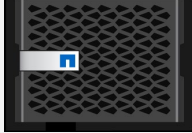

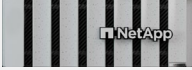




- Across architectures, blocks, or files, and on flash, disk, or cloud
- Across deployment models, from engineered storage arrays to commodity servers
- Across enterprise and emerging applications

Although this course focuses on physical ONTAP clusters, the knowledge also applies to Cloud Volumes ONTAP and ONTAP Select software.



# ONTAP 9.12 based hardware systems

March 2023

AFF A-Series	AFF C-Series	FAS
<b>Enterprise level</b>  <ul style="list-style-type: none"><li>AFF A900</li><li>AFF A800</li><li>AFF A700</li></ul>	<b>Enterprise level</b>  <ul style="list-style-type: none"><li>AFF C800</li></ul>	<b>Enterprise level</b>  <ul style="list-style-type: none"><li>FAS9500</li><li>FAS8700</li></ul>
<b>Mid level</b>  <ul style="list-style-type: none"><li>AFF A400</li></ul>	<b>Mid level</b>  <ul style="list-style-type: none"><li>AFF C400</li></ul>	<b>Mid level</b>  <ul style="list-style-type: none"><li>FAS8300</li></ul>
<b>Entry level</b>  <ul style="list-style-type: none"><li>AFF A250</li><li>AFF A150</li></ul>	<b>Entry level</b>  <ul style="list-style-type: none"><li>AFF C250</li></ul>	<b>Entry level</b>  <ul style="list-style-type: none"><li>FAS2750</li><li>FAS2720</li></ul>

For product specifications, see the Hardware Universe: [hwu.netapp.com](http://hwu.netapp.com).

NetApp 6 © 2023 NetApp, Inc. All rights reserved.

NetApp has a storage system to support the performance needs and budget needs of all customers. The NetApp FAS systems provide a unified storage platform with only HDDs (hard disk drives) or a hybrid configuration with both HDDs and SSDs (solid state drives). The NetApp AFF systems contain only high-speed SSDs to support mission-critical applications with sub-millisecond access times. The NetApp C-Series systems use SSDs with quad-level cell (QLC) flash technology to support tier 1 and tier 2 applications that do not require sub-millisecond performance.

For more detailed information about the supported storage systems for ONTAP 9 software, see the Hardware Universe at <http://hwu.netapp.com>.

## All SAN Array systems

- All SAN Array (ASA) systems are optimized for enterprise tier 1 SAN workloads.
- ASA systems use proven AFF system hardware.
- ASA systems provide uninterrupted access to data during a planned or unplanned storage failover.
- NetApp ONTAP System Manager is streamlined for SAN:
  - Implementation
  - Configuration
  - Management




With ONTAP 9.7 software, NetApp introduced the NetApp All SAN Array (ASA) systems. NetApp ASA systems build on the all-flash AFF systems to deliver continuous SAN availability for enterprises that run mission-critical applications. These systems provide streamlined implementation, configuration, and management through a solution that is dedicated to running only tier 1 SAN workloads.




## Lesson 2

### The ONTAP cluster

 8 © 2023 NetApp, Inc. All rights reserved.

## The cluster



For cluster specifications and node mixing rules, see the Hardware Universe:  
[hwu.netapp.com](http://hwu.netapp.com).

**NetApp** 9 © 2023 NetApp, Inc. All rights reserved.

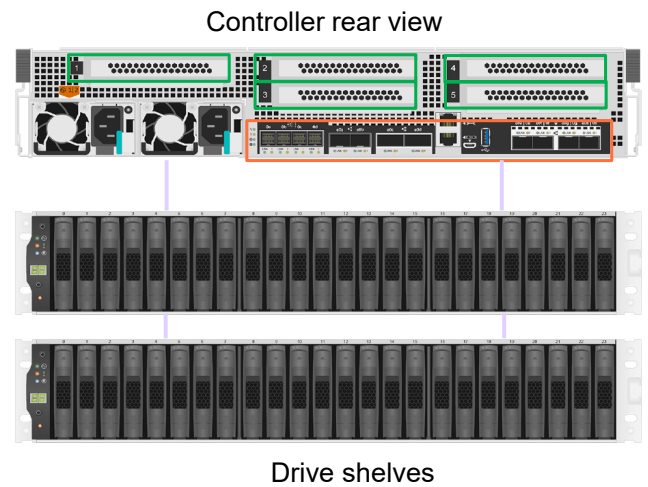
You might wonder, “What *is* a cluster?” This course examines cluster components individually, but first, consider a high-level view.

A cluster is one or more FAS or AFF controllers that run ONTAP software. In ONTAP terminology, a controller is called a node. In clusters with more than one node, a cluster interconnect is required so that the nodes appear as one cluster.

A cluster can be a mix of FAS and AFF models, depending on the workload requirements. Nodes can be added to or removed from a cluster as workload requirements change. For more information about the number and types of nodes, see the Hardware Universe at <http://hwu.netapp.com/>.

## A cluster node

- A FAS, AFF, or ASA storage controller that runs ONTAP software
- Storage and network ports
- Expansion slots  
Not all entry-level systems have expansion slots.
- Nonvolatile write protection that uses NVRAM, NVMEM, or NVDIMMs
- Drive shelves or internal drives or both




For product specifications, see the Hardware Universe:  
[hwu.netapp.com](http://hwu.netapp.com).

For information about specific controller models, see the product documentation on the NetApp Support Site, or see the Hardware Universe at <http://hwu.netapp.com/>.

## HA pairs

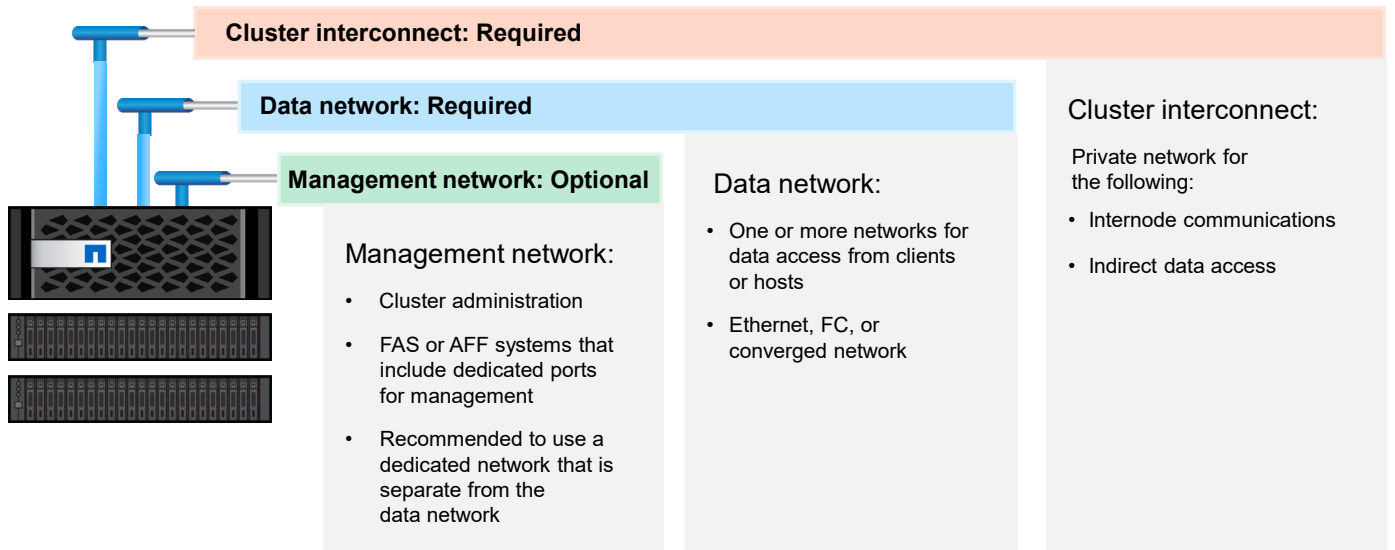
- Characteristics of a high-availability (HA) pair:
  - Two connected nodes are in a partnership.
  - Nodes connect to the same drive shelves.
  - Nodes own some of the drives.
  - If a node fails, the surviving node takes control of the failed partner's drives.
- Components of HA pair connections:
  - HA interconnect
  - Multipath HA shelf connectivity



In multi-node clusters, high-availability (HA) pairs are used.

The controllers in the nodes of an HA pair connect either through an external HA interconnect, which consists of adapters and cables, or through an internal interconnect. The nodes must use redundant paths to connect to the shared drive shelves.

# Network



NetApp 12 © 2023 NetApp, Inc. All rights reserved.

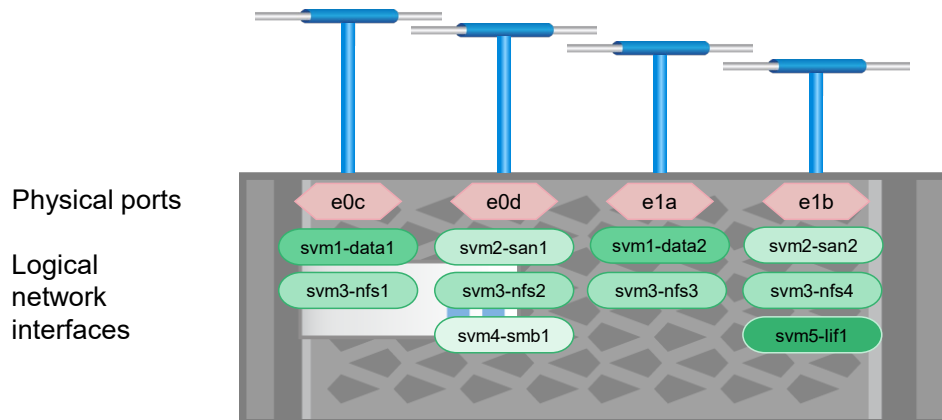
Clusters require two or more networks, depending on the environment.

The nodes communicate with each other over a cluster interconnect, even when the cluster is composed of only one HA pair. In a 2-node cluster, the interconnect can be switchless. Clusters with more than two nodes require a private cluster interconnect that uses switches.

Clients and hosts access data through the data network. The data network includes one or more networks that are primarily used for data access by clients or hosts. Depending on the environment, there might be Ethernet, FC, or converged networks. Data networks can consist of one or more switches or even redundant networks.

The management network is for cluster administration. Redundant connections to the management ports on each node should be provided from the management network. In smaller environments, the management and data networks might be on a shared Ethernet network.

## Ports and LIFs



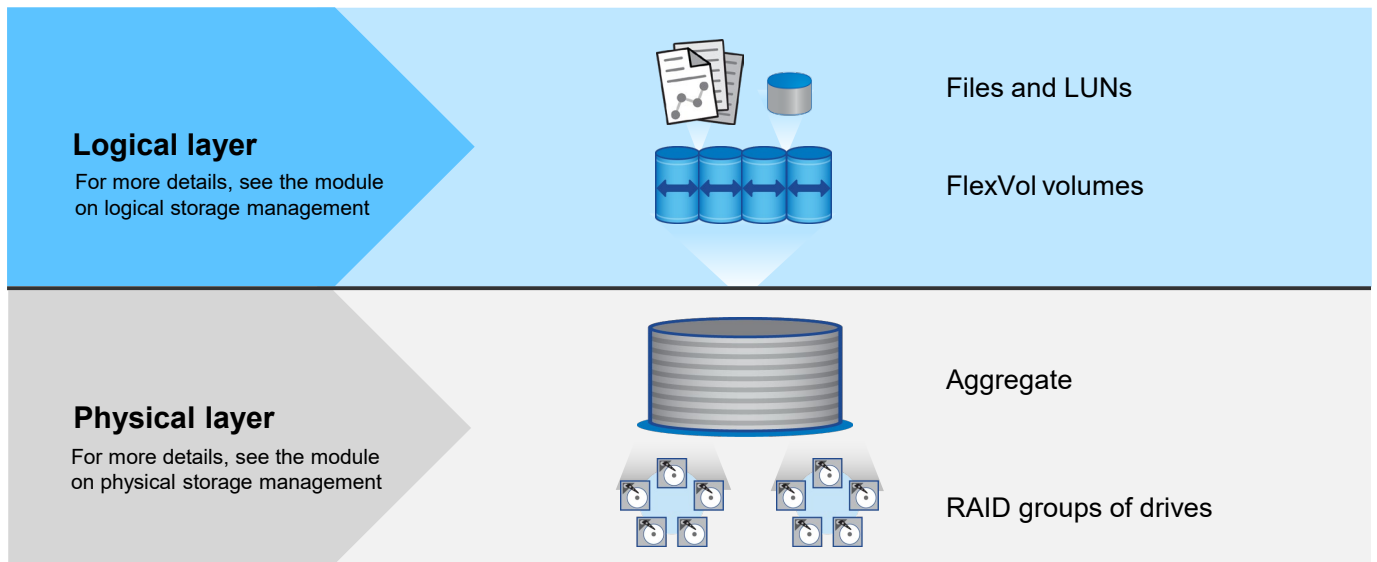
NetApp 13 © 2023 NetApp, Inc. All rights reserved.

[0] Nodes have various physical ports that are available for cluster, management, and data traffic. The ports must be configured appropriately for the environment.

[1] A LIF represents a network access point to a storage VM (storage virtual machine, also known as SVM) in the cluster. A LIF belongs to a storage VM and is associated with a physical port. A physical port can host multiple LIFs that belong to multiple storage VMs. A storage VM likely owns multiple LIFs, with each LIF associated with a different physical port on possibly different cluster nodes.



## ONTAP storage architecture



NetApp 14 © 2023 NetApp, Inc. All rights reserved.

The ONTAP storage architecture dynamically maps physical storage resources to logical containers.

The physical storage layer includes the internal and external drives to which the data is written and from which the data is read.

In ONTAP software, drives are grouped into RAID groups. An aggregate is a collection of physical drive space that contains one or more RAID groups. Each aggregate has a RAID configuration and a set of assigned drives.


The physical storage layer is covered in detail in a different module.

Within each aggregate, you can create one or more FlexVol volumes. A FlexVol volume is an allocation of drive space that is a portion of the available space in the aggregate. A FlexVol volume can contain files, LUNs, or NVMe namespaces. The FlexVol volumes, files, LUNs, and namespaces make up the logical storage layer.


The logical storage layer is covered in detail in a different module.



## Lesson 3 Storage VMs

 15 © 2023 NetApp, Inc. All rights reserved.

## Storage virtual machine



NetApp 16 © 2023 NetApp, Inc. All rights reserved.

You use storage VMs to serve data to clients and hosts. The storage VM behaves like a physical storage system. Users are unaware that they are accessing a virtual storage system. Like a virtual machine running on a hypervisor, a storage VM is a logical entity that abstracts physical resources. Data that is accessed through the storage VM is not bound to a location in storage. Network access to the storage VM is not bound to a physical port.

A storage VM serves data to clients and hosts from one or more volumes through one or more network LIFs. Volumes can be assigned to any data aggregate in the cluster. LIFs can be hosted by any physical or logical port. Both volumes and LIFs can be moved without disrupting data service, whether you are performing hardware upgrades, adding nodes, balancing performance, or optimizing capacity across aggregates.

The same storage VM can have a LIF for NAS traffic and a LIF for SAN traffic. Clients and hosts require only the address of the LIF (IP address for NFS, SMB, S3, or iSCSI; worldwide port name [WWPN] for FC) to access the storage VM. LIFs keep their addresses as they move. Ports can host multiple LIFs.

An ONTAP cluster can host multiple storage VMs. Each storage VM has its own security, administration, and namespace. When a storage VM is created, a root volume is also created. The root volume serves as the NAS client entry point to the namespace that the storage VM provides. NAS client data access depends on the health of the root volume in the namespace. SAN client data access is independent of the root volume health in the namespace.

## Storage VM types

```
svl-nau::> vservers show
```

Vserver	Type	Subtype	Admin State	Operational State	Root Volume	Aggregate
cluster1	admin	-	-	-	-	-
cluster1-01	node	-	-	-	-	-
cluster1-02	node	-	-	-	-	-
svm_1	data	default	running	running	svm_1_root	n1_data_1
svm_2	data	default	running	running	svm_2_root	n1_data_1
svm_3	data	default	running	running	svm_3_root	n1_data_2

6 entries were displayed.

### Admin storage VM:

- Is created during cluster setup
- Represents the cluster
- Exists one time per cluster
- Owns cluster-scoped resources

### Node storage VM:

- Is created during cluster setup
- Represents an individual node
- Exists one time per node in the cluster
- Owns node-scoped resources


### Data storage VM:

- Provides client access to user data
- Includes data volumes, LIFs, protocols, and access control
- Is for multiple use cases:
  - Secure multitenancy
  - Separation of resources and workloads
  - Delegation of management


**NetApp** 17 © 2023 NetApp, Inc. All rights reserved.

A data storage VM contains data volumes and LIFs that serve data to clients. Unless otherwise specified, the term “storage VM” refers to a data storage VM. In the CLI, storage VMs are displayed as “Vservers.” Storage VMs might have one or more FlexVol volumes or scalable NetApp ONTAP FlexGroup volumes.

## Storage VM benefits



- **Unified storage:**
  - NAS protocols: CIFS and NFS
  - SAN protocols: iSCSI, FC (including FCoE), and NVMe over Fibre Channel (NVMe/FC)
  - Object protocols: Simple Storage Service (S3)



- **Secure multitenancy:**
  - Partitioning of a storage system
  - Isolation of data and management
  - No data flow among storage VMs in the cluster




- **Nondisruptive operations (NDO) and nondisruptive upgrade (NDU):**
  - Resource migration
  - Resource availability during hardware and software upgrades



- **Delegation of management:**
  - User authentication and administrator authentication
  - Access that is assigned by the cluster administrator



- **Scalability:**
  - Addition and removal
  - Modification on demand to meet data-throughput and storage requirements



- **Disaster resiliency:**
  - Storage VMs can be copied or moved to other ONTAP clusters.
  - Storage VM failover protects against regional disasters.

Storage VMs provide many benefits.

The first benefit is unified storage. Storage VMs can serve data concurrently through multiple data access protocols. Storage VMs with FlexVol volumes provide file-level data access through NAS protocols, such as CIFS and NFS. They provide block-level data access through SAN protocols, such as iSCSI, FC, or FCoE. Storage VMs with FlexVol volumes can serve data to SAN and NAS clients independently at the same time.

Another benefit is nondisruptive operations (NDO). Storage VMs can operate continuously and nondisruptively. By enabling resources such as volumes and LIFs to move to other nodes, storage VMs help clusters to operate continuously. Continuous operations are advantageous during software and hardware upgrades, the addition and removal of nodes, and all administrative operations.

A third benefit of storage VMs is scalability. Storage VMs can be added, removed, or given more resources as the underlying physical storage grows. Storage VMs can be modified on demand to meet data-throughput requirements and other storage requirements.

Storage VMs are the fundamental unit of secure multitenancy. Storage VMs enable partitioning of the storage infrastructure so that the infrastructure appears as multiple independent storage systems. Partitions isolate data and management. Each storage VM appears as a single independent server, which enables multiple storage VMs to coexist in a cluster and prevents data from flowing among storage VMs.


Storage VMs support delegation of management. Each storage VM can have its own user authentication and administrator authentication. Storage VM administrators can manage the storage VMs that they are authorized to access. Cluster administrators assign privileges to storage VM administrators.

Finally, storage VMs enhance disaster resiliency. Storage VMs can be migrated or replicated to other ONTAP clusters. All the storage VM configuration information can be included, or parts excluded so that the replicant storage VM can take over in a site failure.



## Lesson 4

# Software-defined storage


 19 © 2023 NetApp, Inc. All rights reserved.

## Software-defined storage

NetApp ONTAP Select software

### ONTAP Select software

- Software-defined storage on third-party servers that is referred to as hybrid cloud infrastructure
- Suited for data center or remote office
- Flexible, capacity-based license



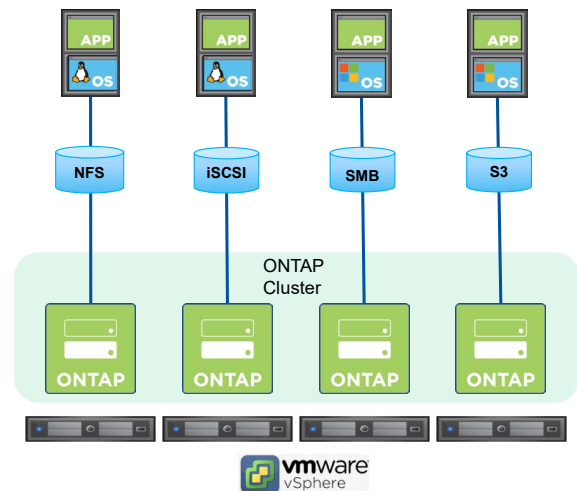
**NetApp** 20 © 2023 NetApp, Inc. All rights reserved.

NetApp ONTAP Select software is ONTAP software that runs on commodity third-party hardware. ONTAP Select software is a core component of the NetApp software-defined storage solutions for on-premises storage needs.

# ONTAP Select software

## Overview

- ONTAP Select software features the following:
  - ONTAP software that runs on commodity hardware  
**Note:** ONTAP Select clusters cannot be mixed with AFF or FAS nodes in a cluster.
  - Enterprise data management services for server direct-attached storage (DAS), external array, and VMware vSAN
- ONTAP Select software provides a cloudlike experience on premises:
  - Flexibility
  - Agility
  - Simplicity



[?](#) More info in Addendum

ONTAP Select software is ONTAP software that runs on commodity hardware (other-vendor hardware).

ONTAP Select software has all the benefits of ONTAP software: cluster-wide namespace, volume moves, workload rebalancing, nondisruptive upgrade (NDU), and nondisruptive operations (NDO).

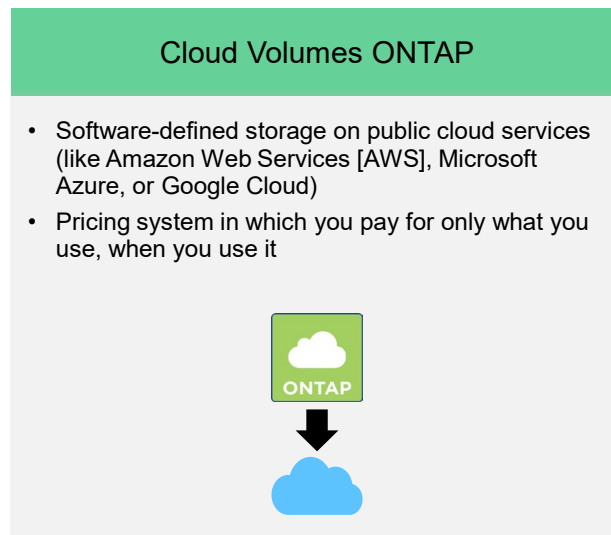
The HA architecture used with ONTAP Select software is based on a non-shared storage model. One node in an HA pair cannot directly access the storage owned by the other node. This design can affect certain ONTAP Select operational characteristics.


**Note:** ONTAP Select clusters cannot be mixed with FAS nodes or clusters.



## Software-defined storage

NetApp Cloud Volumes ONTAP software



 22 © 2023 NetApp, Inc. All rights reserved.

Cloud Volumes ONTAP (CVO) enhances the use of cloud storage resources to deliver consistent, uninterrupted, multiprotocol access across an entire organization. Cloud Volumes ONTAP offers built-in enterprise data management services, cloud data protection, security, and cost efficiency. Cloud Volumes ONTAP provides centrally managed, high-performance file, block, and object storage with enhanced resilience, delivered natively on AWS, Google Cloud, and Azure. Customers can deploy NetApp Cloud Volumes ONTAP systems in public clouds using a “pay-as-you-go” or a “bring your own” ONTAP license model.

NetApp Blue XP enables you to centrally manage your hybrid multi-cloud infrastructure keeping you in control of your data no matter where it is.

You can use Blue XP to set up and use Cloud Volumes ONTAP for efficient, multi-protocol data management across clouds.

You can use Blue XP to set up and use each of the file-storage services.

You can use Blue XP to discover and manage your on-prem ONTAP clusters by creating volumes, backing up to the cloud, replicating data across your hybrid cloud, and tiering cold data to the cloud.

You can also use Blue XP to enable integrated cloud services such as:

- [Cloud Data Sense](#)
- [Cloud Insights](#)
- [Cloud Backup](#)

# Software-defined storage

## NetApp public cloud software ecosystem

### Microsoft Azure >

Azure NetApp Files  
Cloud Manager  
Cloud Volumes ONTAP  
Cloud Data Sense  
Astra  
Cloud Insights  
Cloud Sync  
Cloud Backup  
Global File Cache  
Virtual Desktop Service  
Cloud Analyzer  
Elastigroup  
Ocean  
Eco

### Google Cloud >

Cloud Volumes Service for Google Cloud  
Cloud Manager  
Cloud Volumes ONTAP  
Cloud Data Sense  
Astra  
Cloud Insights  
Cloud Sync  
Cloud Backup  
Global File Cache  
Virtual Desktop Service  
Cloud Analyzer  
Elastigroup  
Ocean  
Eco

### aws >

Amazon FSx for NetApp ONTAP  
Cloud Volumes Service for AWS  
Cloud Manager  
Cloud Volumes ONTAP  
Cloud Data Sense  
Astra  
Cloud Insights  
Cloud Sync  
Cloud Backup  
Global File Cache  
Virtual Desktop Service  
Cloud Analyzer  
Elastigroup  
Ocean  
Eco



### Cloud Volumes ONTAP

- Self-service cloud storage with built-in protection, security, and cost efficiency
- For file (NFS,SMB), block (iSCSI), and object (S3) workloads



### Azure NetApp Files



### Cloud Volumes Service for AWS



### Cloud Volumes Service for Google Cloud

NetApp-managed file services for NFS and SMB workloads



### Amazon FSx for NetApp ONTAP

AWS-managed services for file (NFS,SMB) and block (iSCSI) workloads

Customers can launch and manage their own instances of Cloud Volumes ONTAP. Customers who do not want to manage ONTAP storage systems in the public cloud themselves can use NetApp Cloud Volumes Service. Cloud Volumes Service is a NetApp-managed, cloud-native file service that is available in the public clouds.

The Azure NetApp Files service is an enterprise-class, high-performance, metered file storage service that is highly available by default. You can select service and performance levels, create capacity pools and volumes, and manage data protection. Like other Azure services, Azure NetApp Files can be provisioned and consumed against an existing Azure agreement.

The NetApp-managed Cloud Volumes Service for Google Cloud is available from the Google Cloud marketplace. Likewise, the Cloud Volumes Service for AWS is available from the AWS marketplace.


Amazon FSx for NetApp ONTAP delivers NFS, SMB, and iSCSI storage that is tightly integrated within the AWS ecosystem. FSx for ONTAP is a native service that AWS fully manages, operates, and supports. FSx for ONTAP is easy to operate and requires no installation, no upgrading, and no maintenance. Choose from a variety of AWS and NetApp management tools, including the AWS Management Console and AWS CloudFormation.

You can use Blue XP to set up and use each of these file-storage services.

Visit NetApp Cloud Central at [cloud.netapp.com](https://cloud.netapp.com) to learn how NetApp can help you optimize the operations of your hybrid multicloud data center.

## References

- ONTAP 9 Documentation Center  
<http://docs.netapp.com/ontap-9>
- NetApp TV  
[Data ONTAP channel](https://www.netapp.tv/data-ontap)
- Cloud Central  
<http://cloud.netapp.com>
- Technical Reports  
[ONTAP AFF All SAN Array systems](https://www.netapp.com/media/10379-tr4515.pdf)

 24 © 2023 NetApp, Inc. All rights reserved.

ONTAP 9 Documentation Center: <http://docs.netapp.com/ontap-9/index.jsp>


NetApp Cloud Central: <http://cloud.netapp.com>

NetApp TV [Data ONTAP channel](https://www.netapp.tv/data-ontap): <https://www.netapp.tv/data-ontap>

Technical Reports: [ONTAP AFF All SAN Array systems](https://www.netapp.com/media/10379-tr4515.pdf): <https://www.netapp.com/media/10379-tr4515.pdf>

# Knowledge check

Module 1: NetApp ONTAP 9 clusters

 25 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

### Which two deployment options are software-defined? (Choose two.)

- a. ONTAP software that is deployed on a FAS system
- b. ONTAP software that is deployed on an AFF system
- c. ONTAP software that is deployed on commodity hardware
- d. ONTAP software that is deployed in the cloud
- e. ONTAP software that is deployed by using a heterogeneous enterprise array

## Knowledge check

### Which set of networks is part of a cluster?

- a. data network, management network, and cluster interconnect
- b. data network, high-availability (HA) interconnect, and cluster interconnect
- c. high-availability (HA) interconnect, cluster interconnect, and backup network
- d. data network, cluster interconnect, and backup network

## Knowledge check

**Which pair of components is a major part of data storage VMs?**

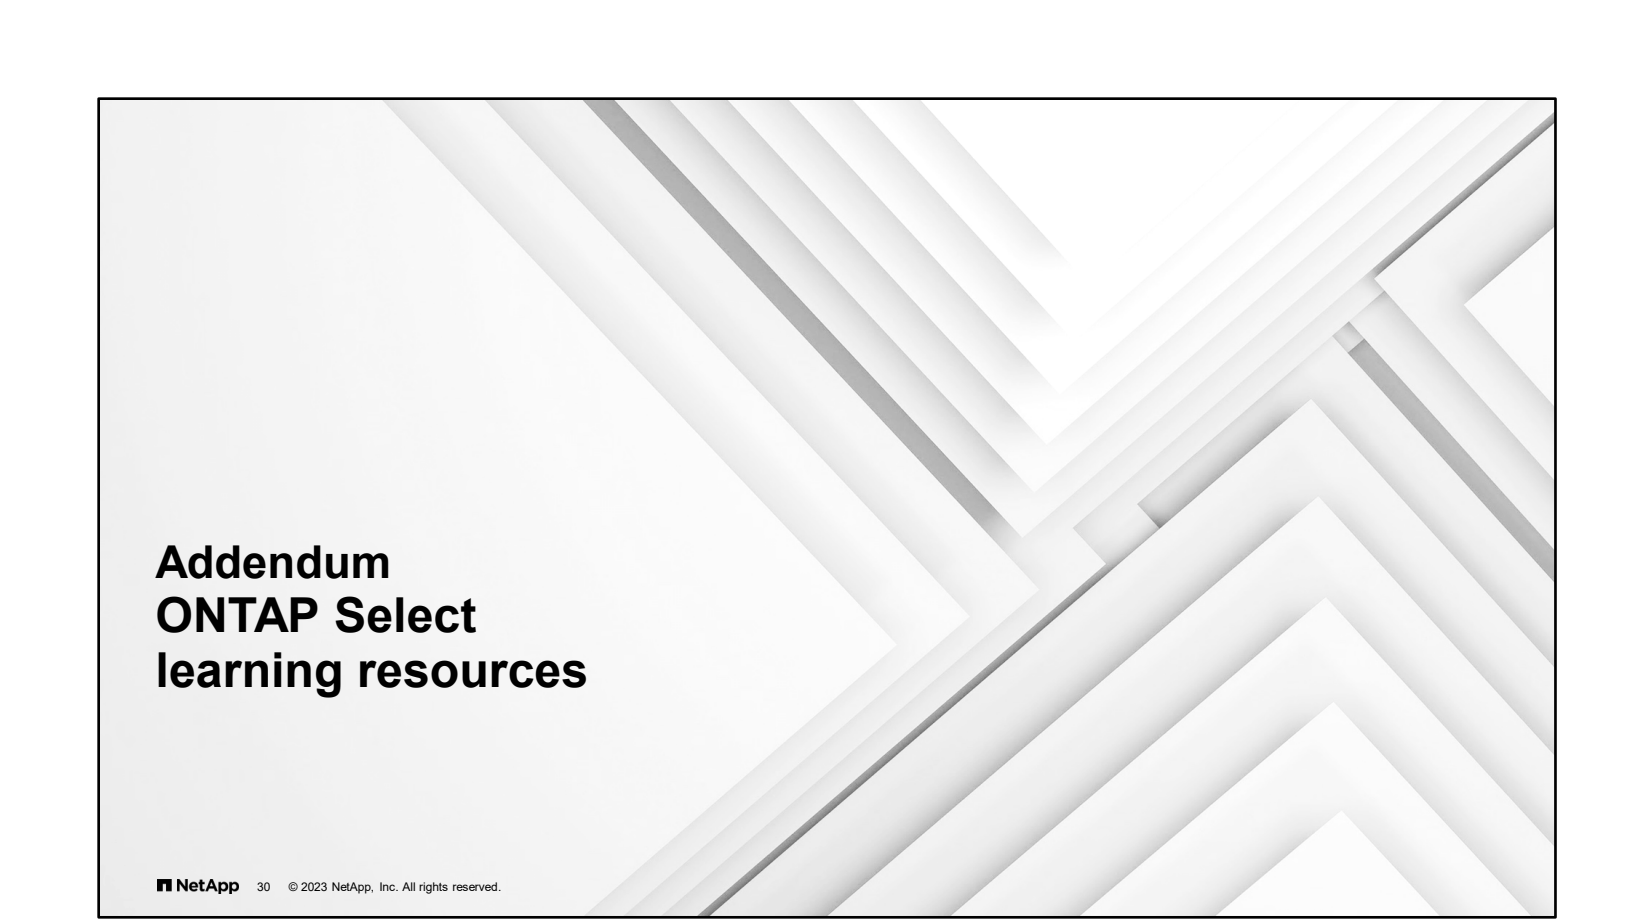
- a. aggregates and network ports
- b. disks and nodes
- c. data LIFs and aggregates
- d. volumes and data LIFs

## Module summary


This module focused on enabling you to do the following:

- Identify ONTAP deployment options
- Define ONTAP cluster components
- Describe the role of a storage VM in the ONTAP storage architecture





# Addendum ONTAP Select learning resources

 30 © 2023 NetApp, Inc. All rights reserved.

## More ONTAP Select learning resources

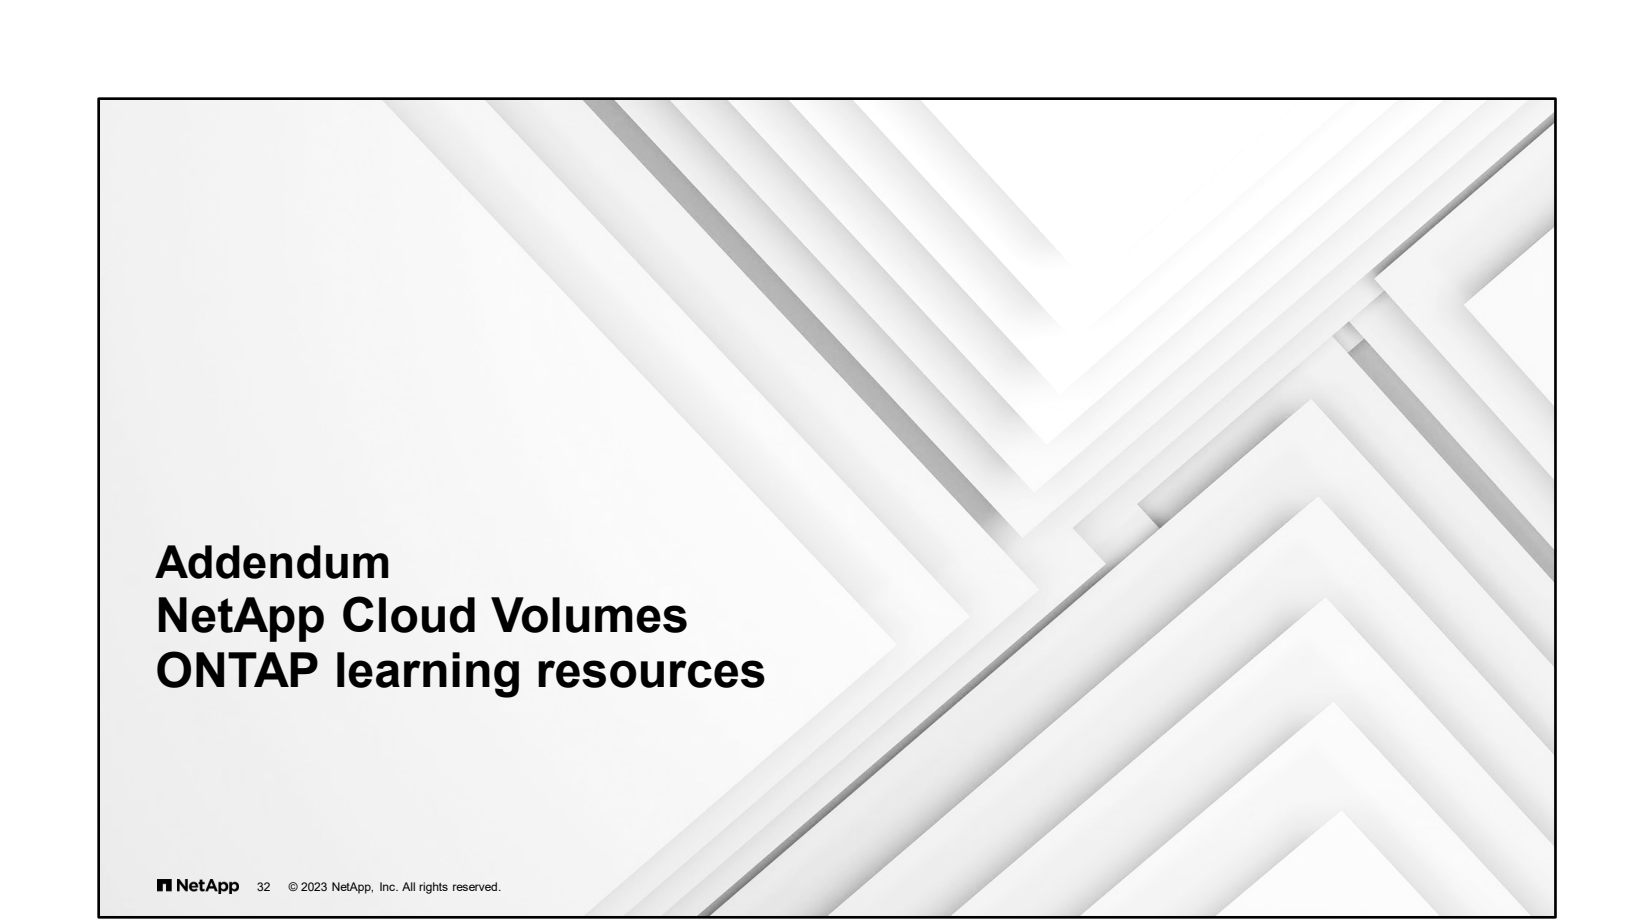
Learn about advanced topics like supported configurations and deploying the software on VMware ESXi or Kernel-Based Virtual Machine (KVM) hosts:

- *ONTAP Select Installation and Deployment* (web-based course)
- [ONTAP Select documentation resources](#)
- Technical Reports:
  - [TR-4517: ONTAP Select on VMware - Product Architecture and Best Practices](#)
  - [TR-4613: ONTAP Select on KVM: Product Architecture and Best Practices](#)


[ONTAP Select documentation resources](#)

[TR-4517: ONTAP Select on VMware - Product Architecture and Best Practices](#)


[TR-4613: ONTAP Select on KVM: Product Architecture and Best Practices](#)



# Addendum NetApp Cloud Volumes ONTAP learning resources

 32 © 2023 NetApp, Inc. All rights reserved.

## Cloud Volumes learning resources

 33 © 2023 NetApp, Inc. All rights reserved.

- [Cloud Volumes ONTAP Fundamentals](#) (online course)
- [Integrating Hybrid Clouds with Amazon Web Services](#) (virtual live course)
- [Integrating Hybrid Clouds with Microsoft Azure](#) (virtual live course)
- [Integrating Hybrid Clouds with Google Cloud](#) (virtual live course)
- [Cloud Volumes ONTAP Documentation](#)
- Technical Reports:
  - [TR-4938: Mount Amazon FSx for ONTAP as a NFS datastore with VMware Cloud on AWS](#)
  - [TR-4897: SQL Server on Azure NetApp Files - Real Deployment View](#)
  - [TR-4816: Performance Characterization of NetApp Cloud Volumes ONTAP in Google Cloud Platform](#)

[\*Cloud Volumes ONTAP Fundamentals\*](#) (online course)

[\*Integrating Hybrid Clouds with Amazon Web Services\*](#) (virtual live course)

[\*Integrating Hybrid Clouds with Microsoft Azure\*](#) (virtual live course)


[\*Cloud Volumes ONTAP Documentation\*](#)

Technical Reports:

- [\*TR-4938: Mount Amazon FSx for ONTAP as a NFS datastore with VMware Cloud on AWS\*](#)
- [\*TR-4897: SQL Server on Azure NetApp Files - Real Deployment View\*](#)
- [\*TR-4816: Performance Characterization of NetApp Cloud Volumes ONTAP in Google Cloud Platform\*](#)

# Module 2

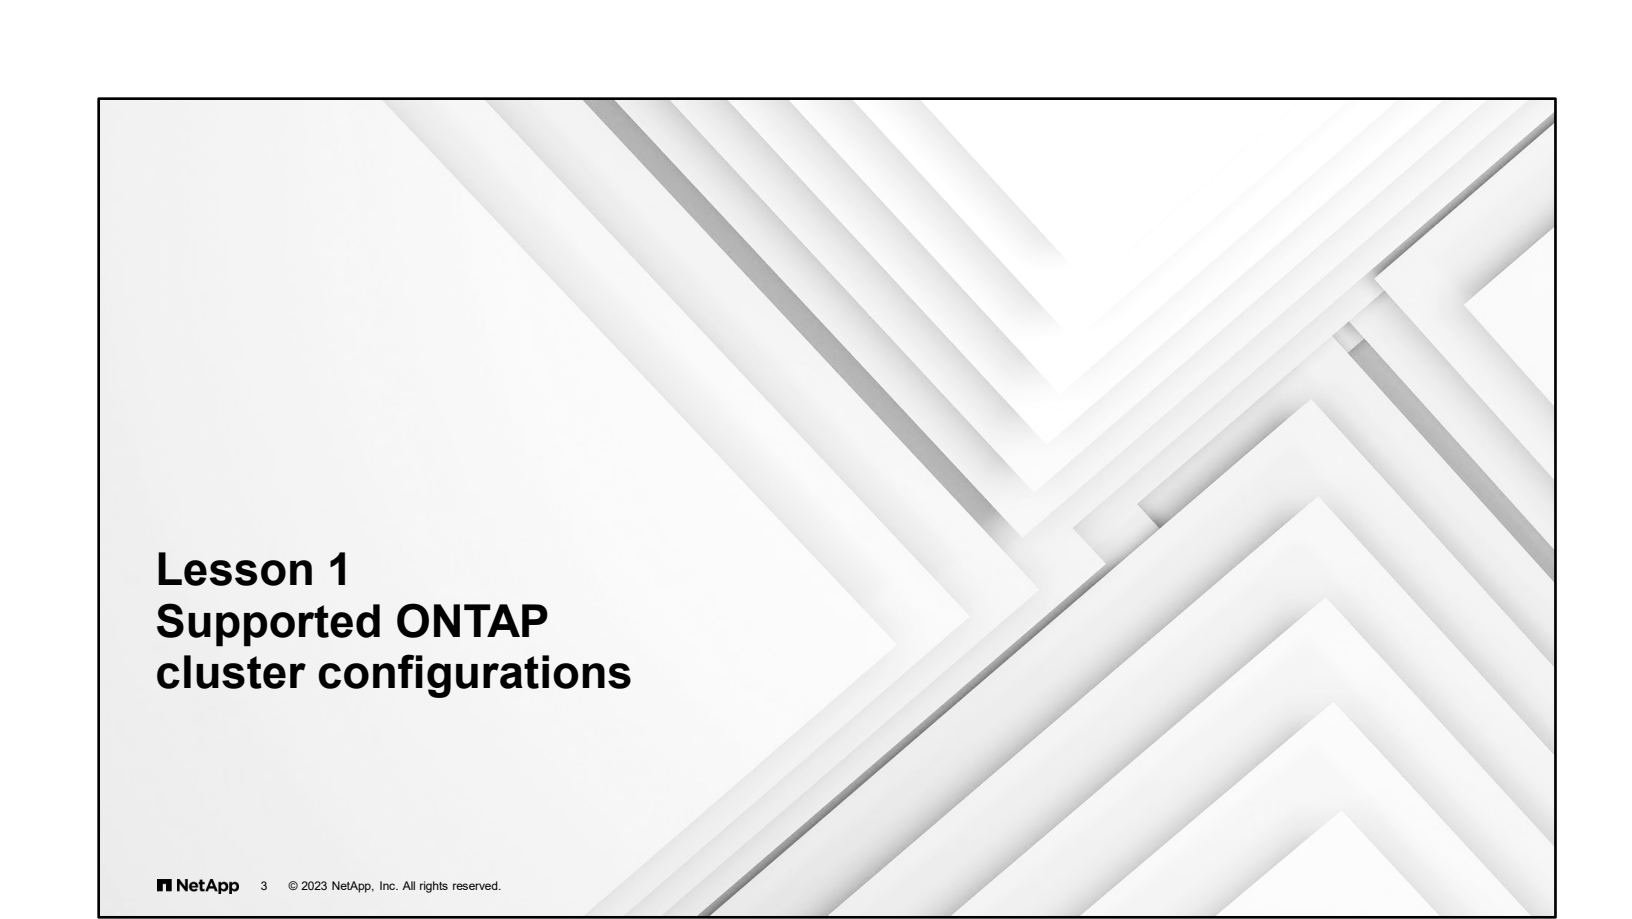
## Cluster setup

 1 © 2023 NetApp, Inc. All rights reserved.

## About this module


This module focuses on enabling you to do the following:

- Identify supported cluster configurations
- List the steps to set up a cluster
- Manage cluster nodes at the hardware level

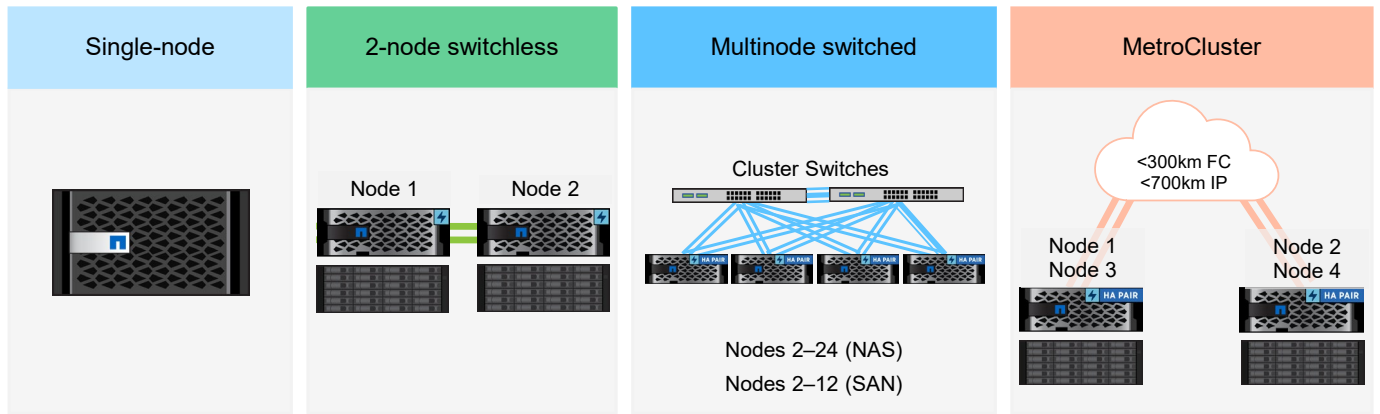


# Lesson 1

## Supported ONTAP cluster configurations

 3 © 2023 NetApp, Inc. All rights reserved.

## Supported cluster configurations



NetApp 4 © 2023 NetApp, Inc. All rights reserved.


Four types of NetApp ONTAP cluster configurations are supported:

- Single-node
- 2-node cluster without network switches
- Multi-node cluster that is made of high-availability (HA) pairs that are connected through network switches
- Geographically separated HA pairs that are connected in a MetroCluster cluster configuration



## Single-node cluster

- Features of a single-node cluster:
  - Special implementation of a cluster that runs on a standalone node
  - An implementation for a workload that requires only one node and does not need nondisruptive operations (NDO)
  - Use case: Data protection for a remote office or test and development
- Features and operations that a single-node cluster does not support:
  - Storage failover (SFO) and high availability
  - Operations that affect multiple nodes



Cluster node


Some features and operations are not supported for single-node clusters. Because single-node clusters operate in a standalone mode, storage failover (SFO) and cluster high availability are unavailable. If the node goes offline, clients cannot access data that is stored in the cluster. Also, any operation that requires more than one node cannot be performed. For example, you cannot move volumes, perform most copy operations, or back up cluster configurations to other nodes.

Because of these limitations, NetApp no longer markets single-node storage systems.

## HA pairs

High-availability (HA) pairs provide hardware redundancy that supports the following features:

- NDO and nondisruptive upgrade (NDU)
- Fault tolerance
- Takeover and giveback of partner storage
- Elimination of most hardware components and cables as single points of failure
- Improved data availability



**NetApp** 6 © 2023 NetApp, Inc. All rights reserved.

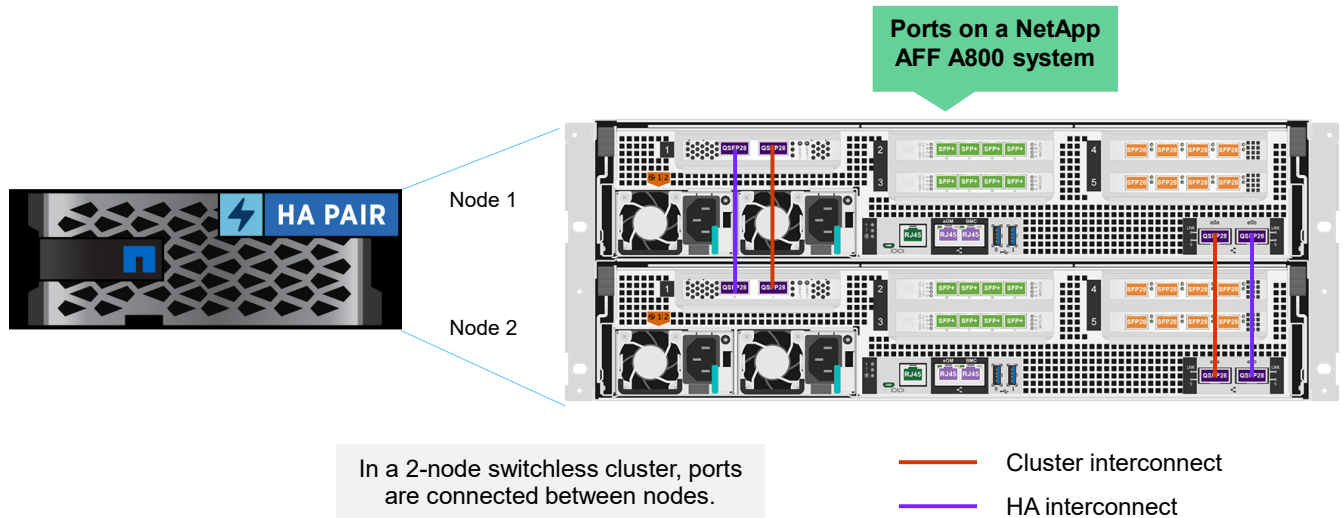
Clusters of two or more nodes are built from HA pairs. HA pairs provide hardware redundancy that is required for nondisruptive operations (NDO) and fault tolerance. The hardware redundancy gives each node in the pair the software functionality to take over and return partner storage. Hardware redundancy also provides the fault tolerance that is required to perform NDO during hardware and software upgrades or maintenance.

A storage system has various single points of failure, such as certain cables or hardware components. An HA pair greatly reduces the number of single points of failure. If a failure occurs, the partner can take over and continue to serve data until the failure is fixed. The controller failover function provides continuous data availability and preserves data integrity for client applications and users.

IDC states: “IDC has reviewed NetApp ONTAP system availability statistics for the period from June 2019 to December 2019, noting that the data indicates a minimum of 99.99993% availability across the tens of thousands of controller pairs running ONTAP 9 software. This population includes NetApp AFF80X0 and AFF A-Series systems as well as FAS25xx, FAS26xx, FAS27xx, FAS8xx0 arrays, and all FAS9000 systems. Clearly, NetApp can deliver ‘six-nines plus’ availability in mixed enterprise workload environments for both block-based and file-based applications.”

Read the full IDC report at <https://fieldportal.netapp.com/content/817678?assetComponentId=819313>.

## 2-node switchless cluster




NetApp 7 © 2023 NetApp, Inc. All rights reserved.

A cluster interconnect is required for cluster communication and data sharing. The example here shows an enterprise-class storage system with two controllers that are installed in the chassis. Each NetApp AFF A800 controller has a set of four onboard Ethernet ports, two that are used for the HA interconnect and two that are used to connect to the cluster interconnect. The AFF A800 is an example of a system that uses external HA interconnect cables. Some NetApp AFF and NetApp FAS models use an internal HA interconnect between HA partner nodes in the same chassis.

In a 2-node switchless cluster, a redundant pair of ports is cabled together, as shown. To enable both HA and SFO functionality in 2-node clusters in which both controllers share the chassis, the HA state must be set by the `ha-config` command in maintenance mode. In most shared chassis systems, the state is set automatically and requires no manual intervention.

## Multinode switched clusters



For more networking details, see the network management module.

NetApp 8 © 2023 NetApp, Inc. All rights reserved.

If a workload requires more than two nodes, the cluster interconnect requires switches. The cluster interconnect requires two dedicated switches for redundancy and load balancing. Inter-Switch Links (ISLs) are required between the two switches. From each node, at least two cluster connections should be present, one connection to each switch. The required connections vary, depending on the controller model and speed of the network ports. Larger systems might require as many as four connections per switch.

After the cluster interconnect is established, you can add more nodes as your workload requires.

For more information about the maximum number and models of controllers that are supported, see the NetApp Hardware Universe ([hwu.netapp.com](https://hwu.netapp.com)).

For more information about the cluster interconnect and connections, see the *ONTAP Networking Reference* [https://docs.netapp.com/us-en/ontap/networking/networking\\_reference.html](https://docs.netapp.com/us-en/ontap/networking/networking_reference.html)

## MetroCluster software

### Benefits of MetroCluster software:

- Geographic separation for business continuity
- Continuous availability and zero data loss
- Set-it-once simplicity
- Zero change management
- Unified solution (support for SAN and NAS)

Learn more about MetroCluster software in the following courses:

First attend:


*ONTAP Data Protection Administration*

Then attend:

*ONTAP MetroCluster Installation (FC)*

– or –

*ONTAP MetroCluster IP Implementation*



**NetApp** 9 © 2023 NetApp, Inc. All rights reserved.

MetroCluster software provides continuously available storage for applications by automatically managing two objectives:

- Zero recovery point objective (RPO) by synchronously mirroring data written to the cluster
- Near zero recovery time objective (RTO) by mirroring configuration and automating access to data at the second site

MetroCluster software provides simplicity with automatic mirroring of data and configuration between the two independent clusters that are in two sites. As storage is provisioned within one cluster, it is automatically mirrored to the second cluster at the second site. MetroCluster software manages the switchover process within the NAS and SAN protocol timeout periods or sooner (typically less than 120 seconds). This management results in a near zero RPO with the applications continuing to access data without incurring failures.


MetroCluster software provides disaster recovery through ONTAP System Manager or by using just one MetroCluster command. The command activates the mirrored data on the surviving site.

MetroCluster sites can connect by FC or IP protocols. Also, 2-node, 4-node, and 8-node MetroCluster configurations are available. Beginning with ONTAP 9.11.1, customers can transition from an existing MetroCluster FC configuration to an AFF A250 or FAS500f MetroCluster IP configuration.



## Lesson 2


### Setting up a cluster

 10 © 2023 NetApp, Inc. All rights reserved.

## Basic hardware setup tasks

Connect the following hardware:

- HA interconnect
- Drive shelf to drive shelf cabling
- Controllers to drive shelves
- Controllers to cluster interconnect
- Controllers to networks
- Any tape devices
- Controllers and drive shelves to redundant power



NetApp 11 © 2023 NetApp, Inc. All rights reserved.

Before ONTAP software can be configured on FAS and AFF systems, the hardware must be installed and set up.

If necessary for your controller type, connect an NVRAM HA cable between partners. The connections can be through the chassis, 10/40/100GbE, or InfiniBand, depending on your storage controllers.

Create shelf stacks by cabling the drive shelves to each other.

Connect controllers to disk shelves. Verify that shelf IDs are set properly.

Connect controllers to networks. Connect any tape devices that you might have. (You can connect tape devices later.)

Connect controllers and disk shelves to power.

This lesson contains only basic information to set up a cluster. For more detailed information, refer to the *Installation and Setup Instructions* (ISI) for your system model and attend the *Universal NetApp FAS Installation* course.

## HA interconnect links

- Are used primarily to mirror the write log
- Provide a channel for certain types of communication traffic between the nodes in an HA pair:
  - Heartbeat
  - Failover
  - Version information
  - Drive firmware updates



**Dual-node chassis**

Uses an internal or node-to-node HA interconnect



**Single-node chassis**

Requires external HA interconnect cables

NetApp 12 © 2023 NetApp, Inc. All rights reserved.

The HA interconnect connects the two controllers of each HA pair for all nodes. The connections are internally provided over the backplane in the chassis of a dual-controller configuration or through node-to-node cabling. For a chassis with a single controller, dedicated HA interconnect cables are required. The dedicated interconnect cable is based on the model and enclosure. Visit the NetApp Support site to see the appropriate hardware configuration guide for your model of storage controller.

The following types of traffic flow over the HA interconnect links:

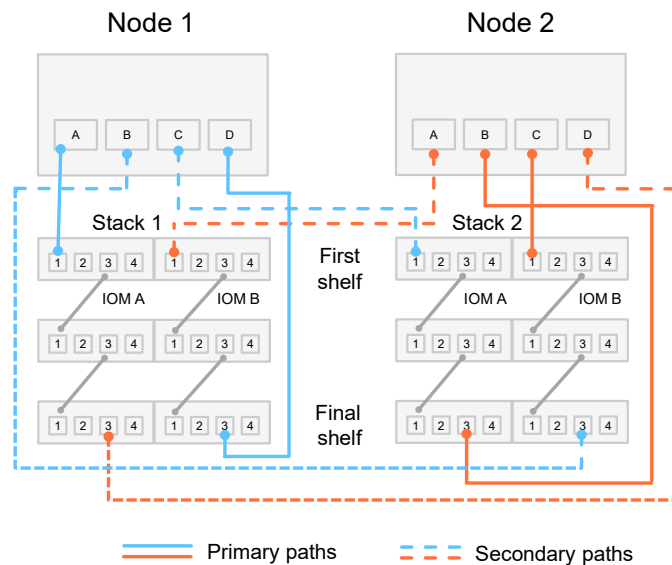
- **Failover:** The directives are related to performing SFO between the two nodes, regardless of which type of failure:
  - Negotiated (planned and in response to an administrator request)
  - Not negotiated (unplanned and in response to an improper system shutdown or booting)
- **Disk firmware:** Nodes in an HA pair coordinate the update of disk firmware. When one node updates the firmware, the other node must not perform any I/O to it.
- **Heartbeat:** Regular messages demonstrate availability.
- **Version information:** The two nodes in an HA pair must be kept at the same major and minor revision levels for all software components.



## Drive shelf cabling

### MPHA configuration

- Multipath high-availability (MPHA) cabling ensures that the storage controllers have redundant paths to all drives in the HA pair.
- MPHA cabling is recommended for HA pair configurations.
- Cabling is mirrored on both nodes to ensure that drive IDs are consistent within the HA pair.



NetApp 13 © 2023 NetApp, Inc. All rights reserved.

To provide fault tolerance, cluster nodes use two connections to every drive in the HA pair. In this example, both storage controllers own a stack of drive shelves.

Both storage controllers use their 0a ports to create the primary path to the first shelf in the shelf stack that is owned by node 1.

Both controllers use the 0d port to create the secondary path from the final shelf in that stack.

To connect to the shelf stack that is owned by node 2, both controllers connect to the first shelf in the stack with port 0c and to the final shelf with port 0b.


The cabling is mirrored so that both nodes generate the same drive ID for all the drives in the pair. If the nodes use different ports, a drive failure is reported on both nodes but with different IDs. This situation can cause confusion, and therefore accidents, when you try to replace a failed drive.

Multipath high-availability (MPHA) cabling is recommended but might not be possible on small systems with fewer storage ports.

## Powering on a system

- Power on network switches.
- Power on drive shelves.
- Power on tape devices (if present).
- Power on storage controllers.




 14 © 2023 NetApp, Inc. All rights reserved.

You should power on the hardware devices in a cluster in the order that is shown.

To power off the entire cluster, power off components in the reverse order.

## Management connections

- Console connections:
  - RJ-45 that uses RS232C ANSI-115.2K-8-None-1
  - Micro USB that uses RS232C ANSI-115.2K-8-None-1
- Remote management device connection: Baseboard Management Controller (BMC) or Service Processor (SP)
- Management network connections (e0M)



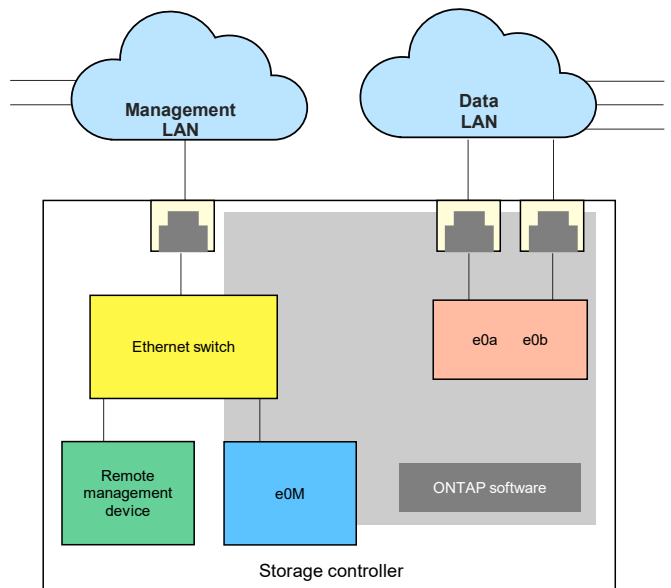
NetApp 15 © 2023 NetApp, Inc. All rights reserved.

Each controller should have a console connection, which is required to get to the firmware and the boot menu. For example, you might use the console connection to the boot menu to access setup, installation, and initialization options. A remote management device connection, although not required, is helpful if you cannot get to the UI or console. Remote management enables remote booting, the forcing of core dumps, and other actions.

The full-sized USB interface is active during only boot device recovery and an ONTAP software update or firmware update.

## Management interfaces

- e0M interface:
  - Is dedicated for management traffic
  - Is used for ONTAP system administration tasks
- BMC or SP interface:
  - Is used to manage and provide remote management capabilities for the storage system
  - Provides remote access to the console and provides monitoring, troubleshooting, logging, and alerting features
  - Remains operational
  - Uses the following setup command: `system service-processor`



NetApp 16 © 2023 NetApp, Inc. All rights reserved.

Some storage system models include an e0M interface. The interface is dedicated to ONTAP management activities. An e0M interface enables you to separate management traffic from data traffic on your storage system for better security and throughput.

To set up a storage system that has the e0M interface, remember the following information:


- The Ethernet port that is indicated by a wrench icon on the rear of the chassis connects to an internal Ethernet switch.
- To manage a LAN in environments in which dedicated LANs isolate management traffic from data traffic, use the e0M interface.
- Configure e0M separately from the Baseboard Management Controller (BMC) or Service Processor (SP) configuration.
- Both configurations require unique IP and MAC addresses to enable the Ethernet switch to direct traffic to either the management interfaces or the BMC or SP.

For more information about configuring remote support, see the *ONTAP System Administration Guide*.

## Console on boot

```
SP node2> system console
Type Ctrl-D to exit.

LOADER> boot_ontap
...
*****
*                               *
* Press Ctrl-C for Boot Menu. *
*                               *
*****
...
```


 17 © 2023 NetApp, Inc. All rights reserved.

Typical boot sequence:

1. Loads the kernel into memory from the boot device
2. Mounts the "/" root image from rootfs.img on the boot device
3. Loads `Init` and runs startup scripts
4. Loads NVRAM kernel modules
5. Creates a `/var` partition on NVRAM (restored from the boot device if a backup copy exists)
6. Starts management processes
7. Loads the data and network modules
8. Mounts the `vol0` root volume
9. Is ready for use

## Boot menu

```
^C
Boot Menu will be available.
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning
(10) Set Onboard Key Manager recovery secrets.
(11) Config node for external key management.
Selection (1-11)? 1
```

 18 © 2023 NetApp, Inc. All rights reserved.

Generally, you allow a node to boot into ONTAP software. The boot menu provides more options that are useful for troubleshooting or maintenance. To access the boot menu, you must press Ctrl+C when you are prompted during the boot sequence.

Select one of the following options by entering the corresponding number:

1. Normal Boot: Continue to boot the node in normal mode.
2. Boot without /etc/rc: This option is obsolete. It does not affect the system.
3. Change password: Change the password of the node, which is also the “admin” account password.
4. Clean configuration and initialize all disks: Initialize the node disks and create a root volume for the node.

**Note:** This menu option erases all data on the disks of the node and resets your node configuration to the factory default settings.

5. Maintenance mode boot: Perform aggregate and disk maintenance operations and obtain detailed aggregate and disk information. To exit Maintenance mode, use the halt command.
6. Update flash from backup config: Restore the configuration information from the root volume of the node to the boot device.
7. Install new software first: Install new software on the node.


**Note:** This menu option is for only installing a newer version of ONTAP software on a node that has no root volume installed. Do not use this menu option to upgrade ONTAP software.

8. Reboot Node: Reboot the node.
9. Configure Advanced Drive Partitioning: For systems that support Advanced Drive Partitioning (ADP), this option enables you to configure ADP of the drives.
10. Define a recovery secret for the Onboard Key Manager.
11. Configure a node to use an external key management system.

## Creating a cluster

### Cluster creation methods:

- **Guided Cluster Setup with NetApp ONTAP System Manager:**
  - Windows discovers new cluster nodes on the network.
  - Double-click a discovered node to launch ONTAP System Manager.
  - Follow the Guided Cluster Setup to configure all the discovered nodes simultaneously.
- **CLI cluster setup wizard:**
  - Use a wizard to create the cluster and join all nodes.
  - Configure the cluster time and NetApp Active IQ functionality.



NetApp 19 © 2023 NetApp, Inc. All rights reserved.

After you install the hardware, you can set up the cluster by using the cluster setup wizard (through the CLI). You can use the Guided Cluster Setup through ONTAP System Manager.

Before you set up a cluster, you should use a cluster setup worksheet and record the values that you need during the setup process. Worksheets are available on the NetApp Support site. If you use the System Setup software, enter the information that you collected on the worksheet as the software prompts you.

To launch the Guided Cluster Setup, start in the Networks page of Windows File Explorer. If the Windows host is attached to the same network as the newly installed ONTAP systems, the Windows host discovers the ONTAP systems automatically. Double-click a discovered node to connect to ONTAP System Manager. When you are prompted, enter the information that you collected on the worksheet. The Guided Cluster Setup discovers all the nodes in the cluster and then configures the nodes simultaneously.

To configure the cluster by using the CLI, enter the cluster setup wizard from a single node in the cluster. The cluster setup wizard prompts you to configure the node management interface. Next, the cluster setup wizard asks whether you want to complete the setup wizard by using the CLI.

If you press Enter, the wizard continues to use the CLI to guide you through the configuration. When you are prompted, enter the information that you collected on the worksheet. After you create the cluster, you use the node setup wizard to join nodes to the cluster one at a time. The node setup wizard helps you to configure the node-management interface of each node.

After you use the CLI to add all nodes, you also need to manually configure a few items. Synchronize the time to ensure that every node in the cluster has the same time and to prevent CIFS and Kerberos failures. You need to decide where to send event notifications: to an email address, a syslog server, or an SNMP trap host. NetApp also recommends that you configure the NetApp Active IQ support tool.

## Cluster setup

Connect to ONTAP System Manager

1. Power on the cluster nodes.
2. Connect a Windows or MacOS host to the same network as the cluster nodes.
3. From the Windows Network page, or the MacOS Finder Networks page, double-click a cluster node to connect to ONTAP System Manager.



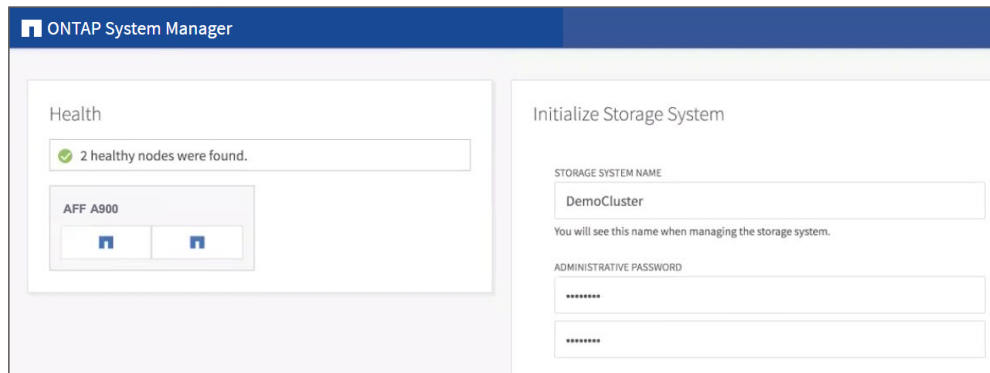
When a Windows or MacOS host is connected to the same network as your new cluster nodes, the nodes appear in the Windows File Explorer Network page or the MacOS Finder Networks page. Double-click one of the cluster nodes to open a connection to ONTAP System Manager.




## Guided Cluster Setup

### System Manager system initialization

- Information about nodes is discovered and displayed.
- Depending on the network configuration, a single-node cluster, a 2-node switchless cluster, or a switched cluster is created.
- Set the cluster name and assign an administrator password.



The screenshot shows the ONTAP System Manager interface during system initialization. The page is divided into two main sections: 'Health' and 'Initialize Storage System'. The 'Health' section on the left shows a green checkmark and the text '2 healthy nodes were found.' Below this, there is a box labeled 'AFF A900' containing two server icons. The 'Initialize Storage System' section on the right has a 'STORAGE SYSTEM NAME' field with the value 'DemoCluster' and a note: 'You will see this name when managing the storage system.' Below that, there are two 'ADMINISTRATIVE PASSWORD' fields, each with a series of asterisks indicating masked input.

 21 © 2023 NetApp, Inc. All rights reserved.

ONTAP System Manager begins in the cluster setup page. Use this page to configure the basic cluster settings, such as the cluster name and administrator password.

## Cluster setup

### Networking section

- In the Networking section, you configure the cluster management and node management network interfaces.
- In the Networking section, you also configure DNS and Network Time Protocol (NTP).
- When you click Submit, the cluster initialization process starts.

#### Networking

CLUSTER IP ADDRESS	SUBNET MASK	GATEWAY
<input type="text" value="10.232.228.132"/>	<input type="text" value="255.255.255.128"/>	<input type="text" value="10.232.228.129"/>
NODE SERIAL NUMBERS	NODE IP ADDRESSES	
<input type="text" value="721639000122"/>	<input type="text" value="10.232.228.134"/>	
<input type="text" value="721639000121"/>	<input type="text" value="10.232.228.131"/>	

Use Domain Name Service (DNS)

DNS DOMAINS

+ Add

NAME SERVERS

+ Add

Use time services (NTP)

NTP SERVERS

+ Add

Use the bottom of the cluster setup page to configure the networking settings. You need to enter the IP addresses for the cluster management and node management network interfaces.

Also, enter your DNS domain name and the IP address of at least one DNS server for host name resolution.

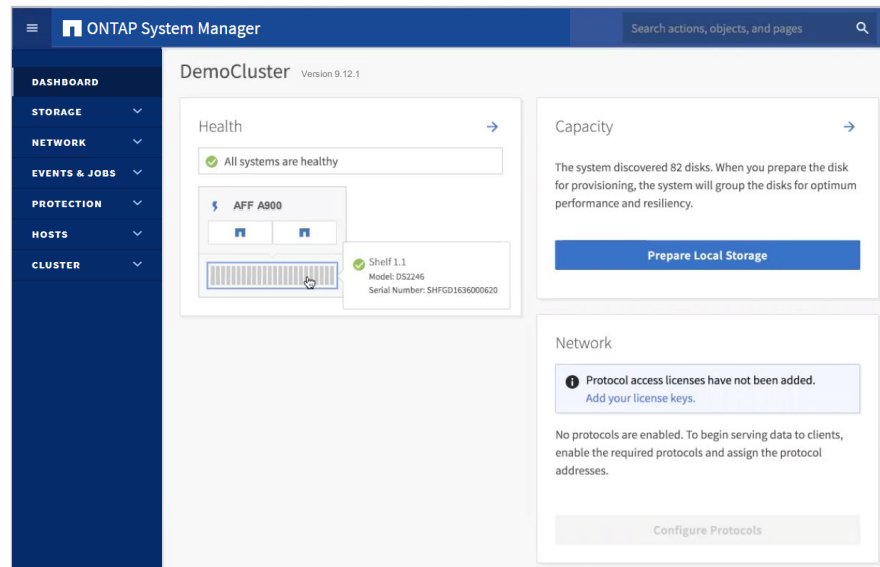
The use of a time service to synchronize the clock of the cluster with the other systems in your environment is strongly recommended. Enter the IP address of at least one time server that supports the NTP protocol.

When you click Submit, the cluster initialization process begins.

## Cluster setup

### Prepare storage

Click **Prepare Local Storage** to configure the storage aggregates (local tiers) according to NetApp best practices.



NetApp 23 © 2023 NetApp, Inc. All rights reserved.

After the cluster initialization is complete, ONTAP System Manager restarts. You might need to use your web browser to approve the connection, because a new security certificate is generated automatically.

The next step is to prepare the storage devices for use. Click **Prepare Local Storage** in the System Manager dashboard to create the data storage aggregates, which are also known as local tiers. The aggregates are configured automatically, following NetApp best practices.

Your feature license keys should come preinstalled by the factory. If not, click **Add your license keys** and enter the license keys manually. You can locate your license keys at the [mysupport.netapp.com](https://mysupport.netapp.com) site.

## Cluster setup

Configure a storage VM for data access

Click **Configure Protocols** to create a storage VM and configure the data access protocols.

The screenshot shows the NetApp System Manager dashboard. The 'Health' section indicates 'All systems are healthy'. The 'Capacity' section shows '0 Bytes USED' and '11.3 TB AVAILABLE' with a '1 to 1 Data Reduction' ratio. The 'Network' section states 'No protocols are enabled. To begin serving data to clients, enable the required protocols and assign the protocol addresses.' A 'Configure Protocols' button is visible at the bottom.

The 'Configure Protocols' page shows the 'STORAGE VM NAME' as 'SANSVM'. Under 'Network Protocols', the 'FC' tab is selected and 'Enable FC' is checked. A table for 'CONFIGURE FC PORTS' is shown below.

Nodes	0e	0f	0g	0h
DemoCluster-01	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
DemoCluster-02	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

'Save' and 'Cancel' buttons are at the bottom.

NetApp 24 © 2023 NetApp, Inc. All rights reserved.

To enable client access to the storage system, you must create a storage VM and configure the data access protocols. Click **Configure Protocols** in the System Manager dashboard.

The tabs in the **Configure Protocols** page enable you to configure the data access protocols for which a compliant license is installed. In this example, only the SAN protocol licenses are present, so no NAS protocols tabs are displayed.


Depending on the data access protocol that is enabled, you might also need to configure LUNs for the SAN protocols or exports or shares for NAS protocols. The methods of doing so are covered in later modules of this course.

If you want to practice a Guided Cluster Setup, see the NetApp Hands-On Labs.

## More training

- Universal AFF and FAS Installation  
(web-based course)



 25 © 2023 NetApp, Inc. All rights reserved.


*Universal AFF and FAS Installation* (web-based course)

[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours00000000027858](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours00000000027858)



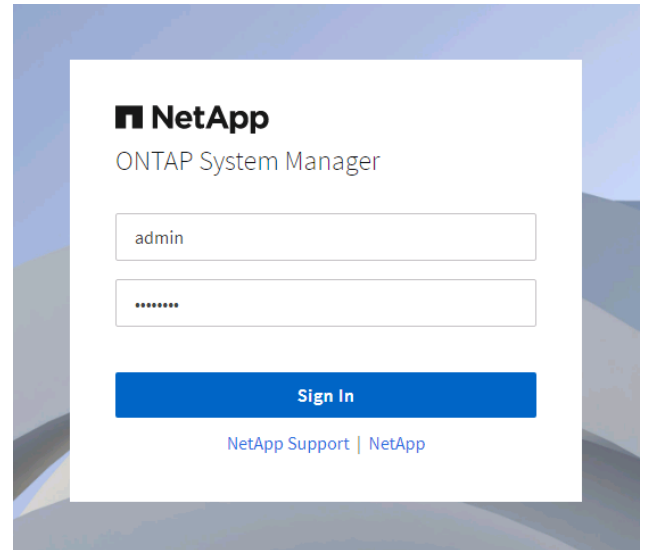
## **Lesson 3**


# **Administration interfaces**

 26 © 2023 NetApp, Inc. All rights reserved.

## Cluster administrators

- Manage the entire cluster:
  - All cluster resources
  - Storage VM (storage virtual machine, also known as SVM) creation and management
  - Access control and roles
  - Resource delegation
- Use login credentials:
  - User name (default): admin
  - Password: password that you created during cluster setup



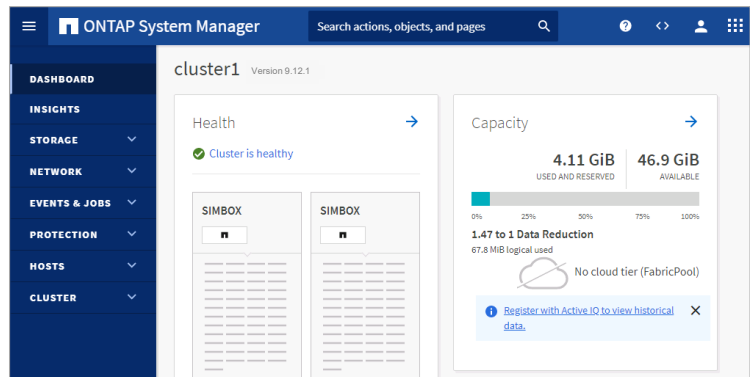
 27 © 2023 NetApp, Inc. All rights reserved.

Use ONTAP System Manager to manage the entire cluster. You manage all cluster resources, the creation and management of storage VMs (storage virtual machines, also known as SVMs), access control and roles, and resource delegation.

To log in to the cluster, use the default user name “admin” and the password that you configured during cluster creation.

## Managing resources in a cluster

- **ONTAP System Manager:**
  - Visual representation of the available resources
  - Wizard-based resource creation
  - Best practice configurations
  - Limited advanced operations
- **The CLI:**
  - Manual or scripted commands
  - Manual resource creation that might require many steps
  - Ability to focus and switch quickly among specific objects
- **Automation tools:**
  - ONTAP RESTful API
  - NetApp PowerShell Toolkit
  - NetApp Python client library
  - Ansible playbooks



```
login as: admin
Using keyboard-interactive authentication.
Password: *****

cluster1:~> cluster show
Node           Health Eligibility
-----
cluster1-01    true   true
cluster1-02    true   true
```

NetApp 28 © 2023 NetApp, Inc. All rights reserved.

You can use many tools to create and manage cluster resources. Each tool has advantages and disadvantages.

ONTAP System Manager is a web-based UI that provides a visual representation of the available resources. Resource creation is wizard-based and adheres to best practices. However, not all operations are available. Some advanced operations might need to be performed by commands in the CLI.

You can use the CLI to create and configure resources. Enter commands manually or through scripts. Instead of the wizards that System Manager uses, the CLI might require many manual commands to create and configure a resource. Although manual commands give the administrator more control, manual commands are also more prone to mistakes that can cause issues. One advantage of using the CLI is that the administrator can quickly switch focus without needing to move through System Manager pages to find different objects.

You can also manage ONTAP clusters through the ONTAP API. Automation tools like Ansible use the ONTAP API to create and manage cluster resources. NetApp and Ansible have partnered to develop a collection of automation and storage configuration management modules to easily provision, deploy, and manage NetApp storage systems. NetApp playbooks combine Ansible modules to deliver full-stack presentations of storage to the host. The recommendation is to use Ansible modules for NetApp, and to use sample playbooks and collections in conjunction with AWX or Ansible Tower for new automation projects. For more information and documentation about NetApp modules for Ansible, go to [thePub](#), and attend the Automate Storage Administration Using ONTAP REST API and Ansible course



## Clustershell

The default CLI, or shell, in ONTAP software is called the clustershell and has the following features:

- Inline help
- Online manual pages
- Command history
- Ability to reissue a command
- Keyboard shortcuts
- Queries and UNIX-style patterns
- Wildcards

```
login as: admin
Using keyboard-interactive authentication.
Password: *****

cluster1::> cluster show
Node                Health  Eligibility
-----
cluster1-01         true   true
cluster1-02         true   true
cluster1::>
```

The cluster has different CLIs or shells for different purposes. This course focuses on the clustershell, which starts automatically when you log in to the cluster.


Clustershell features include inline help, an online manual, `history` and `redo` commands, and keyboard shortcuts. The clustershell also supports queries and UNIX-style patterns. Wildcards enable you to match multiple values in command-parameter arguments.

## Clustershell

Command scope

```
cluster1::> storage aggregate
```

```
cluster1::storage aggregate> modify
```

 30 © 2023 NetApp, Inc. All rights reserved.

Typing the first two levels of the command directory puts you in the command directory. You can then type a command from that level or type a fully qualified command from a different command directory.

## Clustershell

Scope return

```
cluster1::storage disk option> ..  
cluster1::storage disk> top  
cluster1::>
```

Use the `..` command to move up one level in the command hierarchy. Use the `top` command to move to the top level of the command hierarchy.

## Clustershell

### Use of the question mark wildcard

```
cluster1::> storage aggregate
cluster1::storage aggregate> modify ?
[ -aggregate ] <aggregate name>           Aggregate
[ -disktype|-T {ATA | BSAS | FCAL | FSAS | LUN | MSATA | SAS | SATA | SSD | VMDISK} ]
                                           Disk Type
[ -free-space-realloc {on|off|no_redirect} ] Free Space Reallocation
[ -ha-policy {sfo|cfo} ]                   HA Policy
[ -percent-snapshot-space <percent> ]     Space Reserved for Snapshot Copies
[ -space-nearly-full-threshold-percent <percent> ] Aggregate Nearly Full Threshold Percent
[ -space-full-threshold-percent <percent> ] Aggregate Full Threshold Percent
[ -hybrid-enabled {true|false} ]          Hybrid Enabled
[ -force-hybrid-enabled|-f [true] ]       Force Marking of Aggregate as Hybrid Enabled
[ -maxraidsize|-s <integer> ]             Max RAID Size
...
cluster1::storage aggregate> modify
```

At the command line, press the question mark (?) key to show the command directories and commands that are available at that command level.

## Clustershell

### Tab completion

```
cluster1::storage aggregate> modify
  aggr0_n1 aggr0_n2 n1_data_001 n1_data_002
  n1_data_003 n2_data_001

cluster1::storage aggregate> modify -aggregate n2_data_001 -state online
Aggregate online successful on aggregate: n2_data_001

cluster1::storage aggregate>
```

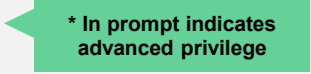
Press the **Tab** key to show available directories, commands, and parameters or to automatically complete a command (or a portion of a command). You can also use the Tab key to complete non-ambiguous substrings of commands, parameters, and values.

## Clustershell

### Privilege levels

```
cluster1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous;
use them only when directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y

cluster1::*>
cluster1::*>
cluster1::*>
cluster1::*> set admin
cluster1::>
```



ONTAP software provides multiple sets of commands that are based on privilege levels. ONTAP software offers administrative, advanced, and diagnostic levels. Use the `priv` command to set the privilege level.

The administrative level provides access to commands that are sufficient for managing your storage system. The advanced and diag levels provide access to the same administrative commands, plus additional troubleshooting and diagnostic commands.

Advanced level and diag level commands should be used only with the guidance of NetApp technical support.

## Clustershell

### More features

The search path enables you to run commands out of scope:

```
cluster1::system node> disk show = storage disk show
```


Abbreviation is permitted (shortest unambiguous sequences of characters):

```
cluster1::> aggr show = storage aggregate show  
cluster1::> net int show = network interface show
```

You can run queries with patterns and wildcards:

```
cluster1::> storage disk show -physical-size >500gb
```

Use the up-arrow key to review command history.

 35 © 2023 NetApp, Inc. All rights reserved.

You can abbreviate commands and parameters in the clustershell if the abbreviation is unambiguous in the current context. You can also run commands out of context, if the command is unavailable in any other context.

## References

- NetApp Hardware Universe:  
<http://hwu.netapp.com>
- ONTAP 9 Documentation Center:  
<http://docs.netapp.com/ontap-9/index.jsp>
- ONTAP System Manager Documentation Center:  
<https://docs.netapp.com/us-en/ontap/index.htm>

NetApp Hardware Universe: <http://hwu.netapp.com>

ONTAP 9 Documentation Center: <http://docs.netapp.com/ontap-9/index.jsp>

ONTAP System Manager Documentation Center: <https://docs.netapp.com/us-en/ontap/index.html>




## Module summary

This module focused on enabling you to do the following:

- Identify supported cluster configurations
- List the steps to set up a cluster
- Manage cluster nodes at the hardware level

# Knowledge check

Module 2: Cluster setup

 38 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

**Client data that is passed between cluster nodes travels over which links?**

- a. HA interconnect
- b. cluster interconnect
- c. management network
- d. data network




## Complete an exercise

Module 2: Cluster setup

### Exploring ONTAP Management UIs

- Access your lab equipment.
- Open your Exercise Guide, Module 2.
- Complete Exercise 1.
- Share your results.

This exercise requires approximately  
**30 minutes.**

 40 © 2023 NetApp, Inc. All rights reserved.

See the instructions in your Exercise Guide.



## Share your experiences

Roundtable discussion

### ONTAP System Manager versus clustershell:


- Which method do you prefer to use?
- Why?

Have a roundtable discussion with the class to answer these questions. Add any comments about experiences or “lessons learned” during the exercises that others might find helpful.

If you encounter an issue, notify your instructor immediately so that it can be resolved promptly.

# Module 3

## Cluster management

 1 © 2023 NetApp, Inc. All rights reserved.

In this module, you learn how to configure key features of NetApp ONTAP software.

## About this module

This module focuses on enabling you to do the following:


- Manage access control
- Set the date and time on cluster nodes
- Manage NetApp ONTAP software licenses
- Manage jobs and schedules

The cluster might require initial configuration, depending on the environment. This module discusses access control, date and time, licenses, jobs, and schedules.



# Lesson 1


## Access control


 3 © 2023 NetApp, Inc. All rights reserved.



## Cluster administrators and storage VM administrators

- **Tasks of cluster administrators:**
  - Administer the entire cluster
  - Administer storage VMs (storage virtual machines, also known as SVMs) on the cluster
  - Create and delegate aggregates for storage VM administrator use
  - Set up data storage VMs and delegate storage VM administration to storage VM administrators
- **Tasks of storage VM administrators:**
  - Administer only their own data storage VMs
  - Set up storage and network resources, such as volumes, protocols, LIFs, and services

 4 © 2023 NetApp, Inc. All rights reserved.



This module focuses on cluster administration. Two types of administrators can manage a cluster.

What a storage VM (storage virtual machine, also known as SVM) administrator can configure is based on how the cluster administrator has configured the storage VM administrator's user account.

## Admin storage VM

Admin storage VM:

- Created automatically during the cluster setup process
- Representation of the cluster
- Not a data server.
  - A cluster must have at least one data storage VM to serve data to clients.
- Primary access point for administration of nodes, resources, and data storage VMs



The cluster management LIF is configured to fail over to any node in the cluster.

The admin storage VM is used to manage the cluster.

There is only one admin storage VM, and it represents the entire cluster. You can use the cluster management LIF to manage any node, resource, or data storage VM.

Unless otherwise specified, the term storage VM typically refers to a data-serving storage VM. Also, in the CLI, storage VMs are displayed as “Vservers,” and many commands use a `-vserver` parameter to specify storage VMs. The term “Vserver” is a holdover from early versions of ONTAP software (formerly clustered Data ONTAP) and is maintained for backward compatibility.

## Admin access

An administrator account is for a predefined cluster administrator:

- Uses the CLI or NetApp ONTAP System Manager (formerly OnCommand System Manager)
- Is associated with cluster or data storage VMs



You can create additional administrator accounts with role-based access control (RBAC):

```
cluster1::> security login
```

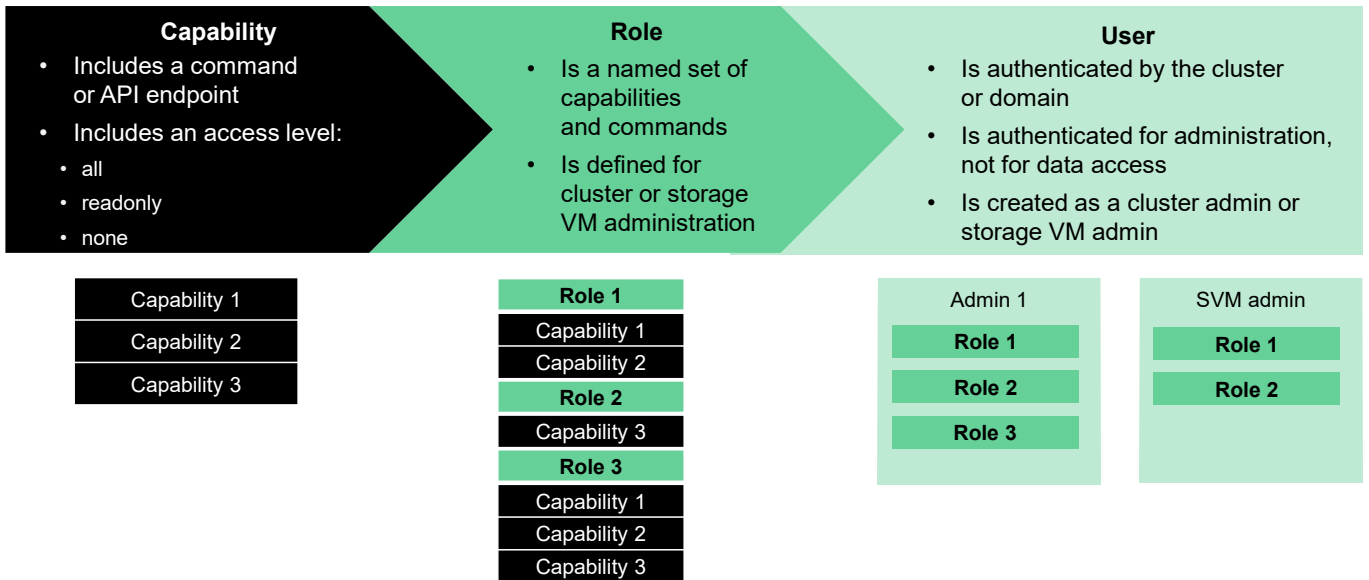
You can use the default system administration account to manage a storage system, or you can create administrator user accounts to manage administrative access to the storage system.

You might want to create an administrator account for the following reasons:

- You can specify administrators and groups of administrators with different degrees of administrative access to your storage systems.
- You can limit an administrator's access to specific storage systems by providing an administrative account on only those systems.
- Creating different administrative users enables you to display information about who is performing which commands on the storage system.

# RBAC

RBAC users, roles, and capabilities



You assign users to roles based on their responsibilities.

Each role is granted a set of rules that enables a set of capabilities. A role is defined as cluster-scoped or SVM-scoped. You can use built-in roles and create custom roles. The capabilities of the predefined roles cannot be changed.

Capabilities are a combination of a command and an access level. A command is a specific instruction or an entire command tree. The three access levels are all, read-only, and none.

Beginning with ONTAP 9.7 software, capabilities can be defined by using API endpoints instead of commands. An API endpoint might be a specific resource or an entire resource category.

Administrator accounts are assigned roles, and roles are assigned capabilities.

## RBAC

Predefined roles in ONTAP software

### Cluster storage VM roles:


- admin
- readonly
- none
- backup
- autosupport

```
::> security login role show -vserver cluster1
```

### Data storage VM roles:

- vsadmin
- vsadmin-volume
- vsadmin-protocol
- vsadmin-backup
- vsadmin-snaplock
- vsadmin-readonly

```
::> security login role show -vserver svm1
```

 8 © 2023 NetApp, Inc. All rights reserved.

ONTAP software includes administrative access-control roles that can be used to subdivide administration duties for storage VM administration tasks.

The vsadmin role is the superuser role for a storage VM. The admin role is the superuser for a cluster.

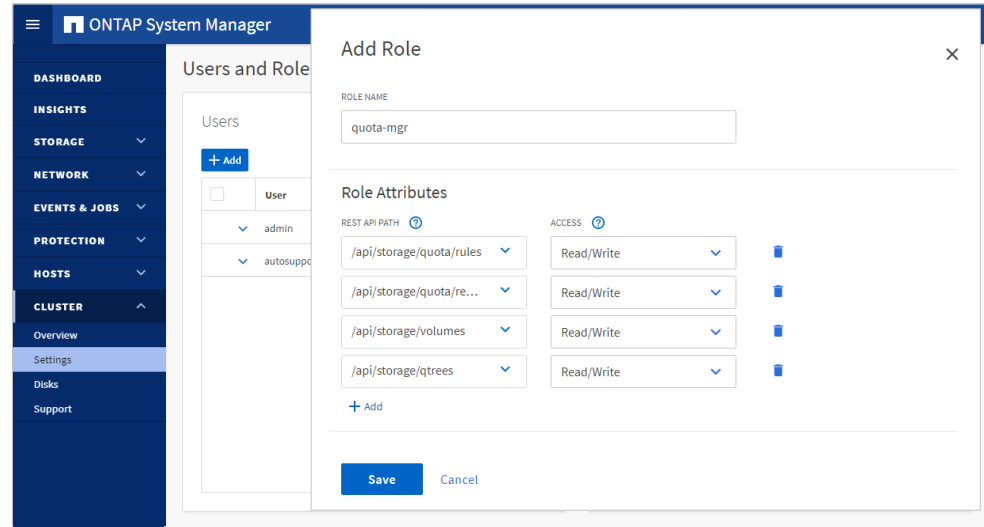
The vsadmin role grants the data storage VM admin full administrative privileges for the storage VM. Additional roles include the vsadmin-protocol role, the vsadmin-readonly role, and the vsadmin-volume role. Each role provides a unique storage VM administration privilege.

A cluster admin with the “readonly” role can grant read-only capabilities. A cluster admin with the “none” role cannot grant capabilities.

## RBAC

### Custom roles

- Role name
- Command directory or API resource
- Optional query or object identifier
- Access level



```
::> security login role create -vserver svml -role svmlvols -cmddirname volume -access all
```

```
::> security login modify -vserver svml -role svmlvols -user ken
```

NetApp 9 © 2023 NetApp, Inc. All rights reserved.

Cluster administrators can create access-control roles to apply to cluster or storage VM administrators. The roles can grant or limit authority to perform certain system administration tasks. An access-control role consists of a role name and a command or a command directory to which the role has access. The role can include an access level (none, readonly, or all) and a query that applies to the specified command or command directory. The example that is shown creates a role that is named `svmlvols` and that grants access to the volume commands but limits access to aggregates that start with the “aggr7” string. The role is assigned to a user who is named Ken.

After the role is created, you can apply the role to individual administrators:

```
c1::> security login role create -vserver svml -role svmlvols -cmddirname volume -query "-aggr aggr7*" -access all
```

```
c1::> security login modify -vserver svml -user ken -role svmlvols
```

Beginning with ONTAP 9.7 software, a REST access-control role consists of a role name and an API resource path name or resource category to which the role has access. Access to specific resources can be controlled by including a unique object identifier in the resource path name. This example creates a role that is named `svmlvols` and that grants access only to volumes owned by `svml` and that reside in aggregates whose names begin with “aggr7.”

Note that the `vol show` command is used to obtain the unique volume identifiers and that the `security login rest-role create` command must be run for each individual volume.

```
c1::> vol show -vserver svml -aggregate aggr7* -fields uuid
```

```
c1::> security login rest-role create -vserver svml -role svmlvols -api /api/storage/volumes/<vol_uuid> -access all
```


Use the `security login role show-ontapi` command to view the mapping of ONTAP command directory names to ONTAP API path names and to convert one type of name to the other. See the *ONTAP REST API Developers Guide* for a description of the API resources.

## Creating ONTAP administrator accounts

- Use the `security login` command to configure role-based administrative access to the cluster.
- Specify the application (access method):  
console, HTTP, SNMP, Secure Shell (SSH), and the API interface.
- Specify the authentication method:  
password, Secure Sockets Layer (SSL) certificate, SNMP community string, Active Directory authentication, Lightweight Directory Access Protocol (LDAP) or Network Information Service (NIS) authentication, public-key authentication, or Security Assertion Markup Language (SAML) authentication.
- Optionally, specify an access-control role.

```
::> security login create -vserver cluster1 -user-or-group-name elsa -role admin  
-application http -authentication-method password
```

```
Please enter a password for user 'elsa': *****  
Please enter it again: *****
```

 10 © 2023 NetApp, Inc. All rights reserved.

The `security` command tree contains all the commands that are necessary to configure role-based administrative access.

The `security login create` command creates a login method for the management utility. A login method consists of a user name, an application (access method), an authentication method, and optionally an access-control role name.

## Active Directory authentication for administrators

- You must configure access to an Active Directory domain controller to authenticate domain accounts.
- If you have already configured a CIFS server for a data storage VM, you can configure the storage VM as a gateway, or *tunnel*, for Active Directory access by the cluster.


```
::> security login domain-tunnel create -vserver svm3
```

- If you have not configured a CIFS server, you can create a computer account for the cluster in the Active Directory domain.

```
::> vserver active-directory create -vserver cluster1 -account-name CLUSTER1 -domain demo.com
```

- When a login is created for an Active Directory group, all group members share access privileges.

```
::> security login create -vserver cluster1 -role admin -application ssh  
-user-or-group-name demo\Administrators -authentication-method domain
```

 11 © 2023 NetApp, Inc. All rights reserved.

You do not need an ONTAP CIFS license to use Active Directory for authentication of cluster and storage VM administrators. Use the `vserver active-directory create` command to create a computer account for the storage VM in Active Directory.

```
::> vserver active-directory create -vserver cluster1 -account-name CLUSTER1 -domain  
demo.com
```

If you have already configured a CIFS server for a data storage VM, you can use the `security login domain-tunnel create` command to configure the storage VM as a gateway, or *tunnel*, for Active Directory access to the cluster.

```
::> security login domain-tunnel create -vserver svm3
```

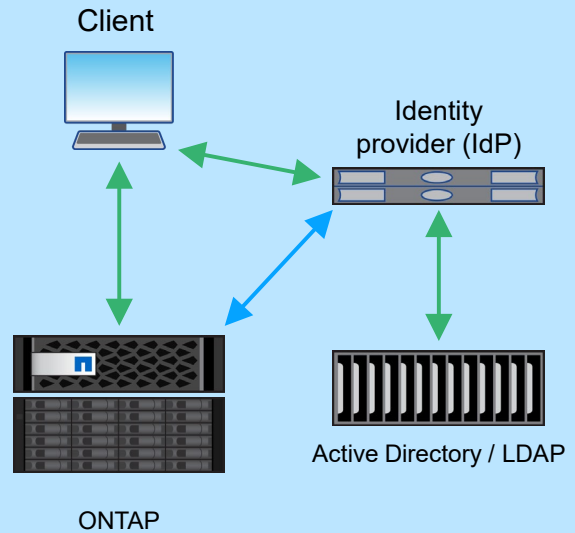
Beginning with ONTAP 9.10.1 software, administrator authentication is attempted by using Kerberos first, and if that fails, with NT LAN Manager (NTLM).



## Securing administrator access

### Multifactor authentication

- Secure access to System Manager and the ONTAP APIs
  - Use an external identity provider to authenticate users and enforce multifactor authentication.
  - Validated with Microsoft Active Directory Federated Services (ADFS) IdP and open-source Shibboleth IdP.
- Secure access to the ONTAP CLI
  - You must have both a SSH public key and a password for multifactor authentication.
  - You must associate the public key with the account before the account can access the storage VM.



NetApp 12 © 2023 NetApp, Inc. All rights reserved.

You can configure multifactor authentication (MFA) for web services by using Security Assertion Markup Language (SAML) authentication. You can use SAML authentication for Service Processor Infrastructure (spi), ONTAP APIs, and NetApp ONTAP System Manager (formerly OnCommand System Manager).

When you configure SAML authentication, users are authenticated by an external Identity Provider (IdP). The IdP is a third-party software such as Microsoft Active Directory Federated Services (ADFS) IdP or open-source Shibboleth IdP. The ONTAP cluster acts as the SAML Service Provider (SP) host. Authentication is performed by exchanging metadata between the IdP and SP.

First, use the `security saml-sp create` command to connect to the external IdP. Then, when you configure the administrator account with the `security login create` command, use the `-authentication-method saml` option to require external authentication of the administrator identity.

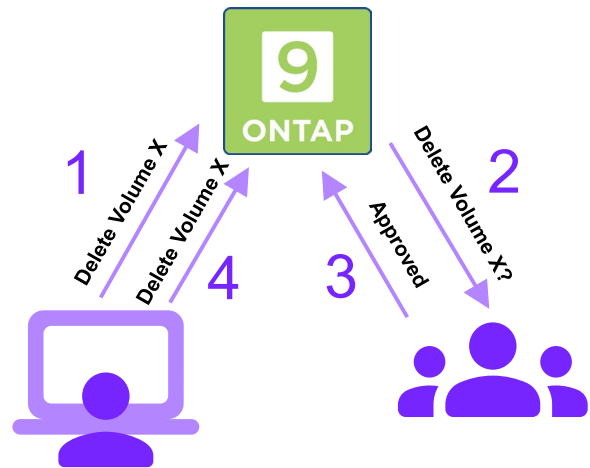
Refer to *TR-4647 Multifactor Authentication in ONTAP 9.3* for more information.

You can use the `security login create` command to require administrators to log in to an admin storage VM or data storage VM with both an SSH public key and a user password.

## Multi-admin approval

### Insider threat mitigation

- Defend against a single rogue or compromised administrator account.
- Require approvals for commands that could result in data loss or data exfiltration.
- Target any action that could result in data destruction, recovery point loss, or data theft.
  - Predefined protected command templates
  - Customer-specific protected command templates



NetApp 13 © 2023 NetApp, Inc. All rights reserved.

The NetApp Multi-Admin Verification feature is an enhanced version of the industry standard Dual Admin Control. Customers enable Multi-Admin Verification to protect themselves against rogue administrators, social engineering attacks, and honest administrator errors. One or more additional administrators must approve actions that could result in data loss or data exfiltration.

When an administrator performs a protected operation, either through the ONTAP System Manager or another management interface, the operation is suspended and queued. ONTAP generates an approval request and optionally sends email notifications to the members of the approving administrator group. After the required number of administrators have granted approval, the requesting administrator can resume the protected operation. If email notifications are configured and the requestor is in the same approval group for the operation, the requestor receives an email when their request is approved. The requesting administrator cannot approve their own operation requests, even if they are a member of the approving administrator group.

Multi-Admin Verification should not be used in environments that automate the protected operations.

## Configuring multi-admin approval

Cluster > Settings > Multi-Admin Approval

Enable multi-admin approval by creating at least one approval group and one protected operation rule.

- Define at least one approver group:
  - Name: Enter a group name.
  - Approvers: Select approvers from a list of administrators.
  - Email address: Enter approver email addresses.
  - Default group: Select a default approver group.
- Optionally define the global settings:
  - Approval expiry period (default 1 hour, max 14 days)
  - Execution expiry period (default 1 hour, max 14 days)
- Define at least one protected operation rule:
  - Operation: Select a supported command from the list.
  - Query: Enter any desired command options and values.
  - Optionally, override the required number of approvers.
  - Optionally, override the default approval groups.

Name	Approvers	Email Address	Default Group
ClusterApprovers	admin	ONTAPApprovers@compa...	<input checked="" type="checkbox"/>
LegalApprovers	harry	LegalApprovers@compa...	<input type="checkbox"/>

Operation	Query	Required Number of Appr...	Approval Groups
set	-privilege diagnostic		Optional
volume delete	-server Legal	2	LegalAp... X
volume snap...	-server Legal	2	LegalAp... X

NetApp 14 © 2023 NetApp, Inc. All rights reserved.

Beginning with ONTAP 9.11.1, you can use multi-admin verification (MAV) to ensure that certain operations, such as deleting volumes or Snapshot copies, can be executed only after approvals from designated administrators.

To configure multi-admin verification, do the following:

1. Create one or more administrator groups with approval and veto powers.
2. Enable the multi-admin approval functionality.
3. Add or modify the set of protected operations or commands in a *rules table*.

When multi-admin verification is enabled, system-defined rules (also known as *guardrail* rules) establish MAV operations to contain the risk of circumventing the MAV process itself. These operations cannot be removed from the rules table. After MAV is enabled, one or more administrators must approve any operation that changes the MAV configuration.

MAV rules are evaluated after role-based access control (RBAC) rules. Therefore, administrators who execute or approve protected operations must already possess the minimum RBAC privileges for those operations.

Enable multi-admin approval by creating at least one approval group and adding at least one rule.

## Administrative auditing

- You can monitor administrator activity for compliance and accountability.
- To enable and disable security audit logging, use the following command:

```
::> security audit modify -cliget on -ontapiget on
```

- Audited commands go to the management log.
- The `security audit log show` command displays cluster-wide audit log messages.

```
::> security audit log show -user elsa
```

- Nodes track local SSH and console commands in the command history log.

All write operations that are performed through the CLI, the web (HTTP) interface, and the ONTAP API are automatically logged. The logging of read operations is disabled by default, but it can be enabled using System Manager or with the `security audit modify` command.

`-httpget`: This option enables auditing of read requests that are made through ONTAP System Manager.

`-cliget`: This option enables auditing of read requests that are made from the command line.

`-ontapiget`: This option enables auditing of read requests that are made through the ONTAP API.

The `security audit log show` command displays cluster-wide audit log messages. Messages from all nodes are shown in chronological order.

Operations that are run from the command line are also recorded in the command history log.

## Security login banner and message of the day

For legal purposes, some computer systems must display a warning to unauthorized users who are connecting to the system.

- This legal warning is configured in ONTAP software by using the `security login banner` command.

```
::> security login banner modify
```

- The message of the day (MOTD) subcommand enables you to show a message to all cluster and storage VM administrators when they open a console session:

```
::> security login motd modify
```

When you connect to any government or corporate system, one of the first things that you see is a warning about the legal consequences of unauthorized access. You can use the `security login banner` command to configure this legal warning on your cluster.

Another feature of the `security login` command is the message of the day subcommand. This command enables you to display a short message to anyone who logs in through the CLI console. You might want to provide a reminder about a meeting, system maintenance, or planned downtime, or you might send a birthday or work anniversary greeting to someone.

## Date and time

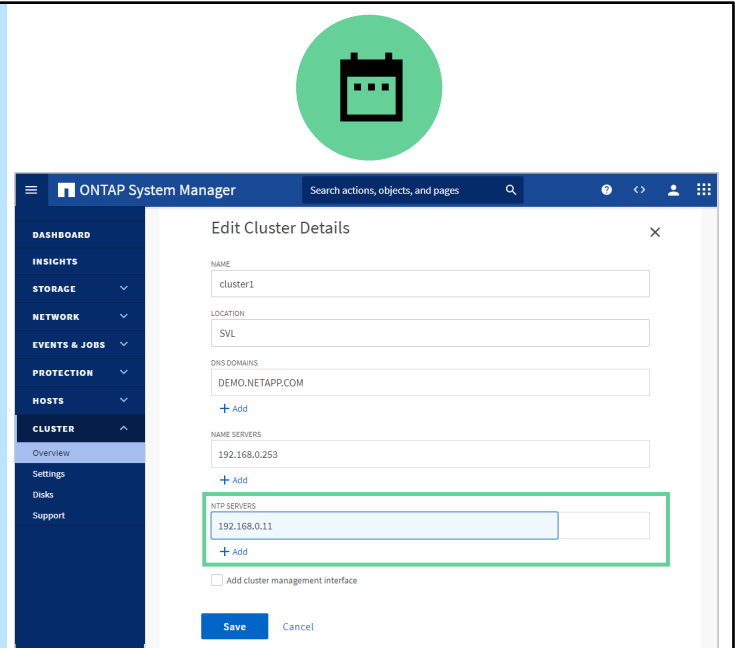
Ways to configure date and time:

- Manually, with the CLI
- Automatically, with Network Time Protocol (NTP) servers

After you add an NTP server, the nodes require time to synchronize.

```
::> cluster time-service ntp server create  
-server xx.xx.xx.xx  
::> date
```

NetApp 17 © 2023 NetApp, Inc. All rights reserved.

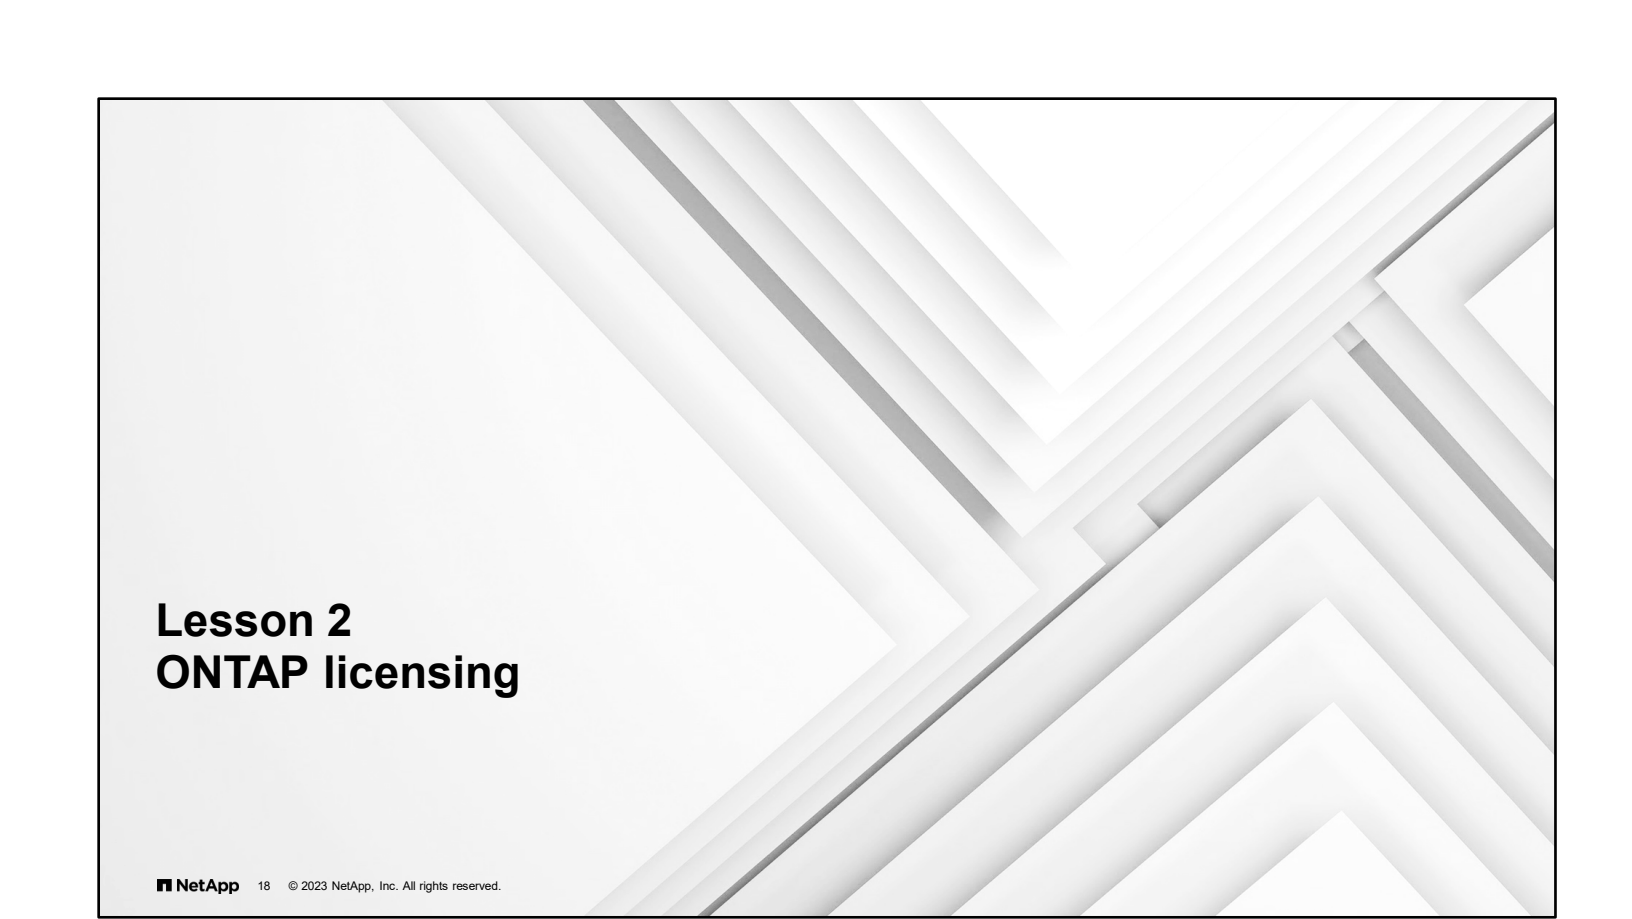


The screenshot shows the ONTAP System Manager interface. At the top right, there is a green circular icon with a calendar symbol. Below it, the main interface is titled 'ONTAP System Manager' with a search bar. A left-hand navigation menu includes 'DASHBOARD', 'INSIGHTS', 'STORAGE', 'NETWORK', 'EVENTS & JOBS', 'PROTECTION', 'HOSTS', and 'CLUSTER'. The 'CLUSTER' section is expanded to show 'Overview', 'Settings', 'Disks', and 'Support'. The 'Settings' section is active, displaying the 'Edit Cluster Details' dialog box. This dialog box has fields for 'NAME' (cluster1), 'LOCATION' (SVL), 'DNS DOMAINS' (DEMO.NETAPP.COM), and 'NAME SERVERS' (192.168.0.253). The 'NTP-SERVERS' field is highlighted with a green border and contains the IP address '192.168.0.11'. There is an 'Add' button below the NTP-SERVERS field and a checkbox for 'Add cluster management interface'. At the bottom of the dialog, there are 'Save' and 'Cancel' buttons.


Problems can occur when the cluster time is inaccurate. ONTAP software enables you to manually set the time zone, date, and time on the cluster. However, you should configure the Network Time Protocol (NTP) servers to synchronize the cluster time.

To configure the date and time, in ONTAP System Manager, on the Cluster Overview page, select **Edit** from the More menu. Click **Add**, and in the NTP Servers field, enter the IP addresses of the time servers, separated by commas. Click **Save**.


Adding the NTP server automatically configures all the nodes in the cluster, but each node needs to synchronize individually. The synchronization for all the nodes in the cluster might require a few minutes.




## Lesson 2 ONTAP licensing

 18 © 2023 NetApp, Inc. All rights reserved.

## License types




Standard and site licenses



Evaluation license



Capacity license

 19 © 2023 NetApp, Inc. All rights reserved.

All the basic features of ONTAP software are included with each cluster that you purchase. Additional features of ONTAP software are enabled by installing additional licenses.

- A site license is for customers with multiple clusters who do not want to administrator individual licenses.
- An evaluation license enables you to use a feature for a few months in order to make a purchasing decision.
- Capacity licenses are a pay-as-you-use-it license used in storage-as-a-service implementations.


In ONTAP 8.2 through ONTAP 9.9.1, license keys were delivered as 28-character strings, and there is one key per ONTAP feature.

Beginning with ONTAP 9.10.1, all license are delivered through NetApp License Files (NLF). NLF licenses can enable one or more ONTAP features, depending on your purchase. You can retrieve NLF licenses from the NetApp Support Site by searching for the system (controller) serial number.



## Standard and site licenses

- Proof of sale is recorded as a license entitlement record.
- License keys are 28 characters long.
- Standard licenses are linked to the controller serial number (node locked).
- Features are licensed on every node and continue to function if one licensed node is running.
- Site licenses enable the feature on the entire cluster.
  - A site license is not carried with nodes that are removed from the cluster.




**NetApp** 20 © 2023 NetApp, Inc. All rights reserved.

**Standard license:** A standard license is issued for a node with a specific system serial number and is valid for only the node that has the matching serial number. Installing a standard, node-locked license entitles a node, but not the entire cluster, to the licensed functionality. For the cluster to be enabled, though not *entitled*, to use the licensed functionality, at least one node must be licensed for the functionality. However, if only one node in a cluster is licensed for a feature, and that node fails, the feature no longer functions on the rest of the cluster until the licensed node is restarted. The `system license show` command displays standard licenses as type “license.”

**Site license:** A site license is not tied to a specific system serial number. When you install a site license, all nodes in the cluster are entitled to the licensed functionality. The `system license show` command displays site licenses under the cluster serial number. If your cluster has a site license, and you remove a node from the cluster, the node does not carry the site license with it. The node is no longer entitled to the licensed functionality. If you add a node to a cluster that has a site license, the node is automatically entitled to the functionality that the license grants.

## Evaluation license


- Enables testing of software functionality before purchasing the license
- Is a time-limited license
- Can be renewed a limited number of times before a purchase is required



An evaluation license enables you to try certain software functionality without purchasing an entitlement. If your cluster has an evaluation license for a package, and you remove a node from the cluster, the node does not carry the evaluation license. The best use of evaluation licenses is for proof-of-concept testing on test and development clusters rather than on a production cluster.

## Capacity licenses

- Capacity licenses are sold individually for increments of storage capacity (500TB, 100TB, 50TB, and so on).
- These licenses are used with NetApp ONTAP Select, NetApp Cloud Volumes ONTAP, and FabricPool functionality.
- Additional capacity can be added to a capacity pool license at any time.
- Enforcement is performed at the aggregate level and relies on an aggregate lease.
- An expired lease prevents users from bringing aggregates back online after a manual reboot.
- License codes are shorter than 28 characters.

 22 © 2023 NetApp, Inc. All rights reserved.



Capacity licenses are additional license requirements on a cluster on which storage capacity is sold in increments. ONTAP Select, Cloud Volumes ONTAP, and FabricPool technology all require capacity licenses.

To increase the amount of storage capacity in the cluster, you must purchase a license for the increment or increments of capacity that you need.

If the lease on an aggregate expires, rebooting the system makes the aggregate inaccessible.

Unlike standard, site, and evaluation licenses, the capacity licenses are not 28 characters long.

## License commands

The screenshot displays the ONTAP System Manager interface. On the left is a navigation sidebar with categories like Dashboard, Insights, Storage, Network, Events & Jobs, Protection, Hosts, and Cluster. The main area shows 'Licenses' with a '+ Add' button highlighted. An 'Add License' dialog box is open, containing a text input field for license keys. Below the dialog, a terminal window shows the following commands and their descriptions:

```
cluster2::> license ?
(system license)
add                Add one or more licenses
capacity>         The capacity directory
clean-up           Remove unnecessary licenses
delete             Delete a license
entitlement-risk> The entitlement-risk directory
show              Display licenses
show-status       Display license status
status>           Display license status
```

All licenses for a given order are included on the system boot media at the factory, and ONTAP software retrieves and applies them automatically during cluster initialization.

ONTAP software also enables you to manage feature licenses in the following ways:

- Add one or more license keys.
- Display information about installed licenses.
- Display the packages that require licenses and the current license status of the packages on the cluster.
- Delete a license from a cluster or from the node with the serial number that you specify.

**Note:** The cluster base license is required for the cluster to operate. ONTAP software does not enable you to delete the license.


- Display or remove expired or unused licenses.

Starting with ONTAP 9.10.1, ONTAP feature license keys will be provided to customers in a license file that can be installed into their ONTAP cluster in a single step.



## Lesson 3

### Policies and schedules

 24 © 2023 NetApp, Inc. All rights reserved.

## Policy-based storage services


### Policy:

- Is a collection of rules that the cluster or storage VM administrator creates and manages
- Are predefined or created for managing data access

### Policy examples:

- Health and security
- Export, quota, file, and data
- Snapshot and SnapMirror
- Quality of service (QoS)



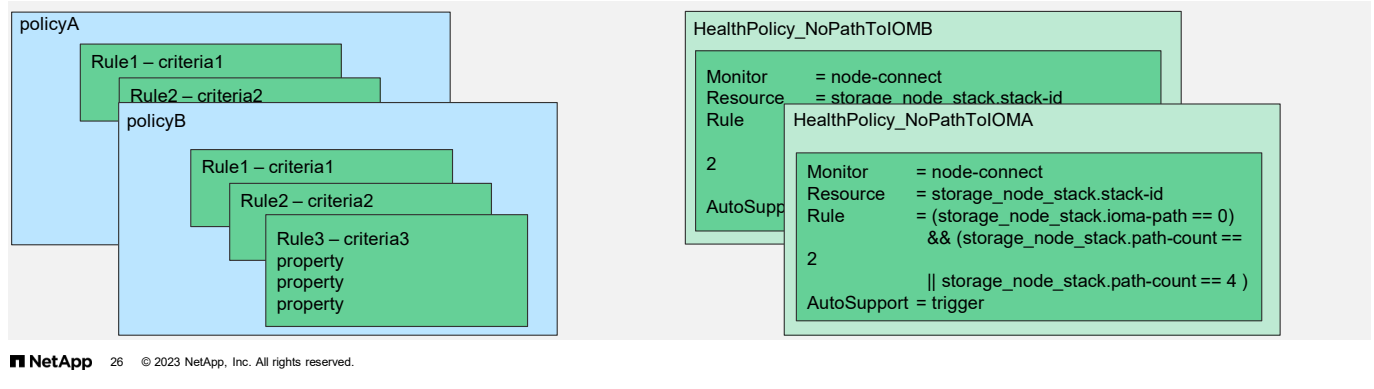
 25 © 2023 NetApp, Inc. All rights reserved.

The following services are policy-based:

- System health
- SnapMirror
- Volume efficiency
- Volume FlexCache
- Volume quota
- Volume Snapshot
- SVM CIFS group
- SVM data
- SVM export
- SVM FPolicy
- SVM security file directory
- Quality of service (QoS) policy group
- Failover

## Policy-based management

- You assign a policy to a service or resource.
- A rule criterion in the policy matches the service or resource.
- The matching rule properties apply to the service or resource.
- The example is a health policy that triggers sending an AutoSupport message when a storage pathway is lost.



Storage VMs use policy-based management for many resources. A policy is a collection of rules or properties that the cluster administrator or storage VM administrator creates and manages. Policies are predefined as defaults or are created to manage various resources. By default, a policy applies to the current resources and to newly created resources, unless otherwise specified.

For example, Snapshot policies can be used to schedule automatic controller-based Snapshot copies. The policy includes such things as the schedule or schedules to use and how many copies to retain. When a volume is created for the storage VM, the policy is applied automatically, but it can be modified later.

The efficiency policy is used to schedule postprocess deduplication operations. The policy might include when and how long deduplication runs.

The examples are only two of the policies that you encounter in ONTAP software. The advantage of policy-based management is that when you create a policy, you can apply the policy to any appropriate resource, either automatically or manually. Without policy-based management, you must enter the settings separately for each individual resource.

# Jobs

- Asynchronous tasks
- Managed by the Job Manager
- Long-running operations
- In a job queue

Job Description	State	Mess...	Start Time	End Time
PATCH /api/storage/volumes/9a2dcb1c-d9...	success	success	2/8/2022, 12:53 AM	2/8/2022, 12:53 AM
PATCH /api/storage/volumes/9a2dcb1c-d9...	success	success	2/8/2022, 12:52 AM	2/8/2022, 12:52 AM
POST /api/cluster/ntp/servers/192.168.0.11	failure	Unabl...	2/15/2022, 9:33 PM	2/15/2022, 9:33 PM

```

::> job show
Job ID Name                                     Owing
-----
1      SnapMirror Service Job.                   cluster1
      Description: SnapMirror Service Job        cluster1-01
      State: Dormant
3      Certificate Expiry Check                  -
      Description: Certificate Expiry Check
5      Auto Balance Aggregate Analyzer           -
      State: Paused
  
```

A job is any asynchronous task that Job Manager manages. Jobs are typically long-running volume operations such as copy, move, and mirror. Jobs are placed in a job queue. Jobs run in the background when resources are available. If a job consumes too many cluster resources, you can stop or pause the job until there is less demand on the cluster. You can also monitor jobs, view job history, and restart jobs.



# Schedules

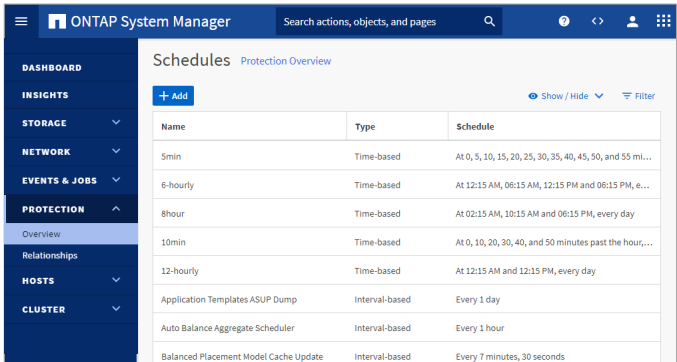
Schedules for tasks:

- Time-based schedules, which run at specific times (similar to UNIX cron schedules)
- Interval-based schedules, which run at intervals

```

::> job schedule show
Name           Type      Description
-----
5min           cron     @:00, :05, :10, :15, :20, :25, :30, :35
8hour         cron     @2:15,10:15,18:15
Auto Balance  Aggregate Scheduler
RepositoryBalanceMonitorJobSchedule
interval      Every 1h
RepositoryBalanceMonitorJobSchedule
interval      Every 10m
daily         cron     @0:10
hourly        cron     @:05
monthly       cron     1@0:20
weekly        cron     Sun@0:15

```



The screenshot shows the ONTAP System Manager interface. The left sidebar contains navigation menus for Dashboard, Insights, Storage, Network, Events & Jobs, Protection, Hosts, and Cluster. The main content area is titled 'Schedules' and includes a table with the following data:

Name	Type	Schedule
5min	Time-based	At 0, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50, and 55 mi...
6-hourly	Time-based	At 12:15 AM, 06:15 AM, 12:15 PM and 06:15 PM, e...
8hour	Time-based	At 02:15 AM, 10:15 AM and 06:15 PM, every day
10min	Time-based	At 0, 10, 20, 30, 40, and 50 minutes past the hour,...
12-hourly	Time-based	At 12:15 AM and 12:15 PM, every day
Application Templates ASUP Dump	Interval-based	Every 1 day
Auto Balance Aggregate Scheduler	Interval-based	Every 1 hour
Balanced Placement Model Cache Update	Interval-based	Every 7 minutes, 30 seconds

Many tasks, such as volume Snapshot copies, can be configured to run on specified schedules. Schedules that run at specific times are called cron schedules. The schedules are like UNIX cron schedules. Schedules that run at intervals are called interval schedules.

To manage schedules in System Manager, on the cluster Configuration tab, click the **Schedules** link. You can create, edit, and delete schedules.


## Module summary

This module focused on enabling you to do the following:

- Manage access control
- Set the date and time on cluster nodes
- Manage NetApp ONTAP software licenses
- Manage jobs and schedules

# Knowledge check

Module 3: Cluster management

 30 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

**The admin storage VM manages the cluster and serves data.**

- a. true
- b. false

## Knowledge check

**What are four valid types of ONTAP licenses? (Choose four.)**

- a. capacity
- b. site
- c. evaluation
- d. expansionary
- e. provisional
- f. standard



## Complete exercises

Module 3  
Cluster management

**Managing ONTAP clusters**

**Managing ONTAP administrators**

**Configuring multi-admin verification**

- Access your lab equipment.
- Open your Exercise Guide, Module 3.
- Complete Exercise 1.
- Share your results.

This exercise requires approximately  
**35 minutes.**

See the instructions in your Exercise Guide.



## Share your experiences


Roundtable discussion

- How did the cluster behave after you specified the NTP server?
- Did the time synchronize immediately?

Have a roundtable discussion with the class to answer these questions. Add any comments about experiences or “lessons learned” during the exercises that others might find helpful.

# Module 4

## Network management


 1 © 2023 NetApp, Inc. All rights reserved.



## About this module


This module focuses on enabling you to do the following:

- Describe the interaction between physical and virtual networking resources in a cluster
- Configure and manage physical and virtual networking resources

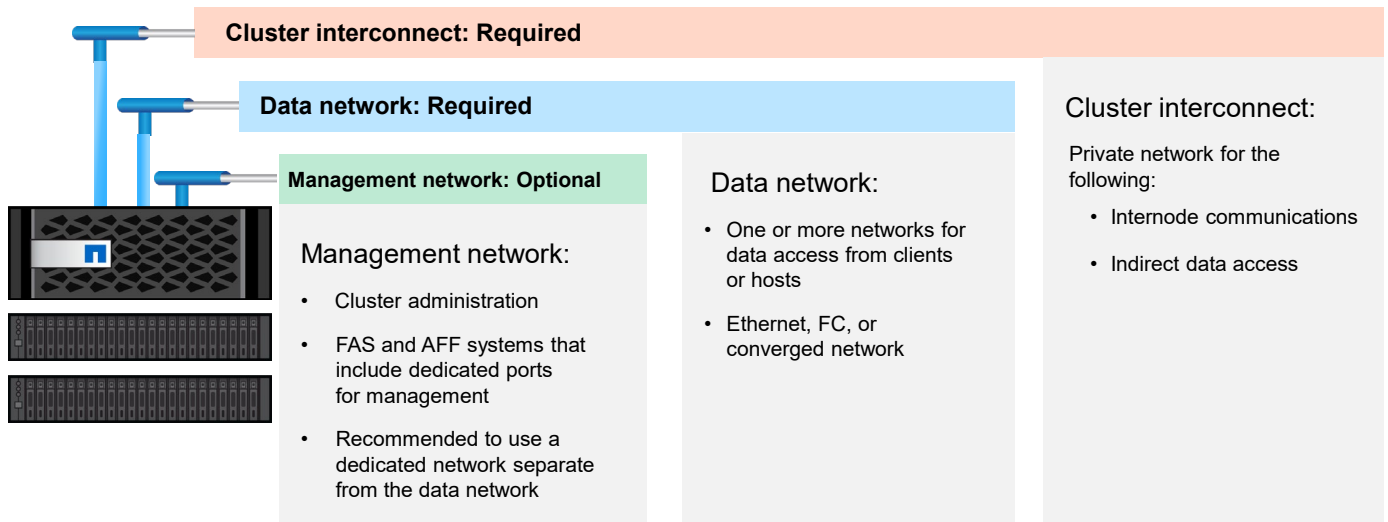


# Lesson 1

## NetApp ONTAP network review

 3 © 2023 NetApp, Inc. All rights reserved.

# Networks



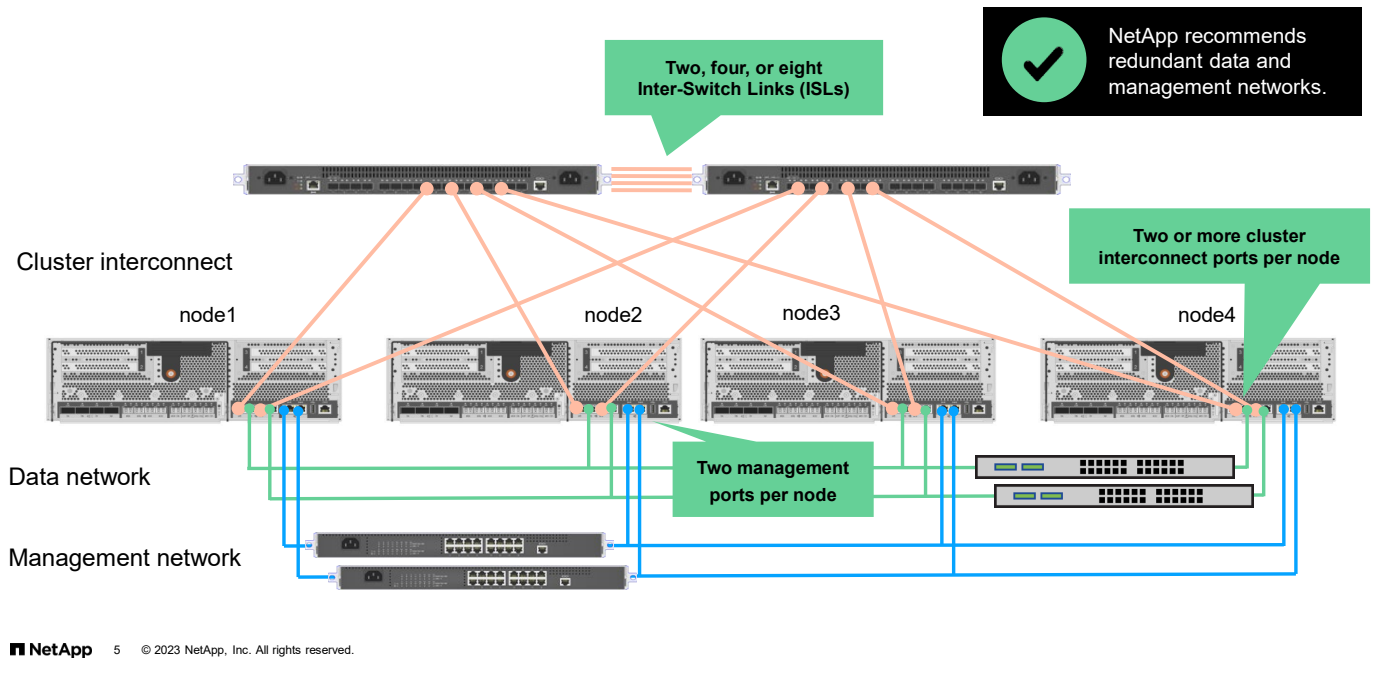
NetApp 4 © 2023 NetApp, Inc. All rights reserved.

In multi-node clusters, nodes need to communicate with each other over a cluster interconnect. In a 2-node cluster, the interconnect can be switchless. When you add more than two nodes to a cluster, a private cluster interconnect that uses switches is required.

The management network is used for cluster administration. Redundant connections to the management ports on each node and management ports on each cluster switch should be provided to the management network. In smaller environments, the management and data networks might be on a shared Ethernet network.

For clients and hosts to access data, a data network is required. The data network can be made up of one or more physical or virtual networks. Depending on the environment, the network might be an Ethernet, FC, or a converged network. Data networks can consist of one or more switches or redundant networks.

## Networks



A NetApp ONTAP software cluster is essentially a cluster of high-availability (HA) pairs. Therefore, you need a cluster interconnect for all the nodes to communicate with one another. If a node cannot see the cluster interconnect, the node is not part of the cluster. Therefore, the cluster interconnect requires adequate bandwidth *and* resiliency.

The figure shows a 4-node cluster and three distinct networks. ONTAP software requires both data and management connectivity, which can coexist on the same data network.


In multinode configurations, ONTAP software also requires a cluster interconnect for cluster traffic. In a 2-node configuration, the cluster interconnect can be as simple as to cable the two nodes or to use switches if expansion is desired. In clusters of more than two nodes, switches are required. For redundancy, you should always have at least one cluster port per switch on each node of the cluster. The number of cluster ports per node depends on the controller model and port speed.

Single-node clusters do not require a cluster interconnect, if the environment does not require high availability and nondisruptive operations (NDO).

For site requirements, switch information, port cabling information, and controller onboard port cabling, see the Hardware Universe at [hwu.netapp.com](http://hwu.netapp.com).



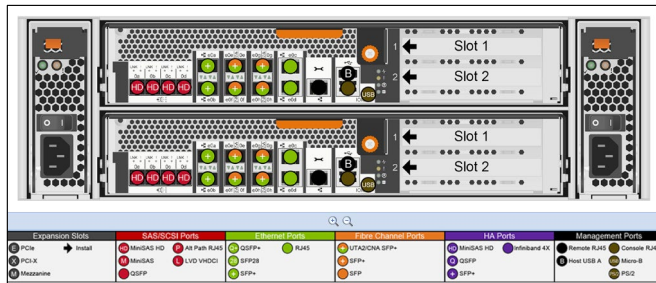
## Lesson 2 Network ports

 6 © 2023 NetApp, Inc. All rights reserved.

## Physical ports example

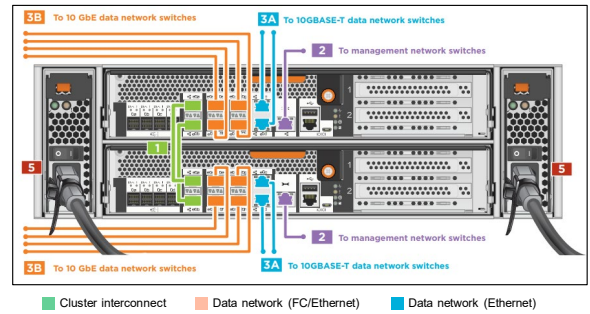
### Hardware Universe

- Focused on port identification
- Downloadable Visio-template-based picture



### Installation and setup instructions (ISI)

- Focused on cabling
- PDF on the NetApp Support site




NetApp 7 © 2023 NetApp, Inc. All rights reserved.

The numbers and types of ports vary by system model, but most systems have dedicated ports for connecting to external drive shelves, Ethernet networks, FC networks, and management networks. There are two primary sources for identifying ports and their use on an AFF or FAS system. In addition to all the technical details, the Hardware Universe includes Visio-template-based diagrams of the front and back of the storage controller. To see how the ports are to be cabled, the Installation and Setup Instructions (ISI) is the best source.

## Physical port identification

- Ethernet ports are named e<location><letter>:
  - e0a is the first port on the controller motherboard.
  - e3a is the first port on the card in slot 3.
- FC ports are named <location><letter>:
  - 0a is the first port on the controller motherboard.
  - 3b is the second port on the card in slot 3.
- Unified target adapter (UTA) ports have both an Ethernet name and an FC name, e <location><letter> / <location><letter>
  - e4a / 4a is the first port on the card in slot 4.
  - Use of show commands returns only FC label names (even in Ethernet mode).



Port names consist of two or three characters that describe the port type and location. You should remember port-naming conventions on the network interfaces.

**Ethernet ports:** The first character describes the port type and is always “e” to represent Ethernet. The second character is a numeral that identifies the slot in which the port adapter is located. The numeral 0 (zero) indicates that the port is on the node motherboard. The third character indicates the port position on a multiport adapter. For example, the port name e0b indicates the second Ethernet port on the motherboard, and the port name e3a indicates the first Ethernet port on an adapter in slot 3.

**FC ports:** The name consists of two characters (dropping the e) but otherwise follows the same naming convention as Ethernet ports. For example, the port name 0b indicates the second FC port on the motherboard. The port name 3a indicates the first FC port on an adapter in slot 3.

**Unified target adapter (UTA) ports:** A UTA port is physically one port but can pass either Ethernet traffic or FC traffic. Therefore, UTA ports are labeled with both the Ethernet name and the FC name. For example, the port name e0b/0b indicates the second UTA port on the motherboard. The port name e3a/3a indicates the first UTA port on an adapter in slot 3.

## Modifying network port attributes

Change the personality of a UTA port




Insert the proper optical module before changing modes.

1. Remove LIFs or migrate LIFs to other ports.
2. Take the port offline.
3. Change the personality of the UTA port.

```
cluster2::> system node hardware unified-connect modify -node cluster2-01 -adapter 0e  
-mode fc|cna
```

4. Reboot the cluster node.

```
cluster2::> system node reboot -node cluster2-01
```

 9 © 2023 NetApp, Inc. All rights reserved.

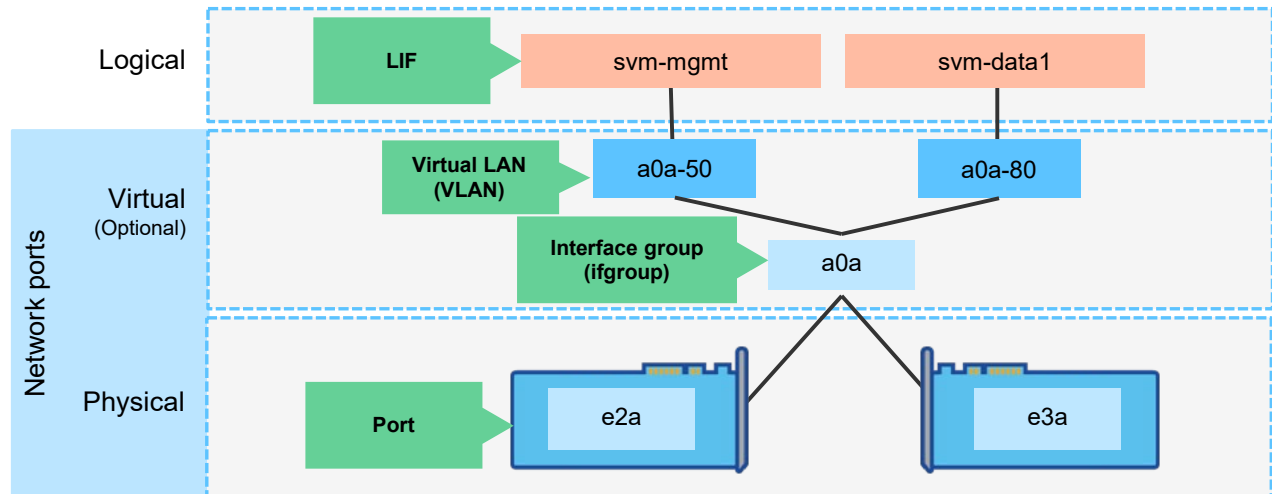
You must use the CLI to change the personality of a UTA port and then reboot the cluster node for the change to take effect. The adapter must also be offline before you can make changes.

- When the adapter type is initiator, use the `run local storage disable adapter` command to take the adapter offline.
- When the adapter type is target, use the `network fcp adapter modify` command to take the adapter offline.

For more information about configuring FC ports, see the *ONTAP SAN Administration Guide* for your release, or attend the NetApp Learning Services *SAN Implementation* course.



## Virtual network ports



NetApp 10 © 2023 NetApp, Inc. All rights reserved.

Nodes have physical ports that are available for cluster traffic, management traffic, and data traffic. The ports must be configured appropriately for the environment. The example shows Ethernet ports. Physical ports also include FC ports and UTA ports.


Physical Ethernet ports can be used directly or combined by using interface groups (ifgroups). Also, physical Ethernet ports and ifgroups can be segmented by using virtual LANs (VLANs). VLANs and ifgroups are considered virtual ports but are treated like physical ports.

Unless specified, the term *network port* includes physical ports, ifgroups, and VLANs.

## Interface groups

- Combination of one or more Ethernet interfaces
- Three interface group (ifgroup) modes:
  - Single-mode (active-standby)
  - Static multimode (active-active)
  - Dynamic multimode with Link Aggregation Control Protocol (LACP)
- Naming syntax:  
a <number> <letter>  
(for example, a0a)

**NOTE:** Vendors might use other terms for combining Ethernet interfaces (for example, Cisco EtherChannel).



An ifgroup combines one or more Ethernet interfaces, which can be implemented in one of three ways.

In single mode, one interface is active, and the other interfaces are inactive until the active link goes down. The standby paths are used only during a link failover.

In static multimode, all links are active. Therefore, static multimode provides link failover and load-balancing features. Static multimode complies with the IEEE 802.1ax (static) standard and works with any switch that supports the combination of Ethernet interfaces. However, static multimode does not have control packet exchange.

Dynamic multimode is similar to static multimode but complies with the IEEE 802.1ax (dynamic) standard. When switches that support Link Aggregation Control Protocol (LACP) are used, the switch can detect a loss of link status and dynamically route data. NetApp recommends that when you configure ifgroups, you use dynamic multimode with LACP and compliant switches.

All modes support the same number of interfaces per ifgroup, but the interfaces in the group should always be the same speed and type. The naming syntax for interface groups is the letter “a,” followed by a number, followed by a letter (for example, a0a).

Vendors might use terms such as link aggregation, port aggregation, trunking, bundling, bonding, teaming, or EtherChannel.

## Creating ifgroups

- Create a link aggregation group.

```
cluster2::> network port ifgrp create
-node cluster2-01 -ifgrp a0a
-distr-func ip -mode multimode
```

- Assign network ports to the link aggregation group.

```
cluster2::> network port ifgrp add-port
-node cluster2-01 -ifgrp a0a -port e0f
```

The name of the ifgroup must be in a<number><letter> format.

**Add Link Aggregation Group**

NODE  
cluster1-02

PORTS TO INCLUDE  
 e0e  e0d  e0f

MODE  
 Single  
Only one port is used at a time.  
 Multiple  
All ports can be used simultaneously.  
 LACP  
The LACP protocol determines the ports that can be used.

LOAD DISTRIBUTION  
 IP based  
Network traffic is distributed based on the destination IP address.  
 MAC based  
Network traffic is distributed based on the next-hop MAC addresses.  
 Sequential  
Network traffic is distributed by round-robin over the outbound links.  
 Port  
Network traffic is distributed based on the transport layer (TCP/UDP) ports.

Save Cancel

NetApp 12 © 2023 NetApp, Inc. All rights reserved.

You can create ifgroups for higher throughput, fault tolerance, and elimination of single points of failure (SPOFs).


You manage ifgroups like you manage physical interfaces, except for the following:

- You must name ifgroups by using the syntax *a<number><letter>*.
- You cannot add a port that is already a member of one ifgroup to another ifgroup.
- Multimode load-balancing methods include the following:
  - MAC: Network traffic is distributed by MAC addresses.
  - IP: Network traffic is distributed by IP addresses.
  - Sequential: Network traffic is distributed as it is received.
  - Port: Network traffic is distributed by the transport layer (TCP/UDP) ports.

A single host does not necessarily achieve larger bandwidth, exceeding the capabilities of any constituent connections. For example, adding four 10GbE ports to a dynamic multimode ifgroup does not result in one 40GbE link for one host. The situation is because of the way that both the switch and the node manage the aggregation of the ports in the ifgroup.

For more information about load balancing, see TR-4182: Ethernet Storage Design Considerations and Best Practices for Clustered Data ONTAP Configurations.

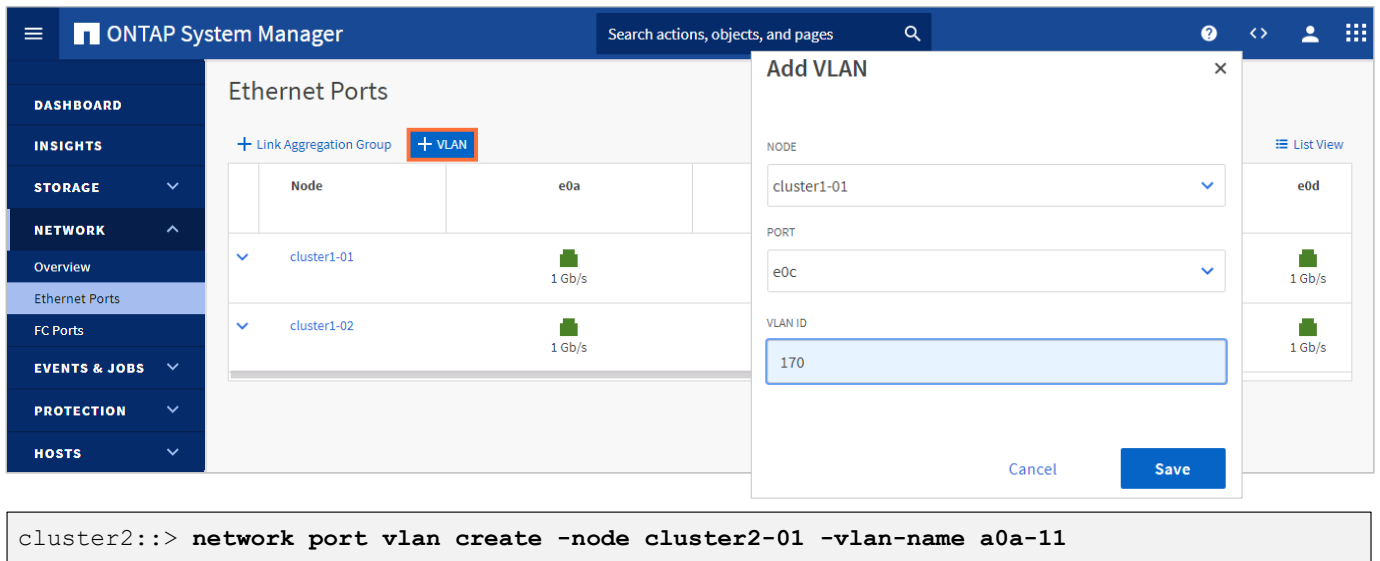
## VLANs



NetApp 13 © 2023 NetApp, Inc. All rights reserved.

A port or ifgroup can be subdivided into multiple VLANs. Each VLAN has a unique tag that is communicated in the header of every packet. The switch must be configured to support VLANs and the tags that are in use. In ONTAP software, a VLAN ID is configured into the name. For example, VLAN e0a-70 is a VLAN with tag 70 that is configured on physical port e0a. VLANs that share a base port can belong to the same IPspace or different IPspaces. The base port can be in a different IPspace than the VLANs that share the base port. IPspaces are discussed later in this module.

## Creating VLANs




The screenshot displays the ONTAP System Manager interface. The main content area shows the 'Ethernet Ports' page with a table listing nodes and their network speeds. A '+ VLAN' button is highlighted. An 'Add VLAN' dialog box is open, allowing the user to specify the node, port, and VLAN ID. Below the dialog, a terminal window shows the command used to create a VLAN on a specific node.

Node	e0a
cluster1-01	1 Gb/s
cluster1-02	1 Gb/s

**cluster2::> network port vlan create -node cluster2-01 -vlan-name a0a-11**

You can create a VLAN for ease of administration, confinement of broadcast domains, reduced network traffic, and enforcement of security policies.

## Ports, ifgroups, and VLAN combinations



NetApp 15 © 2023 NetApp, Inc. All rights reserved.

Most small to medium environments and FC environments use physical ports.

Ethernet environments in which multiple physical networks are impossible often use VLANs to separate management traffic from data traffic. VLANs are also often used to separate differing workloads. For example, you might separate NAS traffic from iSCSI traffic for performance and security reasons.

In Ethernet environments in which many application servers or hosts share switches and ports, dynamic multimode ifgroups of four Ethernet ports per node are frequently used for load balancing.

Environments that use ifgroups typically also use VLANs to segment the network. Segmentation is typical for service providers with multiple clients that require the bandwidth that ifgroups provide and the security that VLANs provide.

Finally, it is not unusual for different types of ports to be used in mixed environments that have various workloads. For example, an environment might use ifgroups with VLANs that are dedicated to NAS protocols, a VLAN that is dedicated to management traffic, and physical ports for FC traffic.

Ifgroups and VLANs cannot be created on cluster interconnect ports.



## Complete an exercise

Module 4  
Network management

### Managing Physical and Logical Network Resources


- Access your lab equipment.
- Open your Exercise Guide, Module 4.
- Complete Exercise 1.
- Share your results.

This exercise requires approximately  
**20 minutes.**



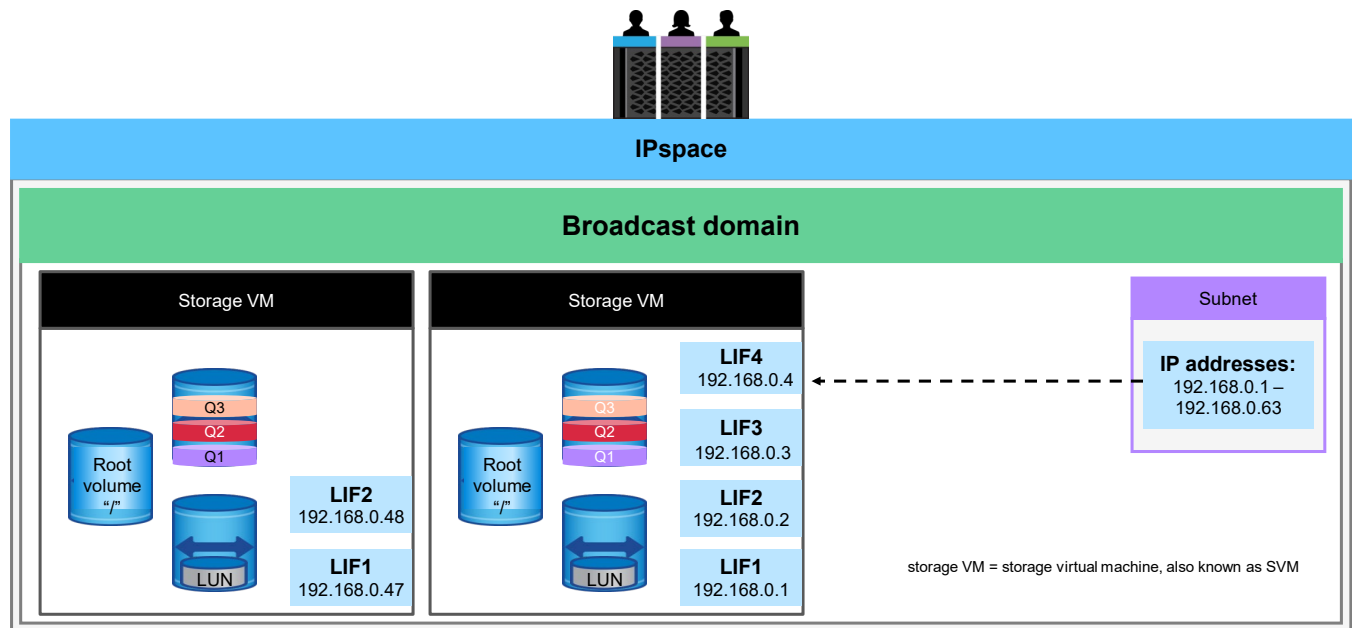
## Lesson 3

# Network traffic segregation

 17 © 2023 NetApp, Inc. All rights reserved.



## IPspace review



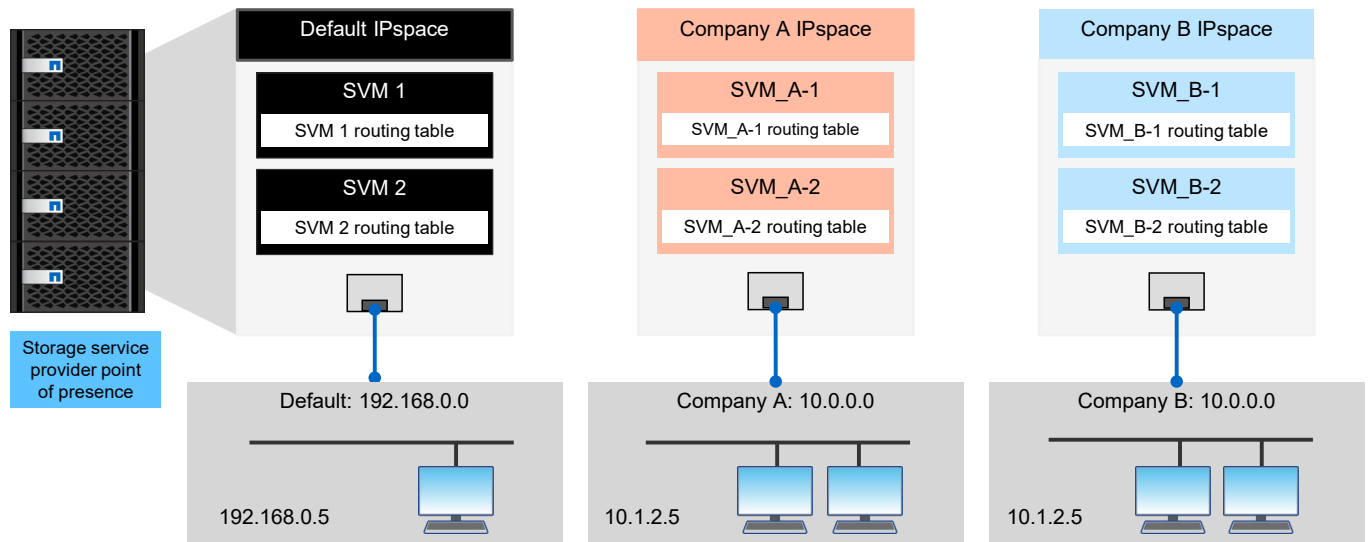
NetApp 18 © 2023 NetApp, Inc. All rights reserved.

ONTAP software has a set of features that work together to enable multitenancy. An IPspace is a logical container that is used to create administratively separate network domains. An IPspace defines a distinct IP address space that contains storage VMs (storage virtual machines, also known as SVMs). The IPspace contains a broadcast domain that determines which network ports are accessible to the storage VMs in the IPspace. The broadcast domain can contain a subnet, which enables you to allocate a pool of IP addresses for your ONTAP network configuration.

When you create a LIF on the storage VM, the LIF represents a network access point to the node. You can manually assign the IP address for the LIF. If a subnet is specified, the IP address is automatically assigned from the pool of addresses in the subnet. This process is much like how a Dynamic Host Configuration Protocol (DHCP) server assigns IP addresses.

# IPspaces

## Segregating networks



NetApp 19 © 2023 NetApp, Inc. All rights reserved.

The IPspace feature enables clients from more than one disconnected network to access a storage system or cluster, even if the clients use the same IP address.

An IPspace defines a distinct IP address space in which virtual storage systems can participate. IP addresses that are defined for an IPspace are applicable only within the IPspace. A distinct routing table is maintained for each IPspace. No cross-IPspace traffic routing occurs. Each IPspace has a unique assigned loopback interface. The loopback traffic on each IPspace is isolated from the loopback traffic on other IPspaces.

### Example

A storage service provider needs to connect customers of companies A and B to a storage system on the storage service provider premises. The storage service provider creates storage VMs on the cluster, one per customer. The storage service provider then provides one dedicated network path from one storage VM to the A network and another dedicated network path from the other storage VM to the B network.

The deployment should work if both companies use nonprivate IP address ranges. However, because the companies use the same private addresses, the storage VMs on the cluster at the storage service provider location have conflicting IP addresses.

To overcome the problem, two IPspaces are defined on the cluster, one per company. Because a distinct routing table is maintained for each IPspace, and no cross-IPspace traffic is routed, the data for each company is securely routed to the respective network. Data is securely routed even if the two storage VMs are configured in the 10.0.0.0 address space.

Also, the IP addresses that are referred to by various configuration files (the `/etc/hosts` file, the `/etc/hosts.equiv` file, the `/etc/rc` file, and so on) are relative to the IPspace. Therefore, the IPspaces enable the storage service provider to use the same IP address for the configuration and authentication data for both storage VMs without conflict.

The ONTAP IPspace feature is similar to Cisco MPLS VPN technology.

# Managing IPspaces

## Create

You can create IPspaces when you need your storage VMs to have distinct and secure storage, administration, and routing.

The screenshot shows the ONTAP System Manager interface. The left sidebar contains navigation options: DASHBOARD, INSIGHTS, STORAGE, NETWORK, Overview, Ethernet Ports, FC Ports, EVENTS & JOBS, PROTECTION, and HOSTS. The main content area is titled 'Overview' and displays 'IPspaces' and 'Broadcast Domains'. An 'Add IPspace' dialog box is open, showing a 'NAME' field with the value 'ipCompanyA' and 'Save' and 'Cancel' buttons. Below the dialog box, a terminal window shows the command: `cluster1::> network ipspace create -ipspace ipCompanyB`

NetApp 20 © 2023 NetApp, Inc. All rights reserved.

IPspaces are distinct IP address spaces in which storage VMs reside. All IPspace names must be unique within a cluster.

- If necessary, you can change the name of an existing IPspace (except for the two system-created IPspaces) by using the `network ipspace rename` command.
- If you no longer need an IPspace, you can delete the IPspace by using the `network ipspace delete` command.

**Note:** No broadcast domains, network interfaces, or storage VMs can be associated with an IPspace that you want to delete. You cannot delete the system-defined Default and cluster IPspaces.

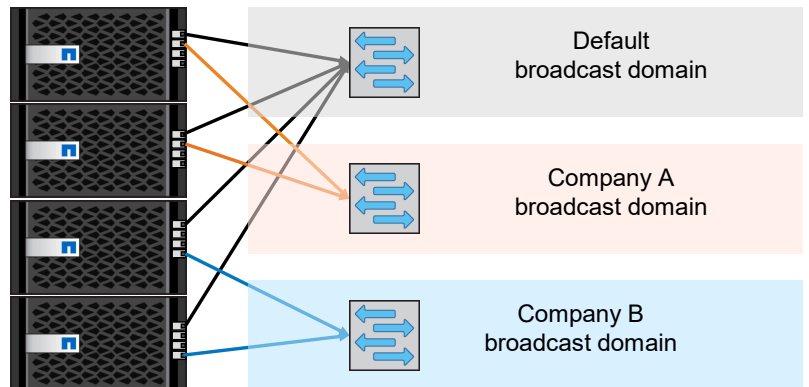
You can display the list of IPspaces that exist in a cluster. You can also view the storage VMs, broadcast domains, and ports that are assigned to each IPspace.

When you create an IPspace, a broadcast domain with the same name is automatically created within the IPspace. You then must assign some network ports to the broadcast domain for the IPspace to become operational.

## Broadcast domains

### Overview

- Broadcast domains enable you to group network ports that belong to the same Layer 2 network.
- A storage VM can then use the ports in the group for data or management traffic.



Broadcast domains can contain physical ports, ifgroups, and VLANs.

NetApp 21 © 2023 NetApp, Inc. All rights reserved.

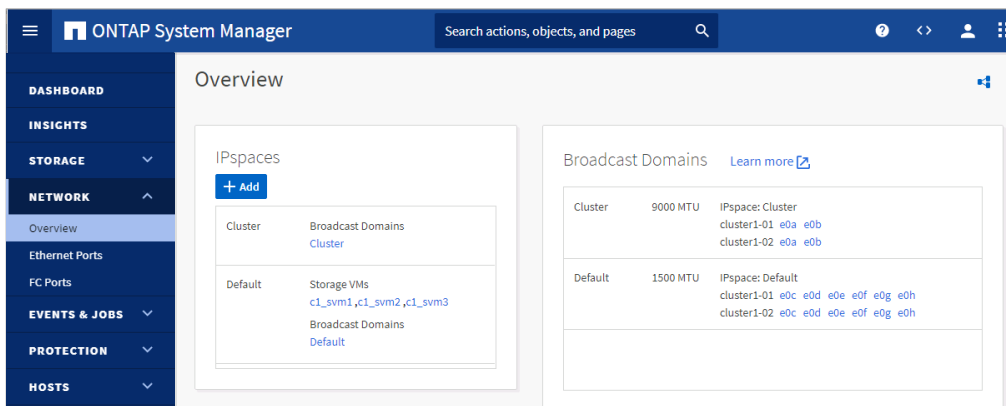
To prevent chaos from using overlapping IP addresses, each IPspace must use different IP networks that are accessed through different network ports. The broadcast domain determines which physical network ports are accessible to the storage VMs in the IPspace. A broadcast domain should include ports from many nodes in the cluster to provide high availability for the connections to storage VMs.

The figure shows the ports that are assigned to three broadcast domains in a 4-node cluster:

- The Default broadcast domain, which was created automatically during cluster initialization, is configured to contain a port from each node in the cluster.
- The Company A broadcast domain was created manually and contains one port each from the nodes in the first HA pair.
- The Company B broadcast domain was created manually and contains one port each from the nodes in the second HA pair.
- The Cluster broadcast domain was created automatically during cluster initialization, but it is not shown in the figure.

The system administrator created the two broadcast domains specifically to support the customer IPspaces.

## Managing broadcast domains



You create broadcast domains to group ports for an IPspace.

```
cluster1::> network port broadcast-domain create -broadcast-domain bdCompanyB -ipspace ipCompanyB  
-mtu 1500 -ports cluster1-01:a0a,cluster1-02:a0a
```

**NetApp** 22 © 2023 NetApp, Inc. All rights reserved.

You create a broadcast domain to group network ports in a cluster that belongs to the same Layer 2 network. Storage VMs can then use the ports.

**Note:** The ports that you plan to add to the broadcast domain must not belong to another broadcast domain.

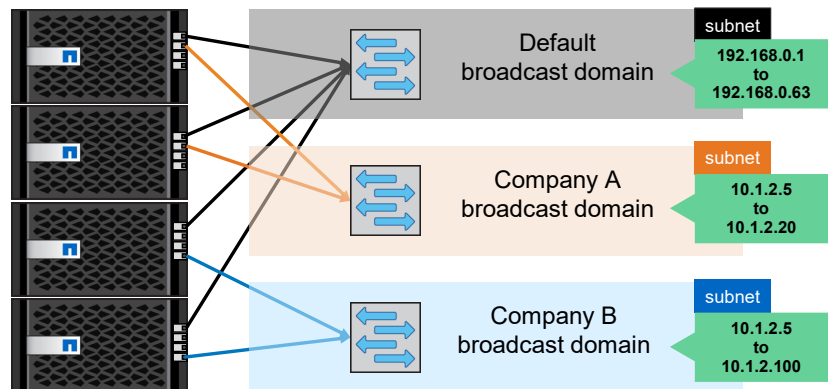
- All broadcast domain names must be unique within an IPspace.
- The ports that you add to a broadcast domain can be network ports, VLANs, or ifgroups.
- You add ports by using the `network port broadcast-domain add-ports` command.
- If the ports that you want to use belong to another broadcast domain but are unused, use the `network port broadcast-domain remove-ports` command to remove the ports from the existing broadcast domain.
- The maximum transmission unit (MTU) value of the ports that you add to a broadcast domain are updated to the MTU value that is set in the broadcast domain.
- The MTU value must match all the devices that are connected to the Layer 2 network.
- If you do not specify an IPspace name, the broadcast domain is created in the Default IPspace.

You can rename or delete broadcast domains that you create but not the system-created Cluster and Default broadcast domains.

To make system configuration easier, a failover group of the same name is created automatically and contains the same ports. All failover groups that relate to the broadcast domain are removed when you delete the broadcast domain.

## Subnets

- Subnets enable the allocation of specific blocks, or pools, of IP addresses for easier LIF creation.
- A subnet is created within a broadcast domain and contains a pool of IP addresses that belong to the same Layer 3 subnet.



Subnets are recommended for easier LIF creation.

**NetApp** 23 © 2023 NetApp, Inc. All rights reserved.

Subnets enable you to allocate specific blocks, or pools, of IP addresses for your ONTAP network configuration. The allocation enables you to create LIFs more easily when you use the `network interface create` command by specifying a subnet name instead of specifying IP address and network mask values.

IP addresses in a subnet are allocated to ports in the broadcast domain when LIFs are created. When LIFs are removed, the IP addresses are returned to the subnet pool and are available for future LIFs.

You should use subnets, because subnets simplify the management of IP addresses and the creation of LIFs when you use the CLI or APIs.

## Creating subnets

- The IPspace and broadcast domain must exist before the subnet can be created.
- Subnet names must be unique within an IPspace.
- IP addresses in the specified range must not be in use by a LIF.  
Use the `-force-update-lif-associations` option to override the rule.

```
cluster1::> network subnet create -subnet-name subnet_A -broadcast-domain bdCompanyB  
-ip-space ipCompanyB -subnet 10.1.2.0/24 -gateway 10.1.2.1  
-ip-ranges "10.1.2.65-10.1.2.126,10.1.2.193-10.1.2.254" -force-update-lif-associations true
```

You can create a subnet to allocate, or reserve, specific blocks of IPv4 or IPv6 addresses for ONTAP network configuration.

When you create subnets, consider the following limitations:

- When you add IP address ranges to a subnet, no IP addresses in the network can overlap (so that different subnets, or hosts, do not attempt to use the same IP address).
- The command fails if any network interfaces currently use the IP addresses in the specified range. The command option `-force-update-lif-associations` can be used to override this behavior and automatically associate any previously assigned addresses with the subnet and prevent them from being assigned to new network interfaces.
- If you enter the address of a gateway, a default route to the gateway is created for the storage VM. If you do not use subnets or do not specify a gateway when you define a subnet, you must use the `route create` command to manually add a route to the storage VM.

## Verifying subnets

To view subnet details:

```
::> network subnet show
```

Subnet Name	Subnet	Broadcast Domain	Gateway	Avail/Total	Ranges
subnet_def	192.168.0.0/24	Default	192.168.0.1	10/50	192.168.0.101-192.168.0.150
subnet_A	10.1.2.0/24	bdCompanyA	10.1.2.1	44/62	10.1.2.65-10.1.2.126
subnet_B	10.1.2.0/24	bdCompanyB	10.1.2.1	37/62	10.1.2.65-10.1.2.126

**Subnets A and B have the same subnet address and gateway but different domains.**

**Notice how subnets A and B use overlapping IP ranges.**


In this example showing subnets, notice how subnet A and subnet B use the same subnet numbers and IP address ranges. Having network interfaces in different subnets share IP addresses does not cause a conflict. The reason is that the subnets are in different broadcast domains in different IPspaces and use different network ports and routing.





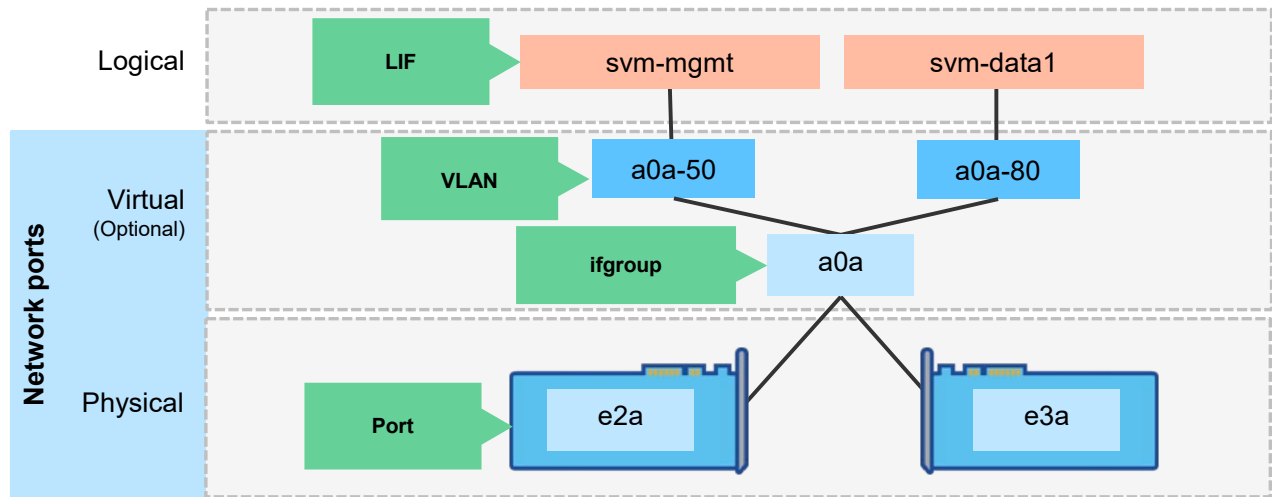
# Lesson 4

## LIFs

 26 © 2023 NetApp, Inc. All rights reserved.

# Network interfaces


Review



NetApp 27 © 2023 NetApp, Inc. All rights reserved.

Recall that a LIF is associated with a physical port, an interface group, or a VLAN. Storage VMs own the LIFs. Multiple LIFs that belong to different storage VMs can reside on a single port.

## Logical network interfaces

 28 © 2023 NetApp, Inc. All rights reserved.

- An IP address or worldwide port name (WWPN) is associated with a LIF:
  - If subnets are configured, an IP address is automatically assigned when a LIF is created. Otherwise, IP addresses must be manually assigned.
  - When an FC LIF is created, WWPNs are automatically assigned.
- One node-management LIF exists per node.
- One cluster-management LIF exists per cluster.
- Multiple data LIFs can be enabled per port (client-facing for NFS, CIFS, S3, iSCSI, FC, and NVMe access).
- For intercluster peering, intercluster LIFs must be created on each node.

Data LIFs can have a many-to-one relationship with network ports. Many data IP addresses can be assigned to a single network port. If the port becomes overburdened, NAS data LIFs can be transparently migrated to different ports or nodes. Clients know the data LIF IP address but do not know which node or port hosts the LIF. If a NAS data LIF is migrated, the client might unknowingly be contacting a different node. The NFS mount point or CIFS share is unchanged.

## Creating data LIFs

- Manually assign an IP address when creating a NAS LIF.  
The subnet mask is determined automatically.
- Specify the subnet name to automatically assign an IP address when creating a NAS LIF.  
You must specify the IP address when subnets are not configured.

Name	Storage VM	Address	Current
cluster_mgmt		192.168.0.101	cluster1
cluster1-02_clus1		169.254.108.211	cluster1
cluster1-02_clus2		169.254.18.79	cluster1
cluster1-02_mgmt1		192.168.0.112	cluster1
c1_svm3_nas_lif2	c1_svm3	192.168.0.120	cluster1
c1_svm1_nas_lif1	c1_svm1	192.168.0.60	cluster1
cluster1-01_mgmt1		192.168.0.111	cluster1
cluster1-01_clus2		169.254.20.16	cluster1
cluster1-01_clus1		169.254.72.91	cluster1
lif_c1_svm6_692	c1_svm6	192.168.0.170	cluster1
c1_svm2_nas_lif1	c1_svm2	192.168.0.61	cluster1

```
cluster1::> network interface create -vserver svm3
-lif svm3_nas_lif2 -data-protocol nfs
-home-node cluster1-01 -home-port e0f -subnet-name sn_svm3
```

NetApp 29 © 2023 NetApp, Inc. All rights reserved.


A LIF is an IP address or worldwide port name (or WWPN) that is associated with a physical port. If a component fails, most LIF types (excluding most SAN types) can fail over to or be migrated to a different physical port. Failover and migration ensure that communication with the cluster continues.

- The underlying physical network port must be configured to the administrative up status.
- If you plan to use a subnet name to allocate the IP address and network mask value for a LIF, the subnet must exist.
- You can create IPv4 and IPv6 LIFs on the same network port.
- You cannot assign both NAS and SAN protocols to a LIF.
  - The supported protocols are CIFS, NFS, Simple Storage Service (S3), FlexCache, iSCSI, FCP, NVMe/FC, and NVMe/TCP.
- The `data-protocol` parameter must be specified when the LIF is created and cannot be modified later.
- The `home-node` parameter is the node to which the LIF returns when the `network interface revert` command is run on the LIF.
- The `home-port` parameter is the port or ifgroup to which the LIF returns when the `network interface revert` command is run on the LIF.
- When you use a subnet to supply the IP address and network mask, if the subnet was defined with a gateway, a default route to that gateway is added automatically to the storage VM when a LIF is created with that subnet.



## Lesson 5


# Nondisruptive LIF configuration

 30 © 2023 NetApp, Inc. All rights reserved.

## Nondisruptive LIF features

### NAS LIF failover

**LIF failover:** Automatic migration that occurs because of a link failure or reboot




NetApp 31 © 2023 NetApp, Inc. All rights reserved.

LIF failover refers to the automatic migration of a NAS LIF to a different network port. A NAS LIF might fail over in response to a node reboot or a link failure on the LIF's current port. LIF failover is a key component for providing highly available connections to storage VMs.

## Nondisruptive LIF features

### NAS LIF migrate

- **LIF failover:** Automatic migration that occurs because of a link failure or reboot
- **LIF migrate:** Manual movement of a LIF to another port



NetApp 32 © 2023 NetApp, Inc. All rights reserved.

LIF migration refers to the manual migration of a NAS LIF to a different network port.


Why migrate a LIF? Migration might be necessary for troubleshooting a faulty port or to offload a node for which data network ports are saturated with other traffic. The LIF fails over automatically if its current node is rebooted.

Unlike storage failover (SFO), LIF failover and migration do not cause a reboot of the node from which the LIF is migrating. After a LIF is migrated, the LIF can remain on the new node for as long as the administrator wants.

## Nondisruptive LIF features

### NAS LIF revert

- **LIF failover:** Automatic migration that occurs because of a link failure or reboot
- **LIF migrate:** Manual movement of a LIF to another port
- **LIF revert:** Manual or automatic sending of a LIF back to the home node and home port



NetApp 33 © 2023 NetApp, Inc. All rights reserved.

A LIF revert moves the NAS LIF back to its home node and home port.

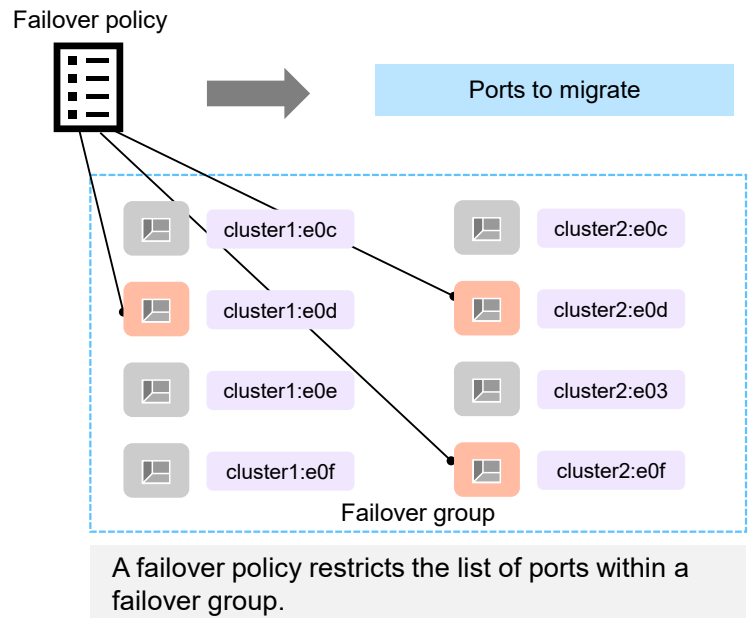
When auto-revert is enabled, a data LIF is automatically reverted to its home node under circumstances such as startup, changes to the status of the management database, or when the network connection is made.



## Failover groups versus failover policies

A failover group is a list of ports (physical or virtual):

- Defines the targets for the LIF
- Is automatically created when you create a broadcast domain
- Does not apply to iSCSI or FC SAN LIFs



NetApp 34 © 2023 NetApp, Inc. All rights reserved.

You can configure a LIF to fail over to a specific group of network ports by applying a failover policy and a failover group to the LIF. A failover group is a collection of network ports. The failover group determines the ports that are available to which the LIF can fail over. The failover policy is evaluated at failover time to select a port from the failover group. You can also disable a LIF from failing over to another port.

**Note:** LIFs for SAN protocols do not support failover. Thus, they are always set to *disabled*.

**Note:** Beginning in ONTAP 9.11.1, on ASA systems only, iSCSI LIFs can fail over to accelerate client recovery.

## Failover groups

Failover groups are created automatically, based on the network ports in the broadcast domain.




A **Cluster failover** group contains all the ports in the Cluster broadcast domain.



A **Default failover** group contains all the ports in the Default broadcast domain.



Additional **failover groups** are created for each broadcast domain that you create.

 35 © 2023 NetApp, Inc. All rights reserved.

There are two types of failover groups. These two groups are groups that the system creates automatically when a broadcast domain is created and groups that a system administrator defines.

The ports in the Cluster broadcast domain are used for cluster communication and include all cluster ports from all nodes in the cluster.

The ports in the Default broadcast domain are used primarily to serve data but also for cluster and node management.

Failover groups have the same name as the broadcast domain and contain the same ports as the groups in the broadcast domain.

## Custom failover groups

You create custom failover groups for specific LIF failover functionality in one or more of the following circumstances:

- The automatic failover groups do not meet your requirements.
- You require only a subset of the ports that are available in the broadcast domain.
- You require consistent performance.

For example, you have configured SnapMirror replication to use high-bandwidth ports. You might create a failover group that consists of only 25GbE ports, to ensure that the LIFs fail over only to other high-bandwidth ports.



More info in Addendum


You can create user-defined failover groups for special failover situations in which the groups that are based on the broadcast domain do not meet your needs.

You create a failover group of network ports so that a LIF can automatically migrate to a different port if a link failure occurs on the LIF's current port. The failover group enables the system to reroute network traffic to other available ports in the cluster.

- The ports that are added to a failover group can be network ports, VLANs, or ifgroups.
- All the ports that are added to the failover group must belong to the same broadcast domain.
- A single port can reside in multiple failover groups.
- If you have LIFs in different VLANs or broadcast domains, you must configure failover groups for each VLAN or broadcast domain.
- Failover groups do not apply in SAN iSCSI or FC environments.

## Failover policies

Failover policy	Available target ports	Details
broadcast-domain-wide	The LIF fails over to any port from any node in the failover group.	Default for cluster-management LIF
system-defined	The LIF fails over only to a port on the home node or a non-storage failover (SFO) partner.	Default for NAS data LIFs <i>Recommended for nondisruptive software updates</i>
local-only	The LIF fails over only to a port on the home node of the LIF.	Default for cluster LIFs, node management LIFs, and intercluster LIFs
sfo-partner-only	The LIF fails over to a port on the home node or HA partner only.	The sfo-partner-only failover policy is not assigned by default.
disabled	Failover is disabled for the LIF.	LIF that is not configured for failover

 37 © 2023 NetApp, Inc. All rights reserved.

The table shows the available failover policies. You cannot create new failover policies.

You can select from several pre-defined failover policies. Usually, you should use the default policy.

- `broadcast-domain-wide`: All ports on all nodes in the failover group
- `system-defined`: Only the ports on the LIF's home node and a non-SFO partner
- `local-only`: Only the ports on the LIF's home node
- `sfo-partner-only`: The LIF fails over to a port on the home node or SFO partner only
- `disabled`: No ports fail over

SAN data LIFs do not failover and are assigned the 'disabled' failover policy.

## Failover policies and groups

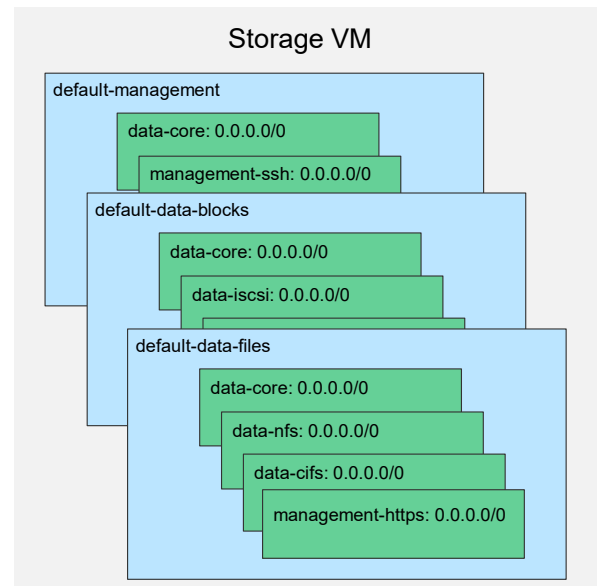
LIF name	LIF type	LIF service policy	Default failover group	Default failover policy
Clus1	cluster	default-cluster	Cluster	local-only
cluster1-01_mgmt1	node management	default-management	Default	local-only
cluster_mgmt	cluster management	default-management	Default	broadcast-domain-wide
svm1_has_lif01	data	default-data-files	Default	system-defined

 38 © 2023 NetApp, Inc. All rights reserved.

The table shows how failover policies and groups work together. Groups include all possible failover targets, whereas policies limit targets within the group.

## Network service policies

- Many ONTAP network services are accessible through logical network interfaces.
- The network service policy defines which network services to provide and from which network addresses to accept requests.
- Each storage VM has its own set of network service policies.
- The network service policy is assigned to a LIF when the LIF is created.



NetApp 39 © 2023 NetApp, Inc. All rights reserved.

Before ONTAP 9.7 software, you needed to specify a role when creating a new logical network interface. The role determined to which network ports the LIF could be bound and the network services that were accessible through that LIF.

Now, it is the broadcast domain that determines to which network ports the LIF can be bound, and the data protocol that determines the network services that are provided.


ONTAP software has a fixed set of network services that it provides. Some network services depend on other network services. For example, the NFS network service depends on the core network service. Service policies group network services so that they can be assigned to LIFs together. Each rule in a service policy specifies the network service to provide and a list of addresses from which requests for that service are accepted.

The admin storage VM and any system storage VM contain service policies that can be used for LIFs in that storage VM, including management and intercluster LIFs. These policies are automatically created by the system when an IPspace is created.

Every data storage VM comes with three built-in network service policies: a service policy for management access, block data access, and file data access. If a list of data services is provided when a storage VM is created, the default data blocks policy, and the default data files policy includes only the network protocols that are specified. If a list of data services is not provided when a storage VM is created, the default data blocks and the default data files policies are set to the network protocols NFS, CIFS/SMB, iSCSI, and FlexCache.


When a LIF is created with a list of data protocols, a service policy equivalent to the specified data protocols is assigned to the LIF. If an equivalent service policy does not exist, a custom service policy is created.

When a LIF is created without a list of data protocols, the default-data-files service policy is assigned to the LIF.




# Lesson 6

## Network security

 40 © 2023 NetApp, Inc. All rights reserved.

## Securing in-flight data

- Data traveling over a network to or from an ONTAP based system is susceptible to interception and theft.
- All NetApp SnapMirror traffic between cluster peers is encrypted.
- Data that is accessed through the HTTPS protocol is encrypted.
- ONTAP software includes the Internet Protocol security (IPsec) for encrypting IP traffic over Ethernet.
  - IPsec supports authentication using pre-shared keys (PSKs) or certificates.
  - IPsec policies that are configured with PSKs require sharing of the key among all clients in the policy.
  - IPsec policies that are configured with certificates do not require sharing of the key among clients because each client can have its own unique certificate for authentication.

 41 © 2023 NetApp, Inc. All rights reserved.



Anytime that data travels over an unsecured network, it is susceptible to interception and theft. The standard defense is to encrypt the data before transmission.

NetApp protocols like SnapMirror software and SnapVault software encrypt all data before the data is transmitted outside the cluster. Some protocols, like HTTPS, have encryption built in, so all data traffic is encrypted. However, older protocols, like NFS and iSCSI, do not include built-in encryption. For these deployments, ONTAP software includes Internet Protocol security (IPsec).

With IPsec, you can encrypt network traffic regardless of the protocol in use. That ability provides simplicity, particularly with NFS environments that choose not to use Kerberos for encryption. You also finally have a way to encrypt iSCSI traffic over the wire.


The ONTAP implementation of IPsec can authenticate the connecting client by using a secret pre-shared key. Beginning with ONTAP 9.10.1, IPsec can also use certificates to authenticate client hosts.

Beginning with ONTAP 9.12.1, NetApp MetroCluster configurations can use IPsec to secure client data traffic.



## Transport Layer Security certificates

- An ONTAP based system uses security certificates to secure access to the following:
  - Command line through the Secure Shell (SSH)
  - System Manager through the HTTPS protocol
- An ONTAP based system creates default certificates during installation:
  - The default certificate is self-signed by the cluster certificate authority (CA).
  - Production environments should use a known commercial CA.
- Client applications might not accept the default certificate.
  - The client might require certificates that are signed by a trusted certificate authority.
  - The client might also require CA certificates.

 42 © 2023 NetApp, Inc. All rights reserved.

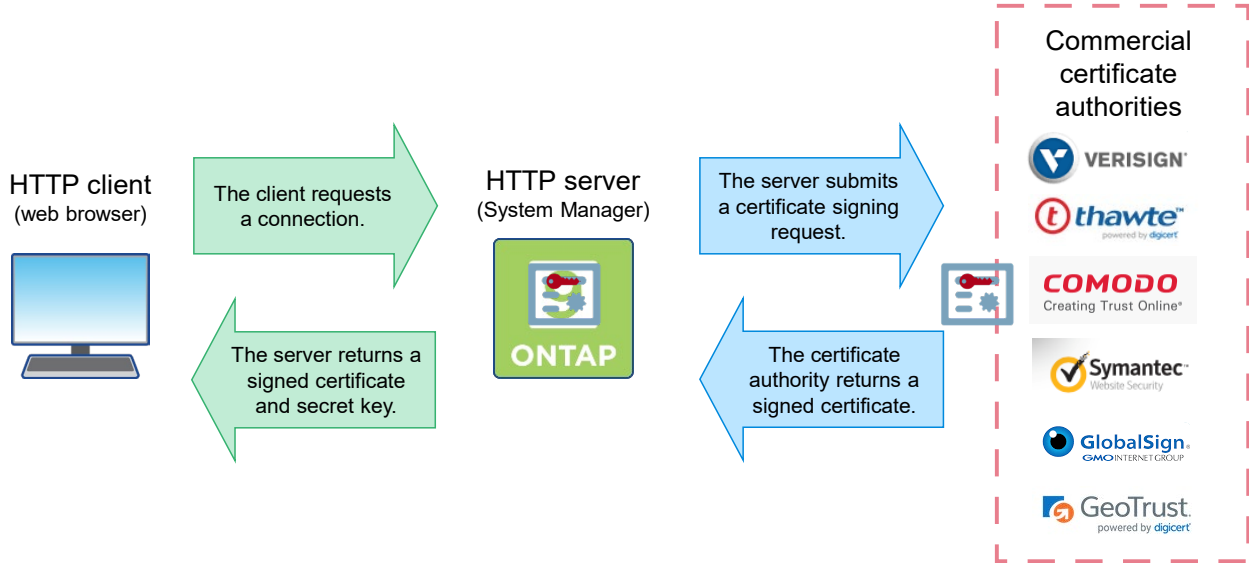


ONTAP based systems can use certificates to secure access to the cluster. During cluster initialization, ONTAP software creates a default security certificate self-signed by the cluster certificate authority.

Many client applications do not accept this self-signed certificate, or they might require you to provide the certificate authority (or CA) certificate.

In production environments, you should use a certificate that is signed by a trusted CA.

## Security certificates



NetApp 43 © 2023 NetApp, Inc. All rights reserved.

The Secure Shell (SSH) and HTTPS protocols use Transport Layer Security (TLS) to secure traffic by encrypting data before transmission and decrypting the data after reception. The encryption and decryption is performed using encryption keys that are known to both the client and server (shared secret key).

The secret encryption key is sent from the server to the client when the secure session is established.

How does the client know that it has contacted the authentic server and not an imposter? To validate its identity, the server includes a digital identity certificate. However, if the certificate is signed by the server itself, the certificate is no better confirmation of the server's identity than the original declaration. Instead, the certificate must be signed by a third party who the client trusts (friend of a friend).

The Public Key Infrastructure (PKI) defines how this process works.

The server should submit a certificate signing request to a trusted commercial certificate authority. The certificate authority verifies the submitter's identity and authority (at a minimum they will verify that the submitter is the true owner of the domain that the certificate covers) and return a signed certificate.

The client host accepts the identity of the server because it is being vouched for by a certificate authority that it trusts. A list of trusted root certificate authorities is included with all modern operating systems.

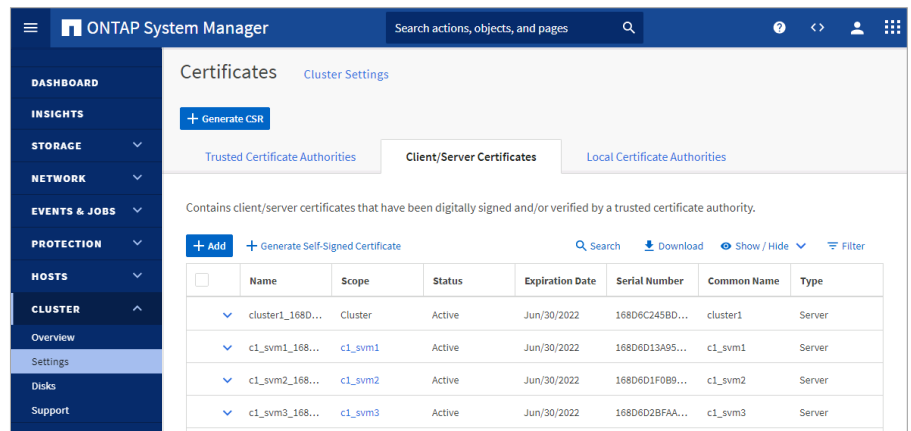
ONTAP based systems are both a TLS client and server. ONTAP software functions that act as a TLS client and require certificate validation are AutoSupport, event management system (EMS), Lightweight Directory Access Protocol (LDAP), audit logging, FabricPool technology, and Key Management Interoperability Protocol (KMIP).

ONTAP software includes a default set of certificates for trusted commercial certificate authorities.

## System Manager Certificates page

Use the ONTAP System Manager Certificates page to do the following:

- Generate a certificate signing request
- View and manage trusted certificate authorities
- View and manage client and server type certificates
- View and manage local certificate authorities that can sign certificates



Starting with ONTAP 9.10.1 software, you can use ONTAP System Manager to manage security certificates.

From the System Manager Certificates page, you can do the following:


- Generate a certificate signing request that can be sent to a commercial certificate authority
- View trusted certificate authorities and add or remove an authority from the ONTAP trust store
- View and manage client and server type certificates  
Certificates can be added or removed, and new self-signed certificates generated.
- View and manage local certificate authorities that can sign certificates

ONTAP software understands the following types of client/server certificates:

- server: Includes server certificates and intermediate certificates
- client-ca: Includes the public key certificate for the root CA of the TLS client
- server-ca: Includes the public key certificate for the root CA of the TLS server to which ONTAP software is a client
- client: Includes a self-signed or CA-signed digital certificate and private key to be used for ONTAP software as a TLS client

## FIPS 140-2 compliance mode

- When it is enabled, ONTAP software requires Federal Information Processing Standard (FIPS) 140-2 compliant encryption.
  - Only TLSv1.3 and TLSv1.2 are permitted.
  - Less secure protocols are rejected.
- When it is disabled, SSLv3 and TLSv1.1 can be added to the supported-protocols list.
- When FIPS-compliant mode is disabled, noncompliant protocols are not automatically re-enabled.
- Use the `security config modify` command to select the accepted security protocols and encryption ciphers.

 45 © 2023 NetApp, Inc. All rights reserved.




If you enable Federal Information Processing Standard (FIPS) compliance, the cluster automatically selects only compliant TLS protocols (currently TLSv1.3 and TLSv1.2). Attempts to use less secure protocols are rejected.

Use the `security config modify` command to change the cluster-wide security configuration. By default, FIPS mode is disabled, and ONTAP software supports the TLSv1.3, TLSv1.2 and TLSv1.1 protocols. For backward compatibility, ONTAP software supports adding SSLv3 and TLSv1 to the supported-protocols list when FIPS compliance is disabled. Use the `-supported-cipher-suites` parameter to control which TLS cipher suites are permitted by the system. Beginning in ONTAP 9.11.1, only the `ecdsa-sha2-nistp256` cipher is permitted in FIPS compliance mode. If you want administrator accounts to access SVMs with an SSH public key, you must ensure that the host key algorithm is supported before enabling SSL FIPS mode.



# Lesson 7

## Routing management

 46 © 2023 NetApp, Inc. All rights reserved.

# Host name resolution

NetApp 47 © 2023 NetApp, Inc. All rights reserved.

Two methods support host-name resolution: DNS and hosts tables.

- You configure DNS and the hosts table in the admin storage VM.
  - When you set up the cluster, you should configure DNS.
  - As nodes join the cluster, configurations are propagated to each node.
  - By default, the order of lookup is hosts table and then DNS.
- Cluster and storage VM administrators can configure DNS in a data storage VM.
- Each storage VM has its own DNS and hosts table configuration.



More info in Addendum

## Host-name resolution for the admin storage VM

Only cluster administrators can configure DNS and the hosts table for host-name lookup in the admin storage VM. All applications except CIFS discovery use the host-name configuration of the admin storage VM. You cannot use NIS configuration for the admin storage VM.

Host-name resolution for the admin storage VM is configured when the cluster is created.

- Hosts table configuration for the admin storage VM: You can use the `vserver services dns hosts` command to configure the hosts table that resides in the root volume of the admin storage VM.
- DNS configuration for the admin storage VM: If you want to configure DNS after you set up the cluster, use the `vserver services dns create` command.

## Host-name resolution for a data storage VM

A cluster or storage VM administrator can configure DNS for host-name lookup in a data SVM. DNS configuration is mandatory when CIFS is used for data access.

DNS services can also be configured on a storage VM for FlexVol volumes by using the Vserver Setup wizard. If you want to configure DNS later, you must use the `vserver services dns create` command.

## Managing the hosts table (cluster administrators only)

A cluster administrator can add, modify, delete, and view the host-name entries in the hosts table of the admin storage VM. A storage VM administrator can configure the host-name entries for only the assigned storage VM.

## Routing management

### Overview

You control the outbound traffic of LIFs by configuring route tables and static routes.

The following is true of route tables:

- Route tables are routes that are automatically created in a storage VM when a service or application is configured for the storage VM.
- Routes are configured for each storage VM, identifying the storage VM, subnet, and destination.
- Route tables are per storage VM, so routing changes to one storage VM do not pose a risk of corrupting another storage VM route table.
- The admin storage VM has its own route table.



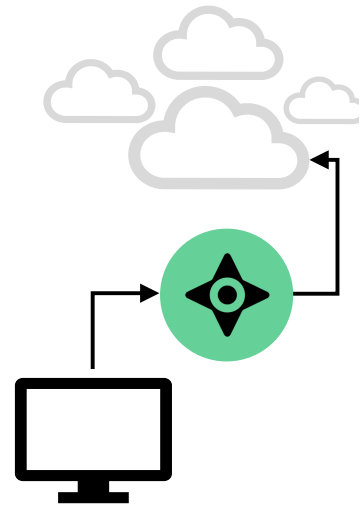
**NetApp** 48 © 2023 NetApp, Inc. All rights reserved.

The admin storage VM can own LIFs and might need route configurations that differ from the configurations on data storage VMs.

## Routing management

### Static routes

- A static route is a defined route between a LIF and a specific destination IP address.
- The route can use a gateway IP address.
- A static default route to the destination gateway is automatically added to the routing table of the storage VM if the following are true:
  1. A default gateway is defined when you create the subnet.
  2. A LIF from the subnet is assigned to a storage VM.



More info in Addendum


A static route is a defined route between a LIF and a specific destination IP address. The route specifies that to reach some destination host or network, it should direct the traffic through a specific interface or gateway router. The destination is typically a network address so that you can reach all the systems on that network.

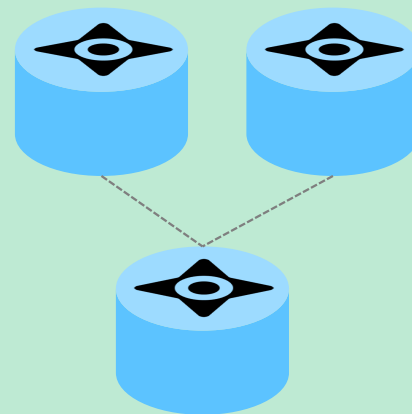
If a default gateway is defined when you create a subnet, the first time that a storage VM is assigned a LIF from the subnet, a default static route to the gateway is automatically added to the routing table of the storage VM.



## Border gateway protocol routing

- ONTAP 9.5 and later software supports Layer 3 routing through the border gateway protocol (BGP).
- Previous ONTAP versions used Layer 2 routing, which creates hash tables of routes based on “distance.” The fewest hops between two points is assumed to be the preferred route.
- Layer 3 routing with BGP uses metrics to pick routes based on measurements like latency and bandwidth availability.
- Support for BGP also enables the separation of LIFs from the physical hardware and makes them entities of the network that are called virtual IPs (VIPs).

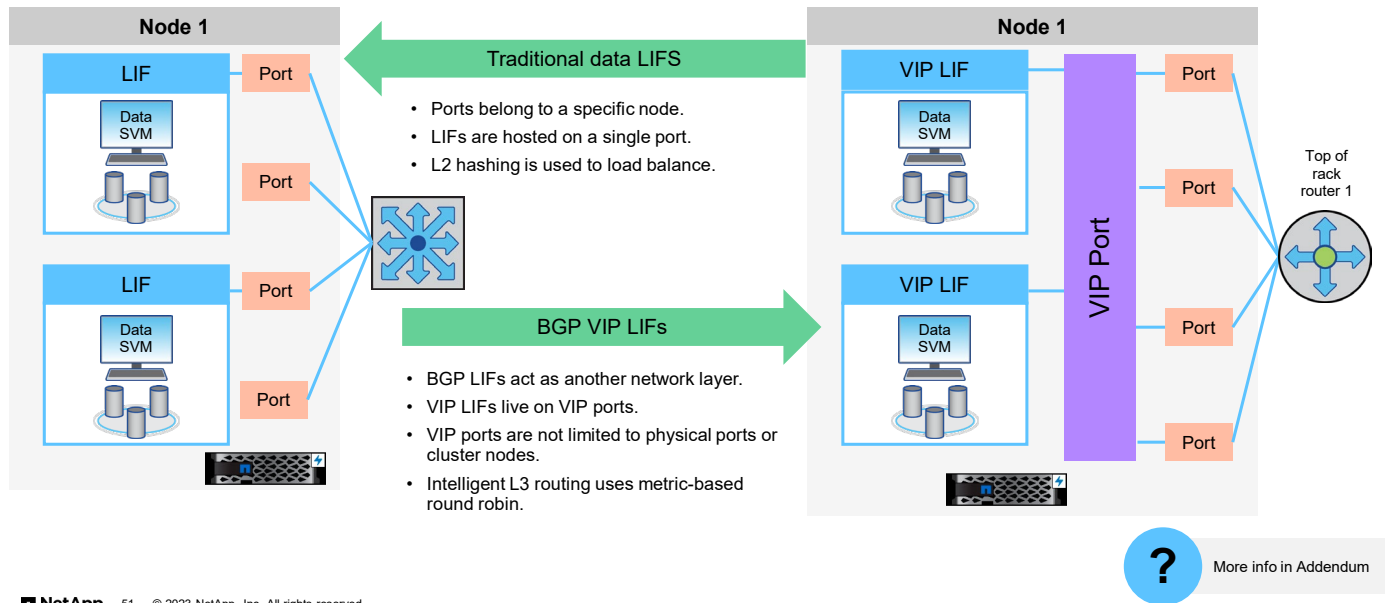
 50 © 2023 NetApp, Inc. All rights reserved.



Support for the border gateway protocol (BGP) brings Layer 3 (L3) routing to ONTAP software. L3 routing is considered more effective at finding the best route between two network locations. Previous versions of ONTAP software supported only Layer 2 (L2) routing, which assumes that fewer hops between locations means that the route is shorter and therefore more efficient. BGP instead relies on metrics to determine which route or routes are operating more efficiently. One of the ways that BGP works is by creating more flexibility in route choices by moving LIFs into the network that are called virtual IPs (VIPs). A standard LIF is tied to physical hardware, which means that its routing options are restricted to physical cabling connections. VIPs provide better redundancy for IP failover events and avoid inactive links.

An analogy for L2 and L3 routing would be the road system for vehicle traffic. In L2 routing, a trip from San Francisco to Los Angeles considers only the major highways that connect the two cities (for example, Interstate 5), regardless of traffic congestion or construction. L3 routing considers all roads and routes over state highways and surface streets, if the metrics indicate that such routes are faster in some areas.

## BGP and VIP LIFs



A LIF is typically assigned to a specific network port on a specific cluster node and is limited by the capabilities of that physical port. Additional network bandwidth can be achieved by assigning the LIF to a link aggregation group and spreading the network traffic across multiple ports on the same cluster node.

VIP data LIFs are like traditional LIFs but are accessible through any of the network ports in the broadcast domain of the storage VM. The network ports might be on different cluster nodes and be attached to different subnetworks. VIP data LIFs can support aggregate throughput that exceeds the bandwidth of any individual port, because the VIP LIFs can send or receive data from multiple subnets or ports simultaneously.

The availability of a VIP data LIF on a port is advertised to the peer router through the BGP LIF. One BGP LIF must be configured on each cluster node that is hosting a VIP data LIF.

Work with your network administrators to determine how to best implement BGP and VIPs into your environment.

## References


- NetApp Hardware Universe  
<http://hwu.netapp.com>
- ONTAP 9 Documentation Center  
<http://docs.netapp.com/ontap-9/index.jsp>

NetApp Hardware Universe: <http://hwu.netapp.com>

ONTAP 9 Documentation Center: <http://docs.netapp.com/ontap-9/index.jsp>

# Knowledge check

Module 4: Network management

 53 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

### Which statement about LIFs is true?

- a. One cluster-management LIF exists per node.
- b. One port can host multiple data LIFs.
- c. Cluster LIFs and data LIFs can share a port.
- d. A data LIF can be associated with multiple storage VMs.

## Knowledge check

### Which statement about LIF failover policies is true?

- a. Failover policies are assigned to failover groups.
- b. Failover policies can be created to control LIF failover behavior.
- c. Failover policies define how ports are selected during LIF failover.
- d. Failover policies define how LIFs are assigned to failover groups.

## Module summary

This module focused on enabling you to do the following:

- Describe the interaction between physical and virtual networking resources in a cluster
- Configure and manage physical and virtual networking resources



## Complete an exercise

Module 4  
Network management

### Managing Virtual Network Resources

- Access your lab equipment.
- Open your Exercise Guide, Module 4.
- Complete Exercise 2.
- Share your results.


This exercise requires approximately  
**35 minutes.**





# Addendum

## Failover group commands

 58 © 2023 NetApp, Inc. All rights reserved.

## Failover

Managing failover groups and LIFs

Create a failover group:

```
::> net int failover-groups create -vserver svm4 -failover-group fg_svm4  
-targets cluster1-01:e0f,cluster1-02:e0f
```

Add or remove targets from a failover group:

```
::> network interface failover-groups add-targets  
::> network interface failover-groups remove-targets
```

Configure failover for an existing LIF:


```
::> net int modify -vserver svm4 -lif svm4_nfs_lif1  
-failover-group fg_svm4 -failover-policy broadcast-wide-domain
```

Failover targets can be any combination of network ports, interface groups, or VLANs on a node.



# Addendum

## Routing management commands

 60 © 2023 NetApp, Inc. All rights reserved.

## Routing management

### Managing routes

Create a static route:


```
::> network route create -vserver svm4 -destination 0.0.0.0/0  
-gateway 192.168.0.1
```

Delete a static route:

```
::> network route delete -vserver svm4 -destination 0.0.0.0/0  
-gateway 192.168.1.1
```

Display static routes:

```
::> network route show  
Vserver  Destination  Gateway  Metric  
-----  -  
svm4     0.0.0.0/0    192.168.0.1  20 ...
```

 61 © 2023 NetApp, Inc. All rights reserved.

Use the `network route` command to manage the static routing table of the storage VM.

The first example creates a default route to all destinations through the gateway with IP address 192.168.0.1.

You can use the optional `-metric` parameter with the `network route create` command to specify a hop count for the route. The default settings for the parameter are 10 for management interfaces, 20 for data interfaces, and 30 for cluster interfaces. The parameter is used for source-IP address selection of user-space applications such as Network Time Protocol (NTP).

## Host name resolution


Configuring data storage VMs

Create a hosts table entry:

```
::> vserver services name-service dns hosts create -vserver svm4  
-address 192.168.0.11 -hostname test.example.com -alias test
```

Configure DNS:

```
::> vserver services name-service dns create -vserver svm4  
-domains example.com -name-servers 192.168.0.11
```

 62 © 2023 NetApp, Inc. All rights reserved.


You can use the `vserver services name-service dns create` command to assign a storage VM to a different DNS domain from the ONTAP cluster.

Each storage VM maintains a local hosts file with host name to IP address mappings. Use the `vserver services name-service dns hosts` command to manage storage VM host table entries.

By default, the storage VM looks up host names in the local hosts file first and then, only if the entry is not present, queries DNS. The `vserver services name-service ns-switch modify` command can be used to change the order or include additional name services, such as Lightweight Directory Access Protocol (LDAP) or Network Information Service (NIS).

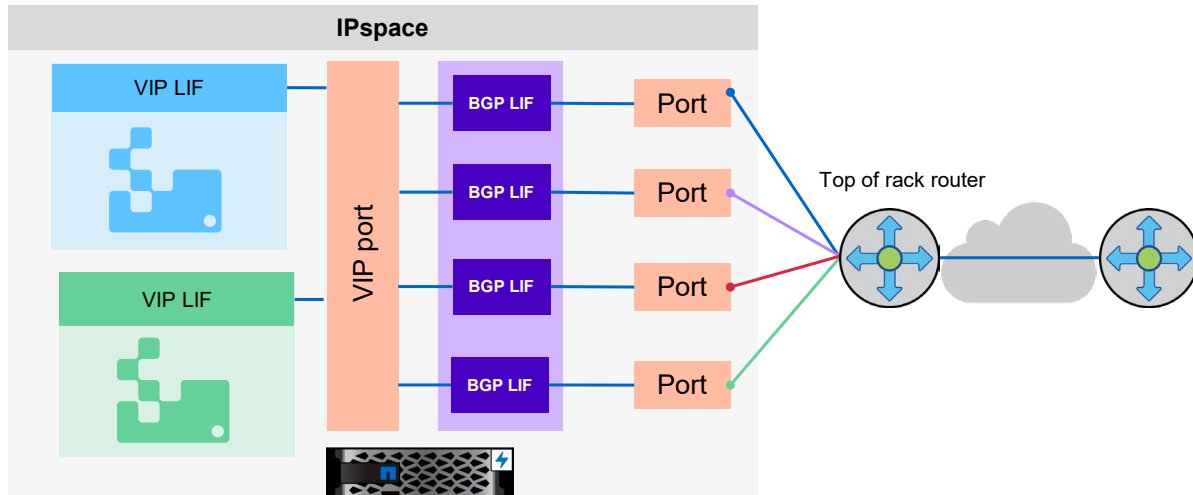


# Addendum BGP and VIP LIFs

 63 © 2023 NetApp, Inc. All rights reserved.

## BGP VIP

Full link use, direct cross-data-center traffic, and failure resiliency with a routed topology, including BGP support



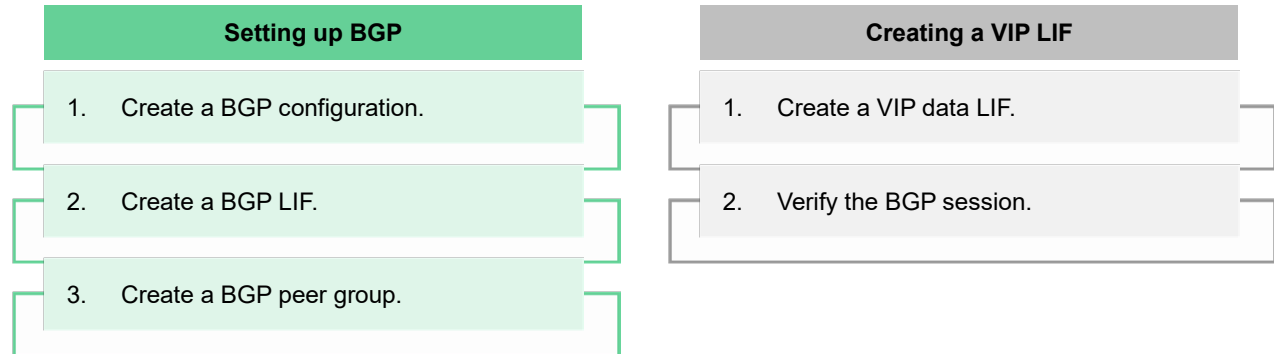
NetApp 64 © 2023 NetApp, Inc. All rights reserved.

The next-generation data centers rely on network Layer 3 and require LIFs to be failed over across subnets. Most of the Massively Scalable Data Center (MSDC) deployments use BGP or Open Shortest Path First (OSPF) as routing protocols to exchange the routes. BGP is most widely used by customers that require virtual IP (VIP) functionality. BGP is used by routers to exchange routing information so that the routers can dynamically update their routing tables with the best available routes to a destination. BGP is a connection-based protocol that runs over TCP. For BGP to work, there must be a connection setup between the two BGP endpoints (generally routers).

VIP enables users to create a data LIF that is not part of any subnet and is reachable from all physical ports of an IPspace on the local node. A VIP LIF is not hosted on any physical interface. It is hosted on a system-created pseudo interface (VIP port).

For ONTAP software, BGP is the routing protocol that is supported for advertising VIP.

## Configuring BGP and VIP LIFs



**NetApp** 65 © 2023 NetApp, Inc. All rights reserved.

Setting up BGP involves optionally creating a BGP configuration, creating a BGP LIF, and creating a BGP peer group. Before you begin, a peer router must be configured to accept BGP connection from the BGP LIF for the configured autonomous system number (ASN). ONTAP software automatically creates a default BGP configuration with default values when the first BGP peer group is created on a given node. A BGP LIF is used to establish BGP TCP sessions with peer routers. For a peer router, a BGP LIF is the next hop to reach a VIP LIF. Failover is disabled for the BGP LIF. A BGP peer group advertises the VIP routes for all the storage VMs in the peer group's IPspace.

When you create a VIP LIF, the VIP port is automatically selected, if you do not specify the home port with the `network interface create` command. By default, the VIP data LIF belongs to the system-created broadcast domain named "VIP" for each IPspace. You cannot modify the VIP broadcast domain.

A VIP data LIF is reachable simultaneously on all ports that host a BGP LIF of an IPspace. If there is no active BGP session for the VIP's storage VM on the local node, the VIP data LIF fails over to the next VIP port on the node that has a BGP session that is established for that storage VM.

For more information about BGP and VIP LIFs, see the *Network Management Guide*.



## Setting up BGP

### Example workflow

Create a BGP configuration (advanced command):

```
cluster1::*> network bgp config create -node node1 -asn 65502  
-holdtime 180 -routerid 1.1.1.1
```

Create a BGP LIF:

```
cluster1::> network interface create -vserver cluster1 -lif bgp1  
-service-policy net-route-announce -home-node cluster1-01 -home-port e0c  
-address 10.10.10.100 -netmask 255.255.255.0
```

Create a BGP peer group (advanced command):

```
cluster1::*> network bgp peer-group create -peer-group group1  
-ip-space Default -local-lif bgp1 -peer-address 10.10.10.1 -peer-asn 65502 -route-preference  
100
```

## Creating VIPs

### Example workflow

Enable multipath routing (advanced command):

```
::*> network options multipath-routing modify -is-enabled true
```

Create a VIP data LIF:

```
::> network interface create -vserver vs34 -lif vip1 -is-vip true  
-data-protocol cifs,nfs,fcache -service-policy default-data-files  
-home-node gw-node1 -address 3.3.3.3
```

Verify that the BGP session is in up status:


```
::> network bgp vserver-status show
```

Node	Vserver	bgp status
node1	vs1	up

For more information about syntax and usage, see the Network Management Guide > Configuring virtual IP (VIP) LIFs section and the Command manual pages for the `network bgp` commands.

# Module 5

## Physical storage management


 1 © 2023 NetApp, Inc. All rights reserved.

## About this module

This module focuses on enabling you to do the following:

- Recognize NetApp ONTAP storage architecture concepts
- Manage physical storage resources, including drives, RAID groups, and aggregates
- Create data aggregates
- Create Flash Pool aggregates
- Set up FabricPool aggregates

## ONTAP storage architecture




NetApp 3 © 2023 NetApp, Inc. All rights reserved.

The NetApp ONTAP software storage architecture uses a dynamic virtualization engine in which data volumes are dynamically mapped to physical space.


In ONTAP software, disks are grouped into RAID groups. An aggregate is a collection of physical disk space that contains one or more RAID groups. Each aggregate has a RAID configuration and a set of assigned disks. The disks, RAID groups, and aggregates make up the physical storage layer.

Within each aggregate, you can create one or more FlexVol volumes. A FlexVol volume is an allocation of disk space that is a portion of the available space in the aggregate. A FlexVol volume can contain files or LUNs. The FlexVol volumes, files, and LUNs make up the logical storage layer.



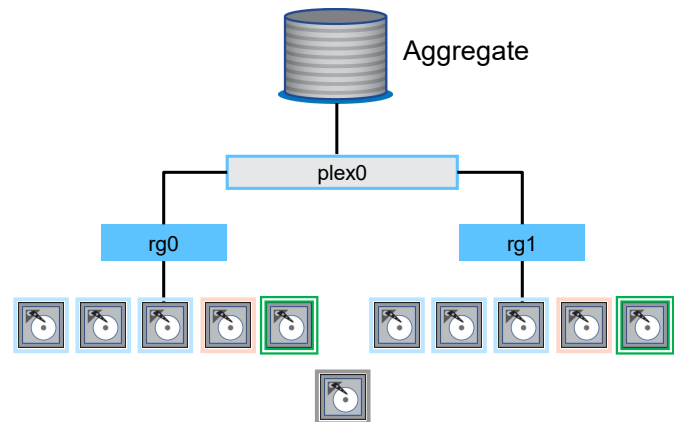
# Lesson 1

## Drives, RAID, and aggregates

 4 © 2023 NetApp, Inc. All rights reserved.

## Physical storage hierarchy

- **Drive:** HDD or SSD
- **RAID group:** Drive-level protection
- **Plex:** Container for RAID groups  
Used by mirrored aggregates
- **Aggregate:** Pool of storage space








NetApp 5 © 2023 NetApp, Inc. All rights reserved.

The ONTAP storage architecture hierarchy contains the following elements:

- **Drives:** Drives play different roles at different times, depending on the state of the drive. Potential drives states include the following:
  - Data
  - Parity
  - Double-parity
  - Triple-parity
  - Spare
  - Broken
  - Unowned
  - Uninitialized (not zeroed)
- **RAID groups:** Each RAID group contains physical disks and is associated with a plex. A RAID group has either a RAID 4, NetApp RAID-DP, or NetApp RAID-TEC configuration.
- **Plexes:** Each plex is associated with an aggregate and contains RAID groups. Typically, an aggregate has only one plex. Aggregates that use SyncMirror technology have two plexes (plex0 and plex1). Plex1 contains a mirror of the plex0 data.
- **Aggregates:** Each aggregate contains a plex or plexes, a RAID configuration, and a set of assigned physical disks to provide storage to the volumes that the aggregate contains.

## Drive types

 <b>NL-SAS HDD</b>	 <b>SAS HDD</b>	 <b>SAS SSD</b>	 <b>NVMe SSD</b>
<p>Near-line serial-attached SCSI hard disk drive</p> <ul style="list-style-type: none"><li>• Same technology that is used in consumer disk drives</li><li>• Dual I/O path</li><li>• High capacity but moderate IOPS</li></ul>	<p>Serial-attached SCSI hard disk drive</p> <ul style="list-style-type: none"><li>• Point-to-point serial protocol</li><li>• Multipath I/O</li><li>• Moderate capacity but high IOPS</li></ul>	<p>Serial-attached SCSI solid-state drive</p> <ul style="list-style-type: none"><li>• No spinning platter</li><li>• Based on flash memory chip technology that is like USB flash drives</li><li>• High IOPS with low latency</li><li>• Can also be used as an aggregate-specific cache</li></ul>	<p>Nonvolatile memory express solid-state drive</p> <ul style="list-style-type: none"><li>• Solid-state flash memory drives that are accessed by using the NVMe protocol</li><li>• Extreme IOPS for demanding workloads</li></ul>

 6 © 2023 NetApp, Inc. All rights reserved.

AFF systems use only SSD drives. FAS systems can use a mix of drive types.

SATA is the disk technology that is used in most consumer-grade PCs. These drives have high capacities but moderate IOPS. NL-SAS drives include SATA-to-SAS adapters, which enable the use of these drives in SAS shelves.

SAS is a point-to-point serial protocol that replaced parallel SCSI to resolve contention issues from multiple devices that share a system bus. SAS disks can use multiple I/O paths. SAS drives typically spin faster than SATA drives and provide higher performance.

SSDs are fast and reliable and use long-lasting technology that is based on the same flash technology that is used for USB flash drives. SSDs can be configured as data storage or as aggregate-specific cache.

NVMe SSDs are a new class of drive. These drives use the same flash memory drives but communicate by using the NVMe protocol over Ethernet. This arrangement enables NVMe drives to operate faster than traditional SSDs.




## Drive ownership

- A drive is unusable until the drive is assigned ownership by a storage controller.
  - By default, ownership is automatically assigned.
  - Ownership can be manually assigned or changed.
  - Software disk ownership is made persistent by writing the ownership information onto the drive.
- Spare drives can be reassigned or unassigned.

```
::> storage disk show -container-type unassigned
```

Disk	Usable Size	Shelf	Bay	Container Type	Position	Aggregate	Owner
9.11.18	-	11	18	unassigned	present	-	-

 7 © 2023 NetApp, Inc. All rights reserved.

ONTAP software automatically assigns drives to a storage controller during the initial hardware setup and checks occasionally to determine whether new drives have been added. When the drive is assigned, the disk ownership information is written to the drive so that the assignment remains persistent.

Ownership can be modified or removed. The data contents of a drive are not destroyed when the drive is marked as unowned. Only the disk ownership information is erased.

Automatic ownership assignment is enabled by default. If your system is not configured to assign ownership automatically, or if your system contains array LUNs, you must assign ownership manually.

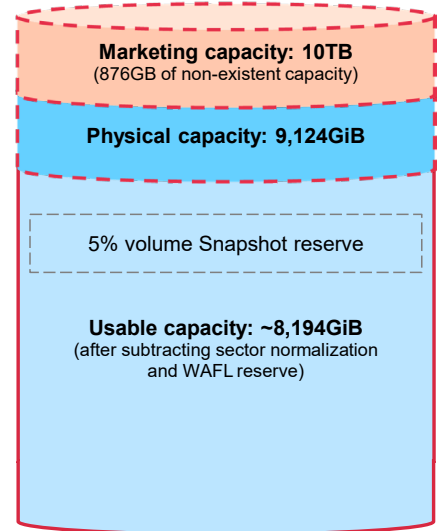
**Note:** The NetApp best practice is to unassign only spare drives.

## Drive capacity

### Marketing and physical capacity

Drive capacity is a confusing and contentious subject for many reasons. Consider a 10TB disk drive:

- **Marketing capacity:** Base-10 number rounded up to an even number
- **Physical or raw capacity:** Available space after sector formatting in base-2 numbering  
After formatting, a 10TB drive is really a 9124GiB drive.
- **Usable capacity:** Space available to NetApp WAFL in base-2 numbering
  - **Sector normalization:** NetApp ONTAP software “rightsizes” all “10TB” drives equally, so they have the same number of sectors. This arrangement might result in 9105GiB.
  - **NetApp WAFL reserve:** This reserve is a percentage of the capacity that is set aside for WAFL metadata. 5% is reserved in AFF aggregates larger than 30TB. Otherwise, 10% is reserved.
  - The space available for data is now ~ 8,194GiB, which appears to an end user as if 2TB has vanished.
- **Snapshot reserve:** FlexVol volumes reserve 5% for storing Snapshot copies, which the customer might perceive as more lost space.



NetApp 8 © 2023 NetApp, Inc. All rights reserved.

Drive capacity is often confusing and contentious. Even on a new system with no data on it, the total capacity reported by the system is significantly smaller than the total of the capacity numbers that are physically shown on the drive carriers.

The root of this issue is how drives are marketed. When drives had very small capacities, it was easier to sell a drive if it was marketed as 100MB rather than 86MB. Vendors and resellers calculated the marketing capacity by using base-10 numbering rather than the base-2 numbering system that is used by computers. Unfortunately, this marketing practice still occurs today. The differences in capacities can be hundreds of gigabytes.

Physical or raw capacity is the actual base-2 computed capacity that the drive has when it leaves the factory.

Usable capacity is the disk blocks that are available to store data after the differences in calculation and overhead are considered. Because not all manufacturers create drives of the same capacity, NetApp normalizes all disks to the size of the smallest available disk capacity. WAFL then reserves the top 10% of capacity for its use. Beginning with ONTAP 9.12.1 software on AFF systems only, if the total aggregate size is larger than 30TB, only 5% of the aggregate space is reserved for use by WAFL.


Now that you know the difference between marketing capacity and usable capacity, you need to define what usable means. NetApp considers all blocks in the active file system and in Snapshot reserves as usable space. The reason is that Snapshot copies hold older copies of data blocks for the purposes of recovering or restoring older versions of files. To offset this perception, ONTAP software supports deduplication and compression to enable customers to pack more data into fewer disk blocks.

NetApp ONTAP System Manager calculates storage capacity based on binary units of 1024 ( $2^{10}$ ) bytes. In ONTAP 9.10.0 and earlier software, these units were labeled in System Manager as KB, MB, GB, TB, and PB. Beginning with ONTAP 9.10.1, they are more accurately labeled KiB, MiB, GiB, TiB, and PiB.

## Drive roles

### Parity drive

**Parity drive:** Stores row parity information that is used for data reconstruction when a single drive fails within the RAID group



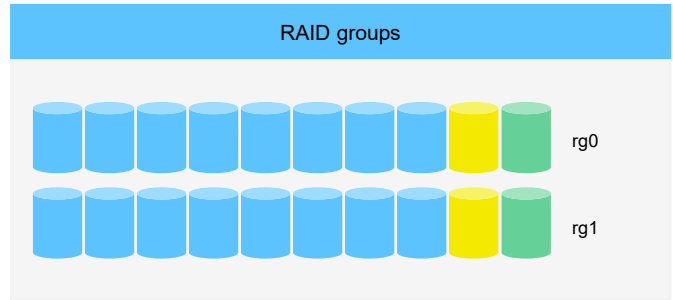
The key component of RAID group functionality is the parity drives. The parity drive stores the results of the parity calculation across all data drives in the RAID group. A parity drive can protect against the loss of a single drive within a RAID group.

If you add a spare drive to an aggregate and the spare is larger than the other data drives, the spare becomes a parity drive. However, the spare does not use the excess capacity unless another drive of similar size is added. The second largest additional drive has full use of additional capacity.

## Drive roles

Double parity drive

**dParity drive:** Stores diagonal parity information that is used for data reconstruction when two drives fail within the RAID group



NetApp 10 © 2023 NetApp, Inc. All rights reserved.



The dParity drive stores diagonal parity values in a RAID DP group. This capability provides protection against two drive failures in a RAID group. Most drive failures are not due to mechanical reasons but are failures in the drive medium. These failures are called soft failures. When drive capacities grew to 1TB, the industry saw an increase in soft failures. More blocks mean more probability of a soft failure. As capacities increased, so did the rebuild times. During the hours it takes to rebuild one failed drive, another drive in the same RAID group could experience a soft failure. This failure could cause the aggregate to go offline to protect against further failures. This situation could result in data loss. This condition is known as a double-disk failure.

It is important to know that a NetApp storage system can safely experience multiple drive failures and remain operational. The distinction is that the failures must occur in the same RAID group to qualify as a double-disk failure.

## Drive roles

Triple parity drive

**tParity drive:** Stores anti-diagonal parity information that is used for data reconstruction when three drives fail within the RAID group



**NetApp** 11 © 2023 NetApp, Inc. All rights reserved.


The tParity drive is the third parity drive that is used by RAID-TEC groups. The tParity drive protects against a third drive failure. RAID-TEC is required when you use drives with capacities larger than 6TB because of the increased probability of soft failures.

## Drive roles

### Spare drive

#### Spare drive:

- Assigned to a storage system but not in use by a RAID group
- Used to create aggregates, add capacity to aggregates, and replace failing drives  
Spare drives must be "zeroed" before use.
- If there are insufficient spare drives available, ONTAP will automatically halt the node after 24 hours.



NetApp 12 © 2023 NetApp, Inc. All rights reserved.

Not all drives are used to store data. To replace failed drives as quickly as possible, storage systems require that a small percentage of drives is set aside as spares. Storage administrators can also use them to grow an aggregate by adding them to a RAID group.


Before a spare drive can be used, all the data blocks must be set to a value of zero. This process is referred to as "zeroing." Newly purchased drives and replacement drives that are sent by the NetApp Support team are already zeroed. If a drive is removed from a RAID group for any reason, it must be zeroed in NetApp ONTAP System Manager (formerly OnCommand System Manager) or the CLI before it is added to the spares pool. Verify that all spare and unused drives are zeroed regularly. An unused drive that is not zeroed is not counted as a spare and is not used to replace a failed drive.

If there are an inadequate number of spare drives available, an ONTAP protection feature will trigger an automatic shutdown in 24 hours.

## ONTAP RAID technologies

### Description

- RAID 4 (row parity)
  - Adds a *row parity* drive
  - Protects against single-disk failure or media error
- RAID DP (double parity) technology
  - Adds a *diagonal parity* disk to a RAID 4 group
  - Protects against two concurrent drive failures within a RAID group
- RAID-TEC (triple erasure coding) technology
  - Adds a *triple-parity* disk to a RAID DP group
  - Protects against three concurrent drive failures



NetApp 13 © 2023 NetApp, Inc. All rights reserved.

### RAID 4

In a RAID 4 group, parity is calculated separately for each row. In the example, the RAID 4 group contains seven disks, with each row containing six data blocks and one parity block.

### RAID DP technology

In a RAID DP group, a diagonal parity set is created in addition to the row parity. Therefore, an extra double-parity drive must be added. In the example, the RAID DP group contains eight drives, with the double parity calculated diagonally by using seven parity blocks.

- The number in each block indicates the diagonal parity set to which the block belongs.
- Each row parity block contains even parity of data blocks in the row, not including the diagonal parity block.
- Each diagonal parity block contains even parity of data and row parity blocks in the same diagonal.

### RAID-TEC technology

In a RAID-TEC group, an anti-diagonal parity set is created in addition to both the row parity and diagonal parity sets. Therefore, an extra third-parity drive must be added. In the example, the RAID-TEC group contains nine drives, with the triple parity calculated anti-diagonally by using seven parity blocks.

- Seven diagonals (parity blocks) exist, but ONTAP software stores six diagonals (p-1).
- The missed diagonal selection is arbitrary.

## RAID group sizes

Default RAID group sizes:

- 21 drives for SATA or NL-SAS drives
- 24 drives for SSD or SAS HDD

Do not wait for an aggregate to fill completely before before expanding the aggregate.

Grow existing RAID groups before adding new RAID groups to the aggregate.

Disk type	Group type	Default	Maximum
NL-SAS	RAID4	7	7
	RAID DP	14	20
	<b>RAID-TEC</b>	<b>21</b>	<b>29</b>
SAS	RAID4	8	14
	RAID DP	16	28
	<b>RAID-TEC</b>	<b>24</b>	<b>29</b>
SSD	RAID4	8	14
	RAID DP	23	28
	<b>RAID-TEC</b>	<b>24</b>	<b>29</b>

To create a RAID-TEC aggregate, you need a minimum of seven drives.


Ideally, you want to use fully populated RAID groups to create aggregates. At a minimum, RAID groups should be half their maximum size. Then when you grow the aggregates, you can add drives to make the RAID groups fully populated.

When drives are added to a RAID group, WAFL directs writes to the new drives and the existing drives evenly. ONTAP automatically rebalances the data across all the drives in the RAID group. This rebalance is performed in the background at a lower priority than user initiated I/O. If the existing drives in the RAID group fill before the rebalance finishes, only the new emptier drives have available space. They are used to store new data, which might decrease the performance of the aggregate until the rebalance completes distributing the free space across all of the drives in the RAID group.



## RAID group recommendations

- Drives must be the same type:
  - SAS, NL-SAS, or SSD
- Drives should be the same size.
- HDD should be the same rotational speed:
  - SAS 10K RPM
  - NL-SAS 7.2K RPM
- You should provide sufficient hot spares.



**NetApp** 15 © 2023 NetApp, Inc. All rights reserved.

A RAID group consists of one or more data drives or array LUNs across which client data is striped and stored.

You change the size of RAID groups on a per-aggregate basis. You cannot change the size of an individual RAID group.

When sizing RAID groups of HDDs or SSDs, observe the following guidelines:

- RAID groups are composed of the same disk type.
- All RAID groups in an aggregate should have the same number of drives.

If you cannot follow the guideline, any RAID group with fewer drives should have only one drive fewer than the largest RAID group.

**Note:** The SSD RAID group size can differ from the RAID group size for the HDD RAID groups in a Flash Pool aggregate. Usually, you should verify that you have only one SSD RAID group for a Flash Pool aggregate, to minimize the number of SSDs that are required for parity.

- The recommended range of RAID group sizes is as follows:
  - Between 12 and 20 for SATA HDDs
  - Between 20 and 28 for SAS HDDs and SSDs

The reliability and smaller size (faster rebuild times) of performance HDDs can support a RAID group size of up to 28, if needed.


- Avoid mixing drives of differing capabilities and capacities in the same aggregate. The performance of the RAID group, and therefore the entire aggregate, is limited by the speed of the slowest drive. Mixing drives of different capacities in the same RAID group results in drive “rightsizing,” in which the extra capacity of the larger drives goes unused.

Recommendations about spares vary by configuration and situation. For information about best practices for working with spares, see Technical Report 3437: *Storage Subsystem Resiliency Guide*.

## Aggregates

- Aggregates are logical containers for the drives that are managed by a node.
- Aggregates consist of one or more RAID groups.
- You can use aggregates to do the following:
  - Isolate workloads with different performance demands
  - Tier data with different access patterns
  - Segregate data for regulatory purposes
- A single node owns an aggregate, but ownership can be transferred to the partner in a high-availability (HA) pair.

During an HA failover, aggregate ownership is temporarily transferred to the surviving partner.



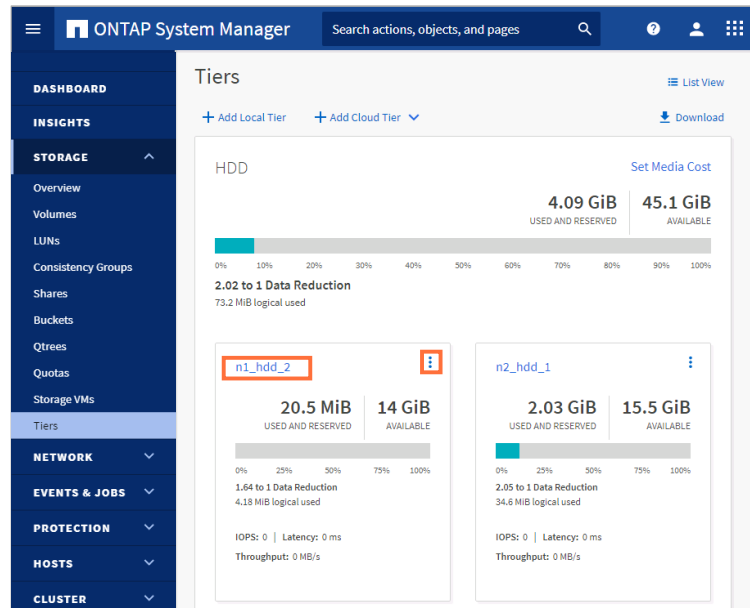
Recall that ONTAP aggregates are storage pools containing all the usable space of all the data drives in every RAID group owned by the aggregate.

You can use aggregates to store FlexVol volumes and FlexGroup volumes containing your user data. You can assign volumes to aggregates based on your performance needs and security requirements.

An aggregate is owned by a single cluster node. However, it can be given to, or taken by, its HA partner node while performing non-disruptive maintenance operations or to recover from a node failure.

## Viewing aggregates in ONTAP System Manager

- Navigate to the NetApp ONTAP System Manager Tiers page:  
**Storage > Tiers**
- ONTAP System Manager refers to aggregates as "local tiers."
- Local tiers are grouped by drive type (HDD, SSD, and Flash Pool).



NetApp 17 © 2023 NetApp, Inc. All rights reserved.

Beginning with ONTAP 9.7, NetApp ONTAP System Manager uses the term "local tier" to refer to a physically attached aggregate. It uses "cloud tier" to refer to a FabricPool aggregate that is tiered to an S3 bucket in an object store.

Click the local tier name link in the Storage Tiers page to view detailed information about the aggregate or use the More menu ("...") to perform administration actions on the aggregate.

## Create an aggregate

- System Manager creates best-practice conformant aggregates.
  - Generally, one aggregate per node per drive type is created.
  - Click **Recommendation details** to view the planned aggregates.
  - Click **Save** to create the aggregates.
- **Switch to Manual Local Tier Creation** or use the CLI or API to create nonconformant aggregates.

```
::> aggregate create -aggregate n2_data_02  
-node cluster1-02 -disktype ssd -diskcount 8
```

The screenshot shows the ONTAP System Manager interface with the 'Add Local Tier' dialog box open. The dialog displays a storage recommendation of 113 GiB, which is usable. It indicates that 2 local tiers can be added on nodes cluster1-01 and cluster1-02. The 'Recommendation details' section shows a table for Local Tier Details:

Local Tier	Node Name	Usable Size	Type	Disks	Layout
cluster1_01_FC_1	cluster1-01	56.3 GiB	HDD	10 X 3.93 GiB ( RAID-DP) 10 X 3.93 GiB ( RAID-DP)	rg1

Below this, the 'Spare Disks' section shows a table:

Node Name	Spare Disks	Type	is Partition
cluster1-01	2 X 3.93 GiB	FCAL	No
cluster1-02	2 X 3.93 GiB	FCAL	No
cluster1-01	20 X 521 MiB	SSD	No

A red box highlights the 'Switch to Manual Local Tier Creation' link at the bottom of the dialog. The 'Save' button is also visible.

The new System Manager simplifies administration by creating aggregates that conform with NetApp best practices. To create nonconformant aggregates, you must switch to manual local tier creation mode or use the CLI.

For most disk types, RAID DP is the default.

RAID-TEC is the only available RAID type if the following are true:

- The drive type of the aggregate drives is FSAS or mSATA.
- The drive size is equal to or larger than 10TB.



## Try this task

**NetApp** 19 © 2023 NetApp, Inc. All rights reserved.

- Use cluster1 in your lab equipment to try the following tasks:
- Open a PuTTY session and use the `aggr show` command.
  - Can you tell which node owns the aggregate?
  - What is the RAID status?
  - How do you determine how many disks are in each aggregate?
- Different commands show similar things in different ways:
  - Enter `aggr show -aggregate aggr0_n1`.
  - Enter `storage disk show -aggr aggr0_n1`.
  - How do the outputs of the commands differ?

1a. The owning node is listed in the Nodes column.

1b. RAID status should be `raid_dp`, normal.

1c. Use the `-instance` switch and review the “number of disks” field or use the `aggr show -fields diskcount` command.


2. The `aggr show` command displays extensive information about the aggregate, including the list of disks. The `storage disk show` command displays a list of disks in the aggregate and information about the disks.

## Adding drives to an aggregate

To add capacity to an aggregate, you add more drives. Careful planning ensures that you use the fewest drives to add the maximum amount of capacity.

**Example:** An aggregate composed of 4TB drives

- Three drives add 12TB of capacity and fill out the RAID group.
- To add 16TB, you need six drives.
  - You have no more spares.
  - Automatic shutdown after 24 hours is triggered
  - The "runt" RAID group provides poor capacity use.



NetApp 20 © 2023 NetApp, Inc. All rights reserved.

You can add drives from the spares pool to an aggregate to increase the aggregate's capacity. When you add drives, consider the size of RAID groups in the aggregate. Plan to fill complete RAID groups to maximize the amount of usable space that is gained in comparison to the number of drives that are used for parity. In the second example, six drives are added to the aggregate. However, only four of the six drives add capacity to the aggregate, because two drives are used for parity drives in a new RAID group.

By using all available spares, you have triggered an ONTAP protection feature that triggers a shutdown in 24 hours unless enough spares are assigned to the storage controller.

When you add drives, also consider the following:

- Addition of drives that the same system owns
- Benefits of keeping your RAID groups homogeneous for drive size and speed
- Types of drives that can be used together
- Checksum rules when drives of more than one checksum type are in use
- Addition of the correct drives to the aggregate (the disk addition operation cannot be undone)
- Method of adding drives to aggregates from heterogeneous storage
- Minimum number of drives that you must add for best performance
- Number of hot spares to provide for protection against drive failures
- Requirements for adding drives from multidisk carrier drive shelves

## Adding capacity to aggregates

Provide the following information:

- Aggregate name
- Disks to add:
  - Disk count
  - Disk type
  - Disk class
  - Disk size
  - Disk list

You cannot shrink aggregates.

```
::> storage disk show -spare -owner cluster1-01
::> storage aggregate add-disks -aggr n1_data_001 -diskclass capacity -diskcount 3
```

NetApp 21 © 2023 NetApp, Inc. All rights reserved.

Tier	Used and Reserved	Available
n1_hdd_2	20.5 GiB 1.64 to 1 Data Reduc 4.18 MiB logical used	15.5 GiB 2.05 to 1 Data Reduc 34.6 MiB logical used

Follow the recommendations that you just learned when adding capacity to an aggregate by using the CLI.


Drives can be selected by size, type, or class, or individually by name.

ONTAP System Manager follows the best practices and needs to know only how many drives to add to the aggregate.



## Lesson 2


# Advanced Disk Partitioning

 22 © 2023 NetApp, Inc. All rights reserved.



## Why slice drives into partitions?

Before ONTAP 8.3 software, the following is how entry-level HA pairs used their drives.




- Of the 24 drives in the chassis, each node can use only six drives to store data:
  - Four drives are used for parity.
  - One drive is reserved as a spare.
  - Six drives are used to store the node root aggregates (one RAID 4 aggregate per node).
  - Six drives are available to store data.
- Efficiency was limited to about 50%.

**NetApp** 23 © 2023 NetApp, Inc. All rights reserved.

Before the introduction of Advanced Disk Partitioning, entry-level systems with internal drives had to split ownership of the drives. Each controller in the HA system requires a root aggregate that consumes three drives to hold a root volume that is generally only 150GB in size. The data aggregate needs two drives for parity, and the system requires at least one spare drive. Only half of the drives in the chassis were available to store data.

Some customers would only assign four drives to node 2, making it an active-standby. Node 1 gained eight more drives but had to do all the work.

## Root-data Advanced Disk Partitioning



- Drives are partitioned into one small root partition and one large data partition.
- Standard aggregate configuration per node is as follows:
  - A root aggregate RAID group of 8 small data partitions + 2 small parity partitions and 2 small spare root partitions
  - A data aggregate RAID group of 9 large data partitions + 2 large parity partitions and 1 large spare data partition
- Total usable capacity is 18 data partitions out of a total of 24 drives, which achieves nearly 75% efficiency.


Advanced Disk Partitioning splits each drive into two sections. The smaller section is used to store the node root aggregates, and the larger section is used to store the node data aggregate.

Advanced Disk Partitioning spreads the smaller node root aggregates across more drives. It does not require dedicated parity drives for the node root aggregates, which both improve storage capacity.

Root-data partitioning is configured at the factory for all FAS HDD systems.

# Root-data-data Advanced Disk Partitioning

ONTAP 9 and later software



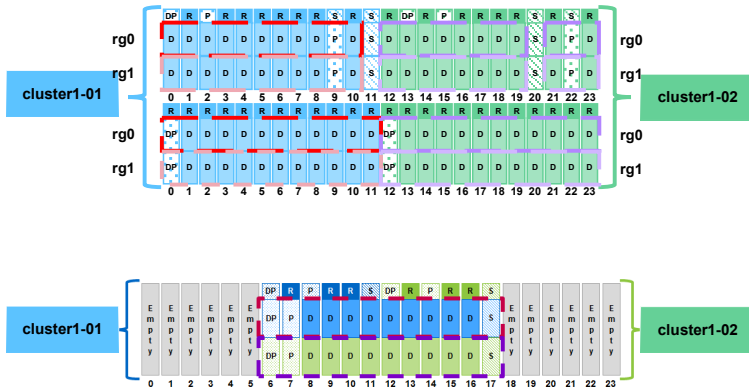
- SSDs are partitioned into one small root and two half-sized data partitions.
- The standard aggregate configuration per node is as follows:
  - A root aggregate RAID group of 8 small data partitions + 2 small parity partitions and 2 small spare root partitions (no change from root-data partition)
  - A data aggregate RAID group of 21 data partitions + 2 parity partitions and 1 spare data partition
- Three fewer drives are used to store parity compared to root-data partitioning.
- Data aggregates are spread across more drives for better throughput.

Root-data-data partitioning goes further by splitting the data partition into two and assigning each half to a different node. Root-data-data partitioning improves performance by spreading the data aggregates across more drives and improves storage capacity by reducing the number of parity and spare drives.

Root-data-data partitioning is configured at the factory for all AFF and SSD-only FAS systems.

# Root-data-data Advanced Disk Partitioning

Additional root-data-data partitioning information



- Root-data-data partitioning is supported on only SSD systems:
  - Default root aggregate provisioning method for AFF systems and SSD-only FAS systems
  - Unsupported on entry-level FAS or AFF MetroCluster FC software
- Data partition assignments with two shelves are like root-data partitioning:
  - Data partitions on an SSD are assigned to the same node.
  - Twice as many RAID groups are used.
- Half-shelf AFF systems have 50% more usable capacity than with root-data partitioning.

The figures show the default configuration for two-shelf and half-shelf AFF systems in ONTAP 9 software.

For root-data partitioning and root-data-data partitioning, RAID uses the partitions in the same way as physical drives. If a partitioned drive is moved to another node or is used in another aggregate, the partitioning persists. You can use the drive only in RAID groups that are composed of partitioned disks. If you add an unpartitioned drive to a RAID group that consists of partitioned drives, the unpartitioned drive is partitioned to match the partition size of the drives in the RAID group. The rest of the drive is unused.



## Complete exercises


Module 5  
Physical storage management

### Managing Physical Storage

#### Exploring RAID-TEC

- Access your lab equipment.
- Open your Exercise Guide, Module 5.
- Complete Exercises 1 and 2.
- Share your results.

This exercise requires approximately  
**20 minutes.**


 27 © 2023 NetApp, Inc. All rights reserved.

Open the exercise guide and complete Module 5 Exercise 1 and Exercise 2.




## **Lesson 3**

# **Flash Cache and Flash Pool features**

 28 © 2023 NetApp, Inc. All rights reserved.

## Accelerate I/O performance

- NetApp ONTAP data caching:
  - Is intelligent policy-based caching of data and metadata
  - Delivers high-speed data access for all protocols
  - Maintains deduplicated blocks in the cache
- Flash Cache feature
  - Is controller-based cache that is shared by all volumes on a node
  - Is the best choice for multiple heterogeneous workloads
  - Is simple to deploy and use
- Flash Pool feature
  - Is aggregate-level cache
  - Delivers high-speed data access for specific workloads
  - Provides cached data persistence through failovers
  - Is optimal for database and transactional applications



NetApp 29 © 2023 NetApp, Inc. All rights reserved.

ONTAP software lets you deploy flash as high-performance cache for your working data set while using lower-cost HDDs for less frequently accessed data.

At the storage level, there are two ways to cache data:

- The controller-based Flash Cache feature accelerates read and write operations and generally provides better performance for file services workloads. Flash Cache intelligent caching combines software and hardware within NetApp storage controllers to increase system performance without increasing the drive count. The Flash Cache controller-based solution is available to all volumes that are hosted on the controller and is the best choice for multiple heterogeneous workloads. A frequently seen use case for Flash Cache is to manage VMware boot storms. Some NetApp systems include Flash Cache (also known as external cache) cards by default.
- The Flash Pool feature is implemented at the aggregate level, enabling SSDs and traditional HDDs to be combined in a single ONTAP aggregate. Flash Pool technology provides read caching and write caching and is well suited for OLTP workloads, which typically have a higher percentage of write operations.

Both ONTAP features improve overall storage performance and efficiency and are simple to deploy and operate.

## Create a Flash Pool aggregate

Convert a traditional HDD aggregate into a Flash Pool aggregate.

Provide the following information:

- Existing aggregate name
- Cache source or drive type
- Number of drives
- RAID type

The screenshot shows the NetApp storage management interface for an aggregate named 'n2\_hdd\_1'. The 'Overview' tab is selected, displaying the following information:


STATUS	Online	FABRICPOOL	Disabled
NODE	cluster1-02	FLASH POOL	Disabled
TYPE	HDD (7 disks)	SNAPLOCK	Disabled
RAID	RAID-DP	MIRROR	Disabled

Capacity information is shown as follows:

USED AND RESERVED	2.03 GiB	AVAILABLE	15.5 GiB	COMMITTED	2.02 GiB
-------------------	----------	-----------	----------	-----------	----------

A bar chart shows the capacity usage, with 1.65 to 1 Data Reduction and 27.3 MIB logical used. A 'More' menu is visible on the right, with 'Add Flash Pool Cache' highlighted.

```
::> aggr modify -aggregate cluster2_01_FC_1 -hybrid-enabled true  
::> aggr add-disks -aggr cluster2_01_FC_1 -disktype SSD -diskcount 8
```



In a Flash Pool aggregate, the SSD RAID group size can be different from the RAID group size for the HDD RAID groups. Usually, you should ensure that you have only one SSD RAID group for a Flash Pool aggregate, to minimize the number of SSDs that are required for parity.

For information about best practices for working with aggregates, see Technical Report 3437: *Storage Subsystem Resiliency Guide*.

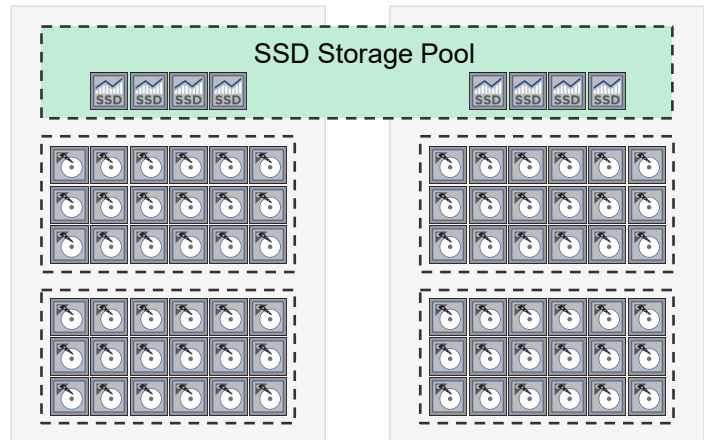
To see the physical and usable capacity for a specific drive, see the Hardware Universe at [hwu.netapp.com](http://hwu.netapp.com).



## SSD storage pool for Flash Pool caching

SSDs can be dedicated to a single Flash Pool aggregate or shared through an SSD storage pool.

- Increased storage use for SSDs in Flash Pool aggregates
- Ability to share an SSD RAID group between HA partners
- Better use of SSD performance



SSD partitioning for Flash Pool intelligent caching enables customers to group SSDs into a shared resource, which is allocated to multiple Flash Pool aggregates. The feature spreads the cost of the parity SSDs over more aggregates, increases SSD allocation flexibility, and maximizes SSD performance.




## Complete an exercise

Module 5  
Physical storage management

### Creating a Flash Pool aggregate

- Access your lab equipment.
- Open your Exercise Guide, Module 5.
- Complete Exercise 3.
- Share your results.

This exercise requires approximately  
**20 minutes.**


 32 © 2023 NetApp, Inc. All rights reserved.

Open the exercise guide and complete Module 5 Exercise 23.



## Lesson 4

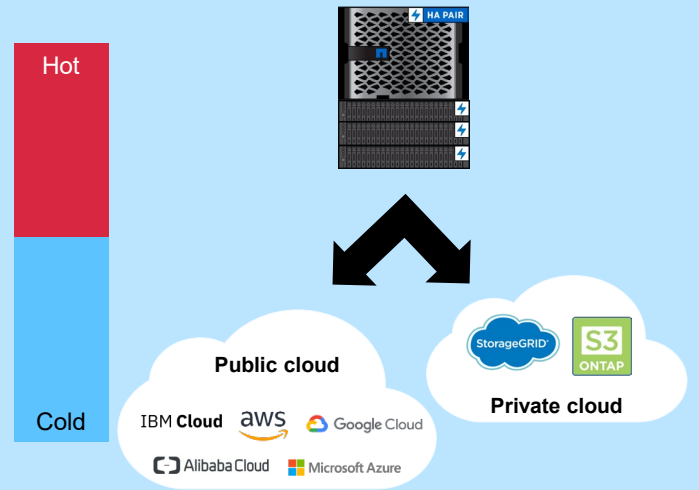
# FabricPool aggregates

 33 © 2023 NetApp, Inc. All rights reserved.

## FabricPool aggregates

### Overview

- What FabricPool aggregates contain:
  - A **performance tier** for frequently accessed (“hot”) data, which is on a local tier
  - A **cloud tier** for infrequently accessed (“cold”) data, which is on an object store
- How FabricPool technology can enhance the efficiency of your storage system:
  - Automatically tier data based on frequency of use
  - Move inactive data to lower-cost cloud storage
  - Make more space available on primary storage for active workloads
  - View how much data in a volume is inactive by using inactive data reporting



NetApp 34 © 2023 NetApp, Inc. All rights reserved.

A FabricPool aggregate contains a performance tier for frequently accessed (“hot”) data, which is typically on an all-SSD aggregate. The FabricPool aggregate also has a cloud tier for infrequently accessed (“cold”) data, which is on an object store. FabricPool technology supports object store types that are in the public cloud by using Amazon Simple Storage Service (Amazon S3). FabricPool technology uses the NetApp StorageGRID solution to support object store types in private clouds.

Storing data in tiers can enhance the efficiency of your storage system. FabricPool technology stores data in a tier based on whether the data is frequently accessed. ONTAP software automatically moves inactive data to lower-cost cloud storage, which makes more space available on primary storage for active workloads.

Use the `-fields` parameter of the `volume show` and `aggregate show` commands to view the amount of data that is inactive.

For more information about FabricPool aggregates, see the *Disks and Aggregates Power Guide*.

## Tiering policies

Define what data is tiered and applied to individual volumes

None	Snapshot-only	Auto	All
Data always remains in the performance tier.	This policy is the default policy.	This policy moves “cold” data blocks that are held in both Snapshot copies and the active file system.	All active and Snapshot data is written directly to the cloud tier.
There is no cooling period.	“Cold” Snapshot copy blocks that are not shared with the active file system are tiered.	There is a 31-day minimum cooling period.	There is no cooling period.
	There is a two-day minimum cooling period.		This policy is designed for SnapMirror or SnapVault target volumes.

NetApp 35 © 2023 NetApp, Inc. All rights reserved.

You can choose from one of these four tiering policies.

Volumes with the None tiering policy never move their data out of the performance tier.

Volumes with the Snapshot-only tiering policy move data blocks inside Snapshot copies which are not shared with the active file system to the cloud tier if they have not been accessed for at least 2 days.

The auto tiering policy maximizes space that is available in the performance storage tier. This policy moves all data blocks to the capacity storage tier when the blocks have not been accessed in the previous 31 days.

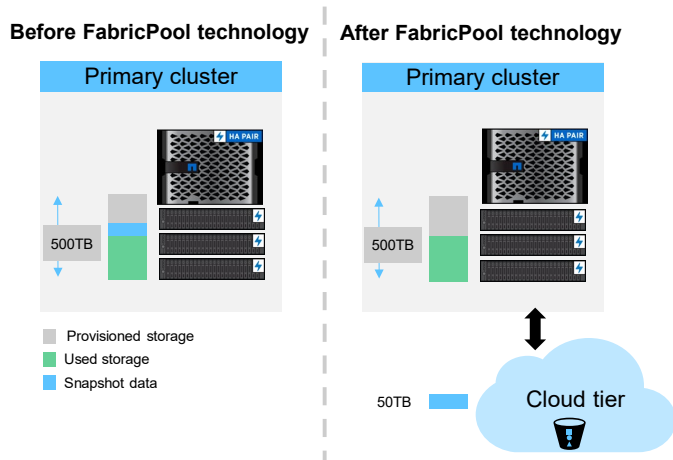
The all tiering policy allows tiering of both Snapshot copy data and active file system user data to the cloud tier as soon as possible without waiting for a cooling period. The all tiering policy was named “backup” before ONTAP 9.6 software. On data protection target volumes, this policy enables all transferred user data blocks to be written to the cloud tier immediately.

You can use the `-tiering-minimum-cooling-days` parameter of the `volume modify` command to change how long data is retained in the performance tier before being moved to the cloud tier.

**Note:** Moving a volume resets the cooling period for all blocks in the volume. This action affects volumes with the Snapshot-only and auto tiering policies, because moved data goes into the performance tier until it cools off.

## Make room for active workloads on primary storage

Snapshot-only tiering to the cloud

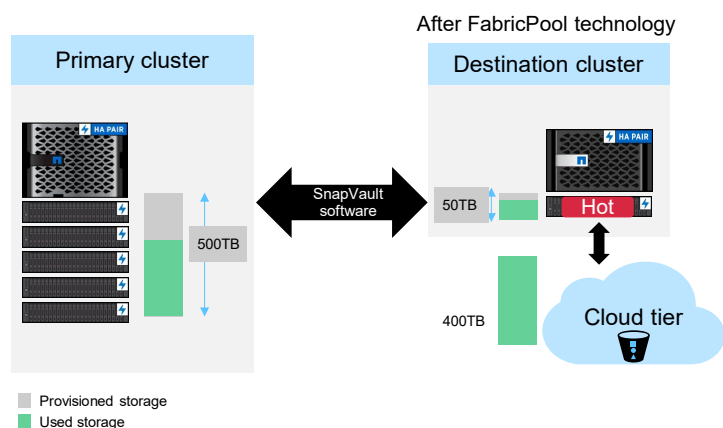


- In this example, Snapshot copies consume ~10% of used capacity.
- Moving Snapshot data enables active workloads to use the performance drives (SSDs) more effectively.

In this example, you see how using FabricPool technology freed up capacity of the primary storage, enabling more active workloads to take advantage of SSDs.

## Shrink your secondary storage footprint

Tier backup data to the cloud



- Expand the capacity of a backup destination cluster by automatically tiering data to the cloud.
- The secondary data center footprint is reduced by up to 90%. Hot data (~10-20%) stays on premises, and the remaining 80-90% goes to the cloud object store.
- This method requires no changes to existing data protection policies. It works seamlessly.



More info in Addendum

NetApp 37 © 2023 NetApp, Inc. All rights reserved.


FabricPool volumes with the “all” tiering policy write data blocks to the cloud tier without waiting for the blocks to cool. The local performance tier is used for staging and to store volume metadata only.

Learn more about FabricPool technology in the module addendum.

If you would like to practice with ONTAP FabricPool technology, see “Storage Tiering in the Cloud with NetApp FabricPool” by the NetApp Lab On Demand team.

## FabricPool mirrors

- Protect against disaster by synchronously replicating data to two object stores.
- Recover from an outage by promoting the mirror object store to primary.
- Enable easy migration of data between an on-site object store and the cloud and between public cloud providers.



NetApp 38 © 2023 NetApp, Inc. All rights reserved.

You can synchronously mirror your FabricPool data to two different object stores. FabricPool mirroring protects you from an object store failure and enables you to easily migrate data from one object store to another.

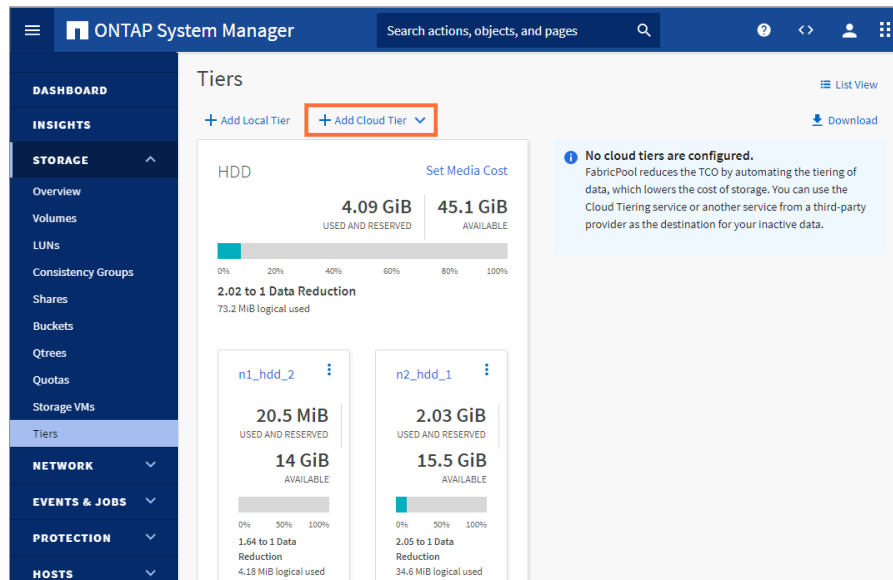
When a second object store is added to a FabricPool aggregate, it becomes the mirror. The original object store is the primary. In the event of an outage, the mirror object store can be promoted to become the primary.

FabricPool mirrors provide a seamless way to migrate data from one object store to another. FabricPool mirrors enable you to easily move data between on-premises storage and cloud-based object stores and between public cloud providers.



# FabricPool technology in ONTAP System Manager

## Adding cloud tiers



NetApp 39 © 2023 NetApp, Inc. All rights reserved.

FabricPool aggregates are aggregates that have an object store attached. You set up an aggregate to use FabricPool technology by first specifying the configuration information of the object store that you plan to use as the cloud tier. Then you attach the object store to an all-flash (all-SSD) aggregate.

Using ONTAP System Manager enables you to create an aggregate and set it up to use FabricPool technology at the same time. (When you use the ONTAP CLI to set up an aggregate for FabricPool technology, the aggregate must exist.)

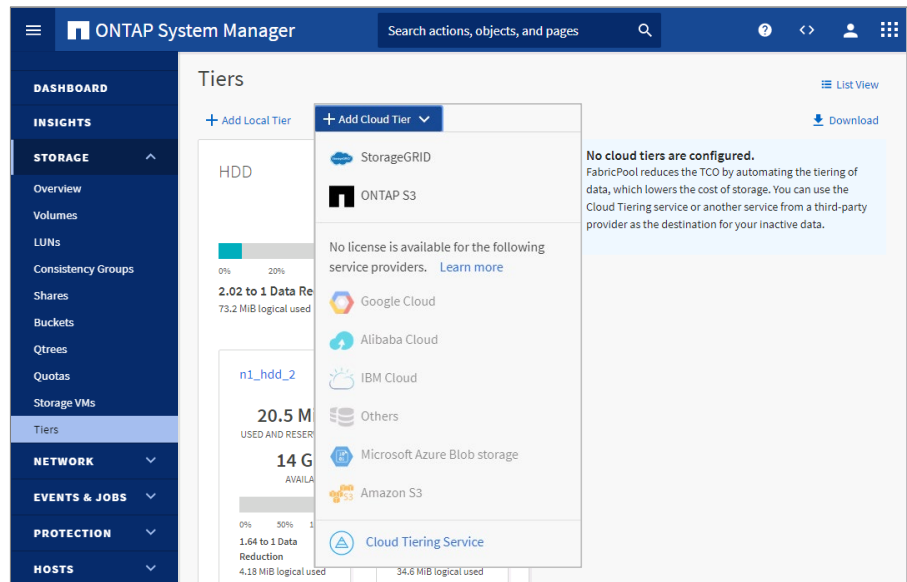
In the Tiers page, use the **Add Cloud Tier** button to add an object store and assign it to an aggregate, to use FabricPool technology.

# FabricPool technology in ONTAP System Manager

## Adding an external capacity tier

Select the object store provider for the cloud tier.

Tiering ONTAP volumes to another ONTAP system, to a StorageGRID solution, or to NetApp Cloud Tiering service does not require an ONTAP FabricPool license.



NetApp 40 © 2023 NetApp, Inc. All rights reserved.

You can use the Add Cloud Tier menu to configure connections to many cloud object store providers.

You can also tier cold data to your StorageGRID private cloud or to other ONTAP systems by using the S3 protocol.

## FabricPool technology in ONTAP System Manager

### Volume tiering policy

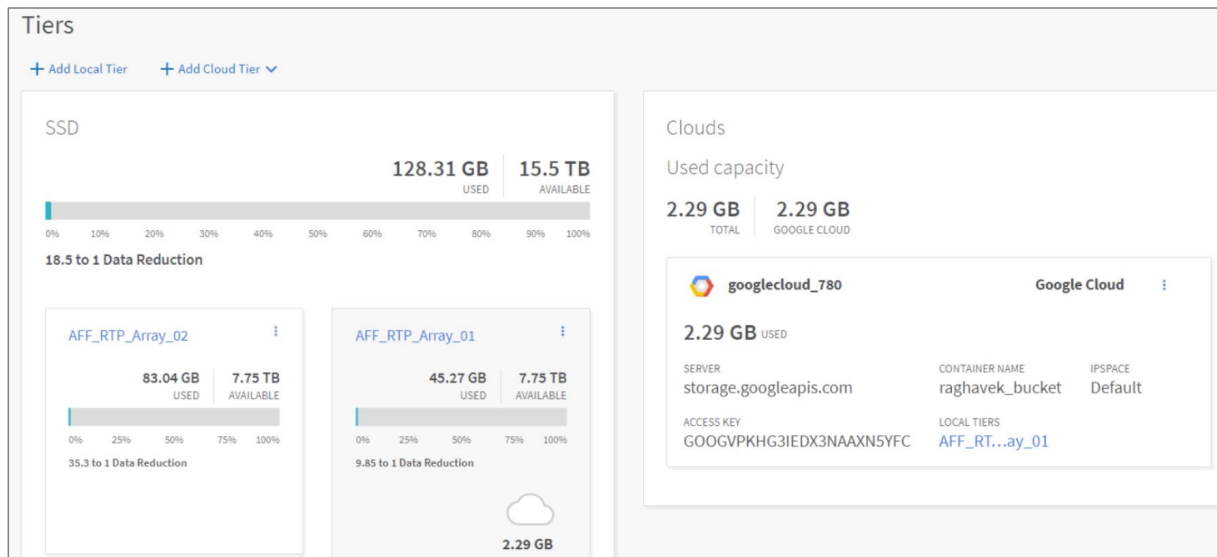
When you add a cloud tier, be aware of the following:

- You should update the volume tiering policy for any volumes that exist in the aggregate.
- The default tiering policy for existing volumes is None.
- Changing the tiering policy of a volume causes data to be migrated to the cloud tier.

Volumes	Storage VM	Inactive Data Capacity	Tiering Policy
<input checked="" type="checkbox"/> S3objectstore	svm3	-	None

When configuring a cloud tier, you are prompted for an existing aggregate to attach. If that aggregate already contains FlexVol volumes, you have an opportunity to update the volume tiering policy. Volumes that were created in non-FabricPool aggregates have their volume tiering policy set to none. When those aggregates are attached to a cloud tier, the data in any existing volumes is not written to the cloud object store. Changing the volume tiering policy causes any data in the existing volumes to be eligible to be written to the cloud object store after the specified cooling period has elapsed.

## FabricPool technology in ONTAP System Manager



NetApp 42 © 2023 NetApp, Inc. All rights reserved.

After you configure a cloud tier, the Tiers page includes local tier and cloud tier information.

You need to know how much data is stored in the local and cloud tiers for FabricPool technology. That information helps you to determine whether you need to change the tiering policy of a volume, increase the FabricPool licensed usage limit, or increase the storage space of the cloud tier.

You can also obtain information about the configuration and use of a cloud tier object store by using the `storage aggregate object-store` commands:

```
cluster1::> storage aggregate object-store show
```


```
cluster1::> storage aggregate object-store show-space
```

## Volume tiering policy

When volumes are created on a FabricPool-enabled aggregate, be aware of the following:

- You should select a tiering policy. The default policy is Snapshot-only.
- Changing the tiering policy of a volume after creation might cause data to be migrated to the cloud tier.
- You can change the cooling off period only for volumes with the Snapshot-only or auto tiering policy.

```
cluster1::> volume modify -vserver svml -volume voll -tiering-policy auto  
-tiering-minimum-cooling-days 15
```

 43 © 2023 NetApp, Inc. All rights reserved.

When you create a volume in FabricPool technology, you can specify a tiering policy. If no tiering policy is specified, the created volume uses the default Snapshot-only tiering policy.

You can change the tiering policy to control how long it takes for data to become cold and be moved to the cloud tier.

Changing the tiering policy from Snapshot-only or None to Auto causes ONTAP software to send active user data blocks that are already cold to the cloud tier. Changing the tiering policy to All causes ONTAP software to move all user blocks in the active file system and in the Snapshot copies to the cloud tier.


Cold data blocks in the cloud tier are returned to the performance tier when they are read, unless the tiering policy is set to All. You can migrate cold data blocks back to the performance tier by setting the volume tiering policy to None and the cloud retrieval policy to Promote.

Any time that you change the tiering policy on a volume, the tiering minimum cooling period is reset to the default value for the policy.



# Knowledge check

Module 5: Physical storage management

 44 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

### Which statement about Advanced Disk Partitioning is true?

- a. Both nodes must have a root partition or a data partition assigned.
- b. Both nodes must have a root partition assigned.
- c. Data partitions can be assigned to any node in a cluster.
- d. Root partitions can be assigned to any node in a cluster.

## Knowledge check

# What does a Flash Pool aggregate contain?

- a. HDDs only
- b. SSDs only
- c. HDDs for data storage and SSDs for caching
- d. HDDs and SSDs that are all used for data caching



## References

- NetApp Hardware Universe  
<http://hwu.netapp.com>
- ONTAP 9 Documentation Center  
<http://docs.netapp.com/ontap-9/index.jsp>
- TR-4070: NetApp Flash Pool Design and Implementation Guide  
<https://www.netapp.com/us/media/tr-4070.pdf>
- TR-4598: FabricPool Best Practices  
<https://www.netapp.com/us/media/tr-4598.pdf>

NetApp Hardware Universe: <http://hwu.netapp.com>

ONTAP 9 Documentation Center: <http://docs.netapp.com/ontap-9/index.jsp>

TR-4070: NetApp Flash Pool Design and Implementation Guide

<https://www.netapp.com/us/media/tr-4070.pdf>

TR-4598: FabricPool Best Practices

<https://www.netapp.com/us/media/tr-4598.pdf>

## Module summary

This module focused on enabling you to do the following:

- Recognize ONTAP storage architecture concepts
- Manage physical storage resources, including disks, RAID groups, and aggregates
- Create data aggregates
- Create Flash Pool aggregates
- Set up FabricPool aggregates




## Complete an exercise

Module 5  
Physical storage management

### Creating a FabricPool aggregate

- Access your lab equipment.
- Open your Exercise Guide, Module 5.
- Complete Exercise 4.
- Share your results.


This exercise requires approximately  
**20 minutes.**

 49 © 2023 NetApp, Inc. All rights reserved.

Open the exercise guide and complete Module 5 Exercise 4.

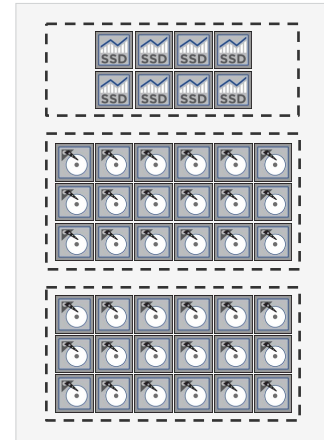


## Addendum Caching policies

 50 © 2023 NetApp, Inc. All rights reserved.

## Setting caching policies

- Caching policies determine how data and metadata are cached in Flash Cache modules and Flash Pool aggregates.
- Caching policies can be applied to storage VMs, volumes, LUNs, or files.
- The caching policy should be changed only if a different policy provides better performance for your workload.
- If the wrong caching policy is configured, volume performance might degrade severely, and performance degradation could increase gradually over time.



```
cluster1::> volume modify -server svm1 -volume vol1 -caching-policy random_read
```


**NetApp** 51 © 2023 NetApp, Inc. All rights reserved.

A caching policy defines how the system caches data in a Flash Pool aggregate or Flash Cache module. Caching policies can be applied to storage VMs (or storage virtual machines, also known as SVMs), volumes, LUNs, or files. Cache policies have the same restrictions on overlap as quality-of-service (QoS) policy groups. If you set a caching policy on a storage VM, that policy is set on all data objects (LUNs, volumes, and so on). You cannot set a different caching policy on a volume, LUN, or file that is owned by the storage VM. The same caching policy is set on all objects that are owned by the storage VM.

The default caching policy works well in most situations. Changes to the caching policy can improve or degrade I/O performance. You should understand the I/O characteristics of your workloads before altering the caching policies and evaluate the effects of any changes.

## Caching policies

Policy name	Insertions using read caching policy				Write insertions
	Random reads	Sequential reads	Random writes	Sequential writes	
auto	Yes	No	No	No	Yes
none	No	No	No	No	No
random_read	Yes	No	Yes	No	No
noread-random_write	No	No	No	No	Yes
meta	Metadata only	No	No	No	No
meta-random_write	Metadata only	No	No	No	Yes
random_read_write	Yes	No	Yes	No	No
random_read_write- random_write	Yes	No	Yes	No	Yes
all_read	Yes	Yes	No	No	No
all_read-random_write	Yes	Yes	Yes	No	Yes
all	Yes	Yes	Yes	Yes	No
all-random_write	Yes	Yes	Yes	Yes	Yes

 52 © 2023 NetApp, Inc. All rights reserved.


If a caching policy is not assigned to an object when it is created, the system uses auto as the default caching policy. Both metadata and user data are eligible for caching. Metadata consists of directories, indirect blocks, and system metafiles. They are eligible for read caching only. When a random write pattern is detected on user data, the first such write is eligible for read caching, but all subsequent overwrites are eligible for write caching.

The available caching policies are the following:

- none: Does not cache any user data or metadata blocks
- auto: Read caches all metadata and randomly read user data blocks and write caches all randomly overwritten user data blocks
- meta: Read caches only metadata blocks
- meta-random\_write: Read caches all metadata and write caches randomly overwritten user data blocks
- random\_read: Read caches all metadata and randomly read user data blocks
- random\_read\_write: Read caches all metadata, randomly read, and randomly written user data blocks
- random\_read\_write-random\_write: Read caches all metadata, randomly read, and randomly written user data blocks. It also write caches randomly overwritten user data blocks.
- all\_read: Read caches all metadata, randomly read, and sequentially read user data blocks
- all\_read-random\_write: Read caches all metadata, randomly read, and sequentially read user data blocks. It also write caches randomly overwritten user data blocks.
- all: Read caches all data blocks that are read and written. It does not do any write caching.
- all-random\_write: Read caches all data blocks that are read and written. It also write caches randomly overwritten user data blocks.
- noread-random\_write: Write caches all randomly overwritten user data blocks. It does not do any read caching.

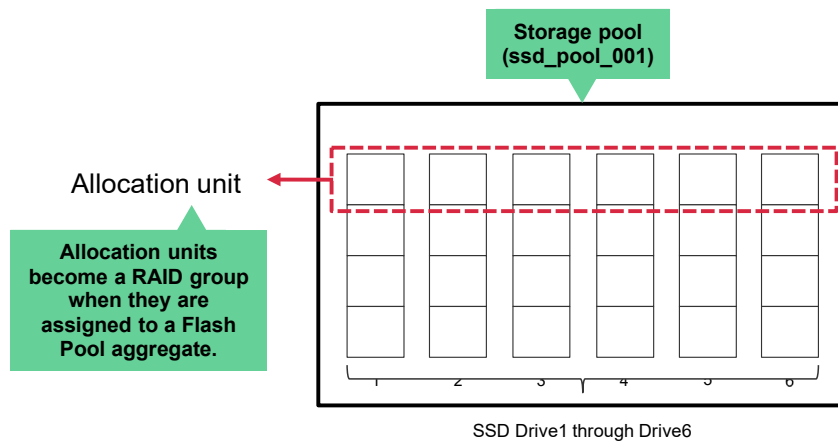


# Addendum Flash Pool SSD storage pools

 53 © 2023 NetApp, Inc. All rights reserved.

## SSD partitioning for Flash Pool cache

Creation



NetApp 54 © 2023 NetApp, Inc. All rights reserved.

SSD storage pools provide SSD caching to two or more Flash Pool aggregates. Creating an SSD storage pool requires between two and 28 spare SSD drives.

In the example, SSD Drive1 through Drive6 are available as spares. The `storage pool create` command is used to create the storage pool. The unit of allocation for an SSD storage pool is equal to a single slice from each SSD drive in the storage pool. The `storage pool create` command slices each SSD drive into four equal pieces, creating an allocation unit that equals one-fourth of all the SSD disks in the storage pool.

An allocation unit becomes a RAID group when the allocation unit is assigned to a Flash Pool aggregate.



## Create an SSD storage pool

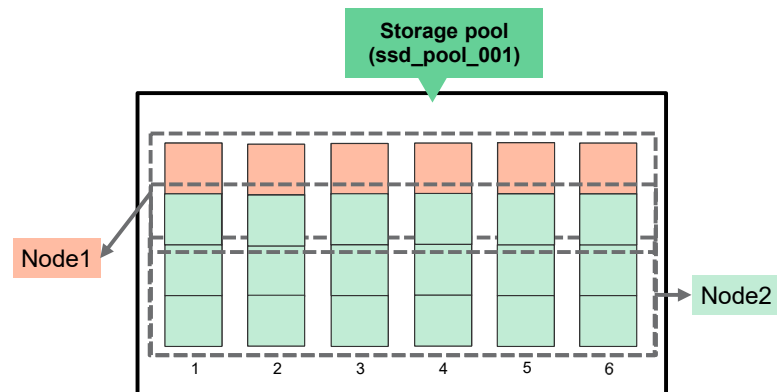
Provide the following information:

- Storage pool name
- Number of drives
- Size of SSDs from the HA pair  
(if multiple sizes are available)

```
::> storage pool create -storage-pool ssd_pool_001 -disk-count 3
```

## SSD partitioning for Flash Pool cache

### Ownership




```
::> storage pool reassign -storage-pool ssd_pool_001  
-from-node cluster1-01 -to-node cluster1-02 -allocation-units 1
```

By default, half of the four allocation units are assigned to each node in the HA pair. To change the ownership of one or more allocation units of a storage pool from one HA partner to the other, use the `storage pool reassign` command. In the example, one allocation unit is reassigned from Node1 to Node2.

# SSD partitioning for Flash Pool cache

## Reassignment



The reassigned allocation group no longer belongs to node 1. The reassigned allocation unit becomes a RAID group when the allocation unit is assigned to a Flash Pool aggregate on node 2.

## Creating a Flash Pool aggregate that uses an SSD storage pool

Use the `storage aggregate add-disk` command to create a Flash Pool aggregate that uses an SSD storage pool.

Provide the following information:

- An existing aggregate name
- A storage pool name
- The number of allocation units to add

```
::> storage aggregate add-disks -aggregate rtp01_fca1_002  
-allocation-units 1 -storage-pool ssd_pool_001
```

 58 © 2023 NetApp, Inc. All rights reserved.


Use the `storage aggregate add-disk` command to create a Flash Pool aggregate.

The `storage-pool` parameter specifies the name of the SSD storage pool from which available allocation units are added to a given aggregate.

The `allocation-units` parameter specifies the number of allocation units to be added to a given aggregate from an SSD storage pool. You can determine the number of allocation units that are available and the size of each unit by using the `storage pool show-available-capacity` command.

# Module 6

## Logical storage management

 1 © 2023 NetApp, Inc. All rights reserved.

## About this module


This module focuses on enabling you to do the following:

- Create and manage FlexVol volumes
- Move a volume within a storage VM (storage virtual machine, also known as SVM)
- Create a NetApp ONTAP FlexGroup volume

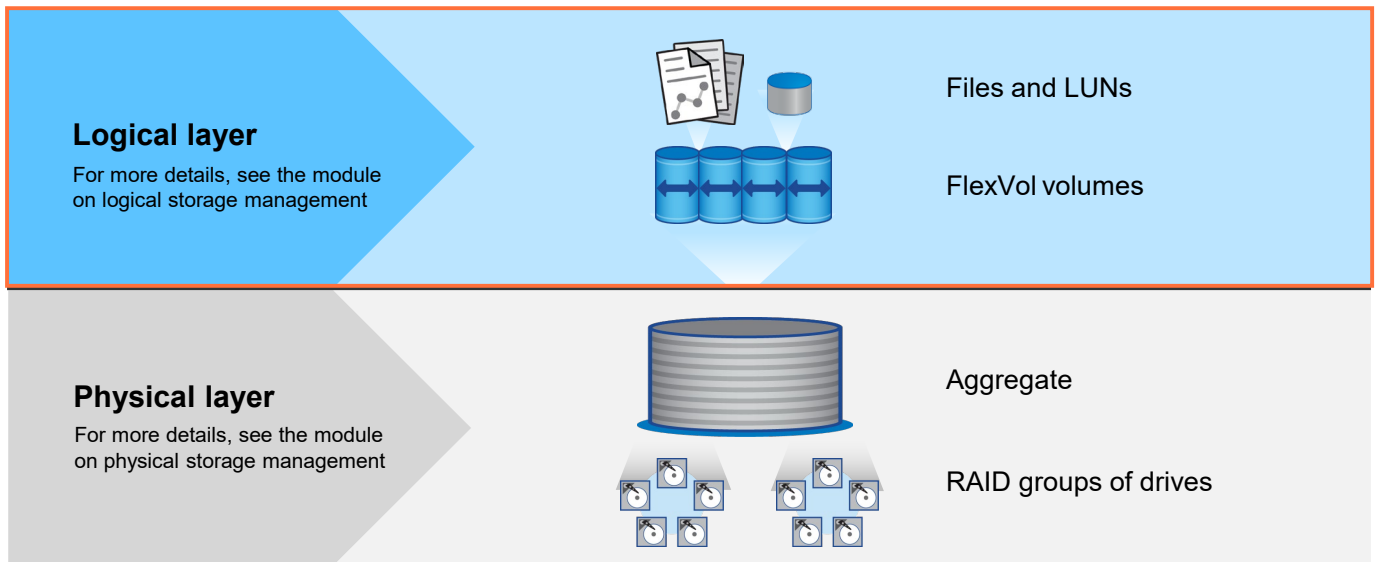


# Lesson 1

## Flexible volumes

 3 © 2023 NetApp, Inc. All rights reserved.

## ONTAP storage architecture



NetApp 4 © 2023 NetApp, Inc. All rights reserved.

As you learned in the physical storage management module, disks are grouped into RAID groups, and RAID groups are grouped into aggregates to form the physical storage layer.


Within each aggregate, you can create one or more FlexVol volumes. A FlexVol volume is an allocation of disk space that is a portion of the available space in the aggregate. A FlexVol volume can contain files or LUNs. The FlexVol volumes, files, and LUNs make up the logical storage layer.



## FlexVol volumes

A FlexVol volume is a container for data.

- Can contain NAS or SAN data
  - Mixing NAS and SAN in the same volume is not recommended.
- Is contained within an aggregate
  - An aggregate can hold multiple FlexVol volumes.
- Can increase or decrease in size, as needed
  - The minimum volume size is 20MB.
  - The maximum volume size is 100TB and 2 billion files.



**NetApp** 5 © 2023 NetApp, Inc. All rights reserved.


A FlexVol volume is loosely coupled to a containing aggregate, which the volume can share with other FlexVol volumes. Therefore, one aggregate can be the shared source of all the storage that is used by all the FlexVol volumes that the aggregate contains.

Because a FlexVol volume is managed separately from the aggregate, you can create small (minimum of 20MB) FlexVol volumes. You can also increase or decrease the size of FlexVol volumes in increments as small as 4KB. A FlexVol volume can grow as large as 100TB and can fill an entire aggregate. A single FlexVol volume can contain up to 2 billion NAS files. In ONTAP 9.8 software, for All SAN Array systems only, the maximum volume size was increased to 300TB.

## FlexVol volumes

### Types

- **System (or node root):**
  - Typically named vol0
  - Contains only configuration and logs
  - Cannot contain user data
  - Owned by the node storage VM
- **Storage VM root volume:**
  - Top level of the namespace
  - Should not contain user data
- **Data volumes:**
  - **NAS:** Contains file systems for user data
  - **SAN:** Contains LUNs for SAN client hosts
  - **Data protection:** Contains backup Snapshot copies




FlexVol volumes are used for the following purposes:

- As node root volumes to hold state data for the node and for the cluster
- As the root of a storage VM (storage virtual machine, also known as SVM) namespace
- To store user data within a storage VM

## Files and LUNs

- A file refers to any data (including text files, spreadsheets, and databases) that is exported to or shared with NAS clients.
- A LUN represents a logical drive that a SCSI protocol (FC or iSCSI) addresses:
  - It is block-level storage.
  - Data is accessible by only a properly mapped SCSI host.



**NetApp** 7 © 2023 NetApp, Inc. All rights reserved.

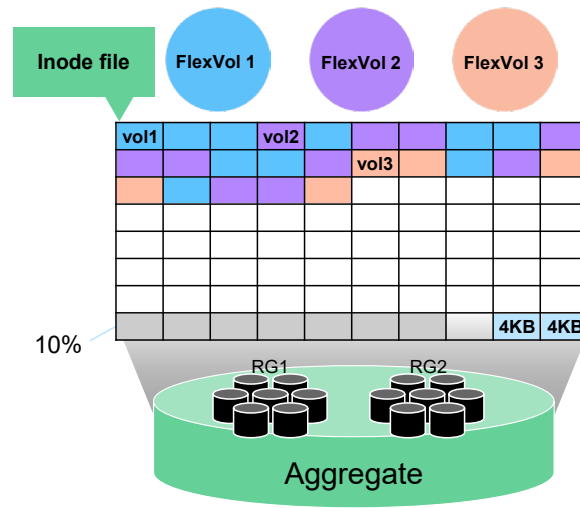
Data that is stored in a volume for a NAS environment is stored as files. Files can be documents, database files and logs, audio and video, or application data. NetApp ONTAP software manages the file system operations, and clients access the data.

Data that is stored in a SAN environment is stored in a logical container that represents a SCSI disk. The container is called a LUN. The LUN is presented to a host, which treats the LUN like a standard SCSI disk and writes data to the LUN in 512-byte logical blocks. Therefore, SAN is often called block-level storage. ONTAP software is “unaware” of the stored files and is “aware” only of the 512-byte blocks to which the host reads or writes.

**NOTE:** Because SAN data (block data) and NAS data (file data) are treated differently, files and LUNs should not be placed in the same FlexVol volume.

## Volumes in aggregates

- **Aggregate:**
  - 4KB blocks
  - NetApp WAFL file system reserving 10%
- **Volume:**
  - Provisioning types:
    - **Thick:** Volume guarantee = volume
    - **Thin:** Volume guarantee = none
  - Dynamic mapping to physical space



NetApp 8 © 2023 NetApp, Inc. All rights reserved.

You can create one or more FlexVol volumes in an aggregate. To understand how space is managed, examine how space is reserved in the aggregate.

The NetApp WAFL file system writes data in 4KB blocks that are contained in the aggregate. (Each 4KB block has an inode pointer. The inode pointers that are assigned to a data file are tracked in the inode file.) When the aggregate is created, the WAFL file system reserves 10% capacity for overhead, but this space is dynamically allocated. The remainder of the aggregate is available for volume creation.

FlexVol volumes are loosely tied to their aggregates. FlexVol volumes are striped across all the drives of the aggregate, regardless of the volume size. In the example, the blue block that is labeled “vol1” represents the inode file for the volume, and the other blue blocks contain the user data.

When a volume is created, the volume guarantee setting must be configured. The volume guarantee setting is the same as the space reservations. If space is reserved for the volume, the volume is thick-provisioned. If space is not reserved during creation, the volume is thin-provisioned. FlexVol volumes are dynamically mapped to physical space. Whether the volume is thick-provisioned or thin-provisioned, blocks are not consumed until data is written to the storage system.

A FlexVol volume can be as small as 20MB or as large as the controller model supports. Also, the volume can grow or shrink, regardless of the provisioning type.

## Volume properties

### Actions that can be taken on volumes




- Create
- Edit
- Resize
- Delete
- Clone
- Move
- Rehost

### Volume options



- Storage efficiency
- Storage quality of service (QoS)

### Tools to protect volumes



- Snapshot copies
- Mirror copies
- Vault copies

FlexVol volumes can be created, modified, resized, deleted, cloned, moved, and rehosted.

Each FlexVol volume maintains a set of configuration properties which governs its behavior. Storage efficiency settings control deduplication and compression activity. Quality of service settings enable guardrails to be placed around client activity.

FlexVol volumes can be protected by making Snapshot copies and backup copies and through replication.

## Creating a flexible volume in a storage VM

Information to provide:

- Volume name
- Storage VM name
- Capacity
- Service level

The screenshot shows the NetApp ONTAP System Manager interface. The 'Volumes' page is active, and the 'Add Volume' dialog is open. The dialog has the following fields and options:

- NAME:** cluster4\_datastore\_08
- STORAGE VM:** svm3
- CAPACITY:** 800 GB
- PERFORMANCE SERVICE LEVEL:** Performance
- OPTIMIZATION OPTIONS:**  Distribute volume data across the cluster

A red box highlights the 'More Options' button in the dialog. The background shows a list of volumes with columns for Name, Storage VM, Capacity, and Status.

```
::> volume create -vserver svm4 -name svm4_vol_002  
-aggr cluster2_01_SAS_1 -size 200gb
```

NetApp 10 © 2023 NetApp, Inc. All rights reserved.

When using NetApp ONTAP System Manager (formerly OnCommand System Manager) to create a volume, you must specify the new volume name and the storage VM that owns it. The capacity of the new volume can be as small as 20MB or as large as the system limits.

System Manager creates the volume in the optimal location based on your provisioning recommendations. New volumes that are assigned the performance service level are provisioned from aggregates that are made of SSDs where possible. New volumes that are assigned the value service level are provisioned from aggregates that are made of HDDs where possible.

The “Add as a cache for a remote volume” checkbox (not shown) causes System Manager to configure the new volume as a FlexCache volume of an existing volume.

The “Distribute volume across the cluster” checkbox causes System Manager to create a NetApp ONTAP FlexGroup volume instead of a FlexVol volume.

FlexCache and FlexGroup volumes are discussed later in this module.

## Balanced placement

Storage service levels



Service level	Application-aligned storage service levels		
	Value	Performance	Extreme
<b>Workload type</b>	Email, web, file shares, and backup	Database and virtualized applications	Latency-sensitive applications
<b>Expected IOPS</b> (IOPS per TB allocated)	128	2048	6144
<b>Maximum service-level objective (SLO)</b> (QoS limit in IOPS per TB stored)	512	4096	12288
<b>Minimum SLA</b> (IOPS per TB allocated)	75	500	1500
<b>Latency (ms)</b>	17	2	1

### Balanced use of cluster resources

- Simplified provisioning
- Intelligent placement that is based on the size of the provisioned components, desired storage service levels, and available system resources
- Predefined storage service levels to match the media with requested performance characteristics (QoS)

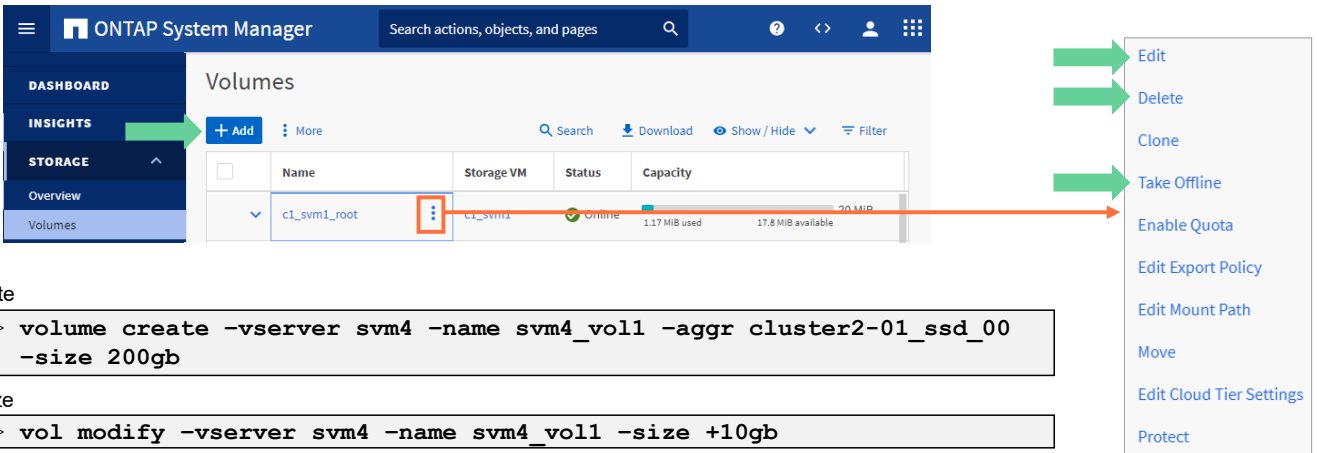
NetApp 11 © 2023 NetApp, Inc. All rights reserved.

Storage service levels help to ensure that limited or expensive cluster resources are dedicated to high-priority workloads. The effects are more noticeable the larger the cluster and the greater the mix of controller models and drive types in the cluster.

You can use the `storage-service show` command at the advanced administrative level to view (but not change) the performance service levels. This command also displays the aggregates that are associated with each service level.

You must use the ONTAP API to define application templates that associate the application with a storage performance service level.

## Managing FlexVol volumes



The screenshot shows the ONTAP System Manager interface. The left sidebar has a 'STORAGE' section with 'Volumes' selected. The main area displays a table of volumes. The first row is highlighted, and a context menu is open over it. The menu items are: Edit, Delete, Clone, Take Offline, Enable Quota, Edit Export Policy, Edit Mount Path, Move, Edit Cloud Tier Settings, and Protect. Green arrows point from the 'Add' button and the 'Volumes' menu item to the screenshot. An orange arrow points from the context menu icon in the table to the context menu itself.

	Name	Storage VM	Status	Capacity
<input type="checkbox"/>	c1_svm1_root	c1_svm1	Online	1.17 MiB used / 17.8 MiB available / 20 MiB

- Edit
- Delete
- Clone
- Take Offline
- Enable Quota
- Edit Export Policy
- Edit Mount Path
- Move
- Edit Cloud Tier Settings
- Protect

### Create

```
::> volume create -vserver svm4 -name svm4_vol1 -aggr cluster2-01_ssd_00 -size 200gb
```

### Resize

```
::> vol modify -vserver svm4 -name svm4_vol1 -size +10gb
```

### Destroy


```
::> vol offline -vserver svm4 -name svm4_vol1  
::> vol delete -vserver svm4 -name svm4_vol1
```

You can perform actions on volumes by using ONTAP System Manager or the CLI volume commands.



## Automatic resizing of volumes

- Automatic resizing of volumes enables a FlexVol volume to automatically grow or shrink the maximum space capacity of the volume.
- You can specify a mode:
  - Off: The volume does not grow or shrink.
  - Grow: The volume automatically grows when space in the volume reaches a threshold.
  - Grow\_shrink: The volume automatically grows or shrinks in response to the amount of used space.
- Also, you can specify the following:
  - Maximum to grow (default is 120% of the initial volume size)
  - Minimum to shrink (default is the initial volume size)
  - Grow and shrink thresholds



**NetApp** 13 © 2023 NetApp, Inc. All rights reserved.

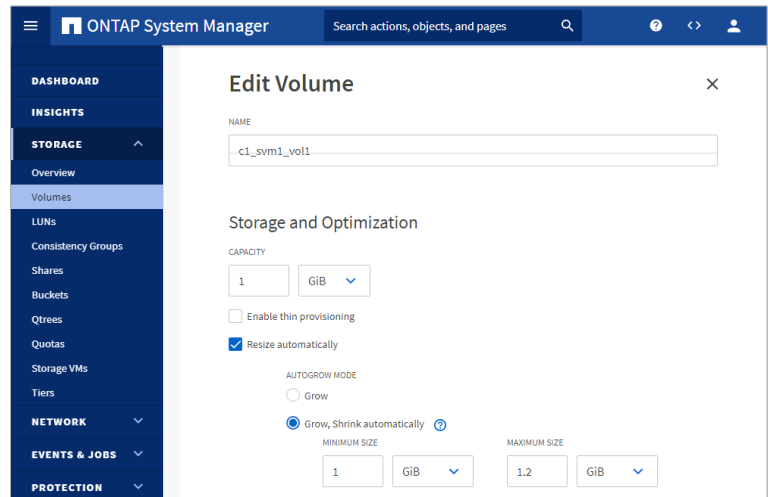
You can enable or disable automatic resizing of volumes. If you enable the capability, ONTAP software automatically increases the capacity of the volume up to a predetermined maximum size. Space must be available in the containing aggregate to support the automatic growth of the volume. Therefore, if you enable automatic resizing, you must monitor the free space in the containing aggregate and add more space when needed.

The capability cannot be triggered to support Snapshot copy creation. If you attempt to create a Snapshot copy and the volume has insufficient space, the Snapshot copy creation fails, even when automatic resizing is enabled.

For more information about automatic resizing, see the *SAN Administration Guide*.

## Enabling automatic resizing


1. In the Volumes page, select the volume and select **Edit** from the menu.
2. Select the **Resize automatically** checkbox.
3. Select an Autogrow Mode option.
4. Specify the Maximum Size value.



```
::> volume autosize -vserver svm4 -volume svm4_vol_002  
-mode grow -maximum-size 200GB
```



## Try this task

 15 © 2023 NetApp, Inc. All rights reserved.

Use cluster1 in your lab environment:

1. Enter the `vol show` command.
2. Enter the `vol show -instance` command.
3. Enter the `vol show -fields comment` command.
4. Answer the following questions:
  - What was different about the output?
  - Can you think of other reasons to use `-fields`?
  - How can you get a list of all the fields that are available for a command?


Review the following answers:

- A different amount of information is displayed about each volume.
- Use the `-fields` parameter to customize the command output for your requirements.
- Type a question mark (?) after the `-fields` parameter to get a list of available fields.



## Lesson 2

# FlexGroup volumes


 16 © 2023 NetApp, Inc. All rights reserved.

## What is a FlexGroup volume?

- A scale-out file system that is created from a group of FlexVol volumes
- A system that you and NAS clients can interact with as you interact with a FlexVol volume

### FlexGroup volumes solve three problems with modern NAS in scale-out storage:

- **Performance:** FlexGroup volumes provide consistently low latency.
- **Capacity:** FlexGroup volumes provide almost unlimited capacity.
- **Management:** A FlexGroup volume looks like a FlexVol volume.



NetApp 17 © 2023 NetApp, Inc. All rights reserved.


With NetApp ONTAP FlexGroup volumes, you can easily provision a massive single namespace in seconds. A FlexGroup volume has a 20PB capacity limit with as many as 400 billion files in 200 constituent volumes. The constituent volumes in a FlexGroup volume collaborate to dynamically balance load and space allocation among themselves.

A FlexGroup volume requires no maintenance or management overhead. You simply create the FlexGroup volume and share the volume with your NAS clients. ONTAP software does the rest.

For more information about FlexGroup volumes, see *NetApp ONTAP FlexGroup Volumes (TR-4557)* and *Scalability and Performance Using FlexGroup Volumes Power Guide*.


# FlexVol volumes versus FlexGroup volumes

How they differ at a high level



## FlexVol volumes

- Are owned by one node
- Span one aggregate
- Isolate reads and writes to one node and aggregate
- Are limited to storing 100TB (system-dependent)
- Are within one namespace, but with limits



## FlexGroup volumes


- Are shared pools of FlexVol volumes
- Have component volumes that span multiple aggregates
- Balance reads and writes across all nodes and aggregates
- Can store up to 20PB (200 FlexVol volumes)
- Are within one namespace, almost without limits

Although FlexGroup volumes are positioned as a capacity feature, the volumes are also a high-performance feature. With a FlexGroup volume, you can have massive capacity, predictably low latency, and high throughput for the same storage container. A FlexGroup volume adds concurrency to workloads with no need for increased management.


## Managing FlexGroup volumes

You manage FlexGroup volumes the same way that you manage FlexVol volumes.

- You create the FlexGroup volume, and ONTAP software manages the rest:
  - When you create the FlexGroup volume, you specify the size, aggregates, storage VM, and file system path.
  - ONTAP software creates equally sized constituent volumes.
- If you need more space, you can add a constituent volume anywhere in the cluster.  
ONTAP automatically redistributes files to the new constituent volumes to balance usage.



What ONTAP software sees



What clients see



More info in Addendum


At a high level, a FlexGroup volume is simply a collection of FlexVol volumes that act as one entity. NAS clients access the FlexGroup volume just as they access a FlexVol volume: from an export or a CIFS (or SMB) share.

Although FlexGroup volumes are conceptually like FlexVol volumes, FlexGroup volumes offer several benefits that FlexVol volumes cannot match.


Files in a FlexGroup volume are allocated to individual member FlexVol volumes and are not striped across volumes or nodes. FlexGroup volumes provide throughput gains by performing concurrent operations across multiple FlexVol volumes, aggregates, and nodes. A series of operations can occur in parallel across all hardware on which the FlexGroup volume resides. FlexGroup volumes are the perfect complement to the ONTAP scale-out architecture.


## FlexGroup rebalance

- For optimal performance, distribute data equally across the constituent member volumes.
- Following are use cases for NetApp ONTAP FlexGroup rebalancing:
  - Unequal loading of FlexGroup volume members due to data ingest or deletion patterns



- Expansion of FlexGroup volumes with new volumes



 20 © 2023 NetApp, Inc. All rights reserved.

Restore optimal performance with FlexGroup rebalance:


- Redistribution of files across all member volumes of the FlexGroup volume
- A transparent and nondisruptive event to clients that are connected to the FlexGroup shares
- Simple administration and control from NetApp ONTAP System Manager and NetApp BlueXP

ONTAP software automatically attempts to evenly distribute file data across all the FlexGroup constituent FlexVol volumes. However, the data blocks of any single file must be contained within the same FlexVol constituent volume. A single file can not span multiple constituent volumes. This situation can sometimes result in an imbalance when a few large files are mingled with many small files in the same FlexGroup volume, or when there are unusual file deletion patterns.

A volume use imbalance also occurs when a FlexGroup volume is expanded. Newly added FlexVol constituent volumes are initially empty.


ONTAP 9.11 software includes FlexGroup rebalancing, in which files are automatically redistributed across the constituent FlexVol volumes to make use even. FlexGroup rebalancing avoids overusing new constituent volumes when a nearly full FlexGroup is expanded.





## Lesson 3

# FlexCache volumes

 21 © 2023 NetApp, Inc. All rights reserved.

# FlexCache volumes

Accelerate hot volumes

A NetApp FlexCache volume is a sparsely populated FlexGroup volume that is used to cache data from a particular volume that is called the origin.


Data is copied from the origin to the FlexCache volume the first time a client reads it.

- Cache read data for I/O intensive workloads
- Prepopulate the FlexCache volume with frequently accessed data
- Cache data within the cluster (intracluster) or to remote sites (including the cloud)

Understand FlexCache volume limitations:

- There are no Snapshot copies or NetApp FlexClone support on the cache volume.
- Deduplication, compression, and compaction are supported.
- FlexCache technology supports NFS version 3 (NFSv3), NFS version 4 (NFSv4), and SMB access to the cache volume.

NetApp 22 © 2023 NetApp, Inc. All rights reserved.



NetApp FlexCache volumes cache frequently accessed NAS data to reduce latency within a cluster or between clusters. For local hot volumes, FlexCache volumes provide more copies of the data to spread the I/O demands across the cluster.

The FlexCache volume is a sparsely populated FlexGroup volume. The origin volume can be either a FlexVol volume or a FlexGroup volume.

FlexCache volumes contain temporary copies of some of the data in the source volume. For this reason, the volumes do not support many of the features of a typical FlexVol volume. One limitation is that you cannot clone a FlexCache volume.



With ONTAP 9.8 software, FlexCache volumes can serve both SMB clients and NFSv3 clients. Starting in ONTAP 9.10.1 software, NFSv4 clients are also served.

## FlexCache volumes

Accelerate data access to remote users

Data distribution across data centers:

- Caches across multiple data centers to reduce WAN latencies
- Brings data closer to compute or users or both
- Populates the cache to reduce initial read latencies
- Works among NetApp AFF, FAS, ONTAP Select, and Cloud Volumes ONTAP based systems



NetApp 23 © 2023 NetApp, Inc. All rights reserved.

Caching frequently accessed remote data closer to the users reduces WAN traffic and latency. Caching works between ONTAP based systems that are on premises and in the cloud.

ONTAP 9.8 software also enables the population of the FlexCache volume with data. This ability eliminates the delay that clients would normally experience when a data block is first accessed, and the data is fetched from the origin.


By default, FlexCache volumes favor availability over consistency. Without global file locking, any modification to the origin volume is distributed to FlexCache volumes, but they might not be updated simultaneously. Global file locking, new in ONTAP 9.9.1 software, favors consistency across volumes over availability. With global file locking enabled, modifications to the origin volume are suspended until all FlexCache volumes are online. Starting in ONTAP 9.10 software, exclusive read locks are enforced globally instead of only at the origin volume or FlexCache volume.

The module addendum contains more information about FlexGroup volumes and FlexCache volumes.



## Lesson 4

# Moving storage resources

 24 © 2023 NetApp, Inc. All rights reserved.


## Volume move

- Rules:

- Move only within the storage VM.
- Move to any aggregate to which the storage VM has permission.
- Move is nondisruptive to the client.

- Use cases:

- **Capacity:** Move a volume to an aggregate with more space.
- **Performance:** Move to an aggregate with different performance characteristics.
- **Servicing:** Move to newly added nodes or from nodes that are being retired.



FlexVol volumes can be moved from one aggregate or node to another within the same storage VM. A volume move does not disrupt client access during the move. The data is moved over the cluster interconnect.

You can move volumes for capacity use, such as when more space is needed. You can move volumes to change performance characteristics, such as from a controller with HDDs to one that uses SSDs. You can also move volumes during service periods.

## How a volume move works

1. A volume is created on the destination aggregate.
2. A Snapshot copy of the source volume is created.
3. The Snapshot copy is replicated to the destination volume.
4. When replication is complete, client access is temporarily blocked.
5. A final replication is performed to reach consistency.
6. Cutover is initiated.
7. Clients access the destination volume, and the source volume is cleaned up.



**NetApp** 26 © 2023 NetApp, Inc. All rights reserved.

When a volume move is initiated, a Snapshot copy of the source volume is created and is used as the basis to populate the destination volume. Client systems continue to access the volume from the source destination until all data is moved. At the end of the move process, client access is temporarily blocked. Meanwhile, the system performs a final replication from the source volume to the destination volume. The system swaps the identities of the source and destination volumes and changes the destination volume to the source volume. When the move is complete, the system routes client traffic to the new source volume and resumes client access.

## The volume move command

**Volume move**

```
cluster2::> vol move start -vserver svm4 -vol svm4_vol_002
               -destination-aggr cluster2_fc1_002
cluster2::> vol move trigger-cutover -vserver svm4 -vol svm4_vol_002
```

Use the `vol move start` command to initiate the volume transfer. If the cutover action is `defer on_failure` and the cutover state moves to “cutover deferred,” use the `vol move trigger-cutover` command to complete the move. To bypass any confirmation before cutover, use `-force true` on the `vol move start` command. The bypass can cause client I/O disruptions.

## Autobalancing aggregates

### Default settings

- If you frequently move volumes to free up space, you can use the `autobalance aggregate` command to configure ONTAP software to autobalance automatically for all aggregates.
- The autobalance aggregate feature is turned off by default. See the addendum for more information.

```
::*> autobalance aggregate config show
      Is the Auto Balance Aggregate Feature Enabled: false
      Threshold When Aggregate Is Considered Unbalanced (%): 70
      Threshold When Aggregate Is Considered Balanced (%): 40
```



If you find you spend a lot of time moving volumes to manage free space and performance, consider enabling the autobalance aggregate functionality. After it is enabled, it works on all the aggregates in the cluster.

You can display information about aggregates that are being considered for automatic rebalancing by using the `autobalance aggregate show-unbalanced-aggregate-state` command.

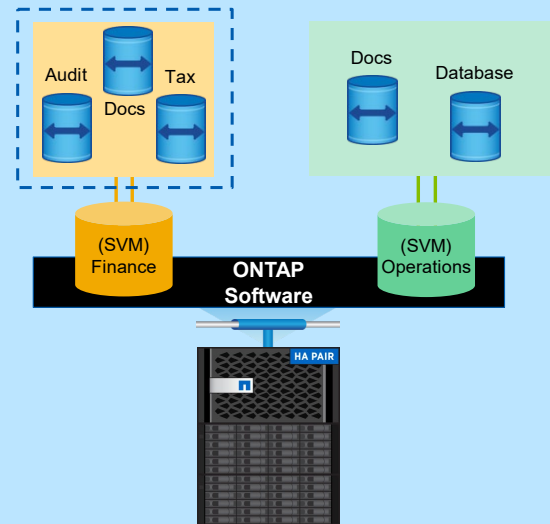
You can display information about volumes that are being considered for relocation by using the `autobalance aggregate show-unbalanced-volume-state` command.



## Volume rehost within a cluster

Steps to rehost a FlexVol volume:

1. Identify the source volume and storage VM.
2. Identify the destination storage VM within the cluster.
3. Prevent access to the volume that is being rehosted.
4. Use the `rehost` command to reassign the volume to the destination storage VM.
5. Configure access to the volume in the destination storage VM.



NetApp 29 © 2023 NetApp, Inc. All rights reserved.

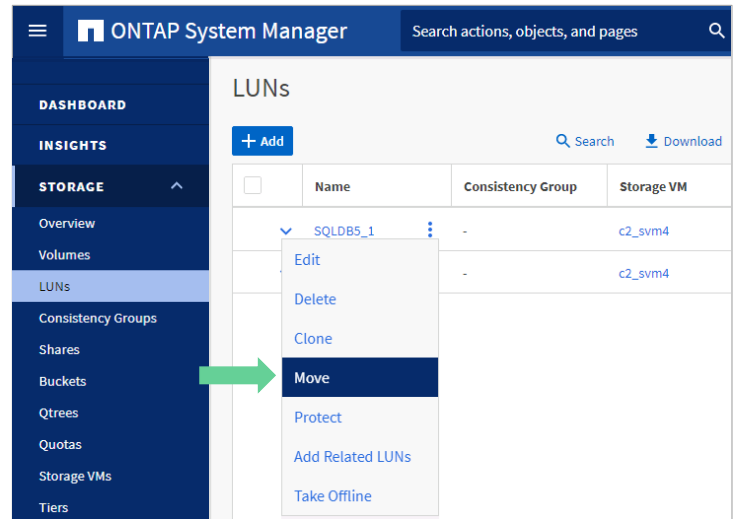
The `volume rehost` command changes the owner of a FlexVol volume from a source storage VM to a destination storage VM. A volume with the same name cannot already exist in the destination storage VM.

If the volume contains a LUN, you can specify that the LUN needs to be unmapped. In addition, you can specify whether you want the LUN to be automatically remapped on the destination storage VM.

**Note:** Volume rehost is a disruptive operation and requires you to reconfigure access to the volume at the destination. Access to the volume must be prevented before a rehost to prevent data loss or inconsistency.

## LUN move

- The `lun move` set of commands enables you to move a LUN to a volume on another node or even the same node.
- The LUN can move only within the same storage VM.
- Snapshot policies are at the volume level, so they do not follow the LUN to the new volume. Therefore, storage efficiency features must be reapplied.
- Use the `lun move-in-volume` command to rename a LUN without moving the LUN.




NetApp 30 © 2023 NetApp, Inc. All rights reserved.

To move a LUN for capacity or performance reasons, use the `lun move` command set rather than moving the container volume. LUNs can be moved only to another volume in the storage VM. You need to set the Snapshot policies on the destination volume. Storage efficiency features, such as deduplication, compression, and compaction, are not preserved during a LUN move. The features must be reapplied after the move is complete.

If you need to rename a LUN, use the `lun move-in-volume` command to “move” the LUN, with a new name, to the current location.

# Knowledge check

Module 6:  
Logical storage management

 31 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

**Which items are modes of the volume automatic resize feature?  
(Choose three)**

- a. off
- b. grow
- c. shrink
- d. grow\_shrink

## References

- [ONTAP 9 Documentation Center](#)
- [TR-4557: NetApp ONTAP FlexGroup Volumes](#)
- [TR-4571: NetApp ONTAP FlexGroup volumes: Best Practices and Implementation Guide](#)
- [TR-4743: FlexCache in ONTAP](#)

[ONTAP 9 Documentation Center](#)

[TR-4557: NetApp ONTAP FlexGroup Volumes](#)

[TR-4571: NetApp ONTAP FlexGroup Volumes: Top Best Practices](#)

[TR-4743: FlexCache in ONTAP](#)

## References Videos

- ONTAP 9 Feature Overview: FlexGroup  
<https://www.youtube.com/watch?v=Wp6jEd4VkgI>
- Manage FlexGroup using OnCommand System Manager 9.4  
<https://www.youtube.com/watch?v=mLpVjoll4GY>

ONTAP 9 Feature Overview: FlexGroup  
<https://www.youtube.com/watch?v=Wp6jEd4VkgI>

Manage FlexGroup using OnCommand System Manager 9.4  
<https://www.youtube.com/watch?v=mLpVjoll4GY>

## Module summary

This module focused on enabling you to do the following:

- Create and manage FlexVol volumes
- Move a volume within a storage VM
- Create a FlexGroup volume



## Complete an exercise

Module 6: Logical storage management

### Managing Data Volumes

#### Creating a FlexGroup volume

- Access your lab equipment.
- Open your Exercise Guide, Module 6.
- Complete Exercise 1 and Exercise 2.
- Share your results.


This exercise requires approximately  
**30 minutes.**

See the instructions in your Exercise Guide.





# Addendum FlexGroup volumes

 37 © 2023 NetApp, Inc. All rights reserved.

## FlexGroup pre-deployment

### Recommended practices

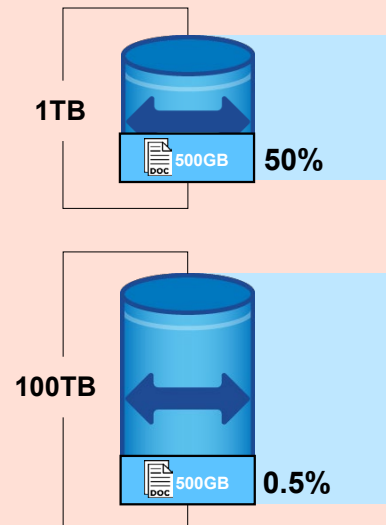
- Confirm homogenous hardware and capacity.
  - Disks, nodes, and available capacity should be identical for predictable performance.
  - Relocate volumes in aggregates by using a nondisruptive volume move, if necessary.
- Use a reliable network that is 10Gb or greater.
  - Flow control is unnecessary across high-bandwidth networks.
- Know the average file size of the workload.
  - Avoid creating small member volumes with large file workloads.
  - Use eight member volumes per node for low-end platforms. Use 16 volumes per node for higher-end platforms.
- Use two aggregates per node to maximize affinities.
  - Advanced Disk Partitioning avoids concerns with wasting drive space.
- Verify that your applications can process 64-bit file IDs.
  - This capacity is needed for more than 2 billion files.



## File size considerations

FlexGroup volumes work best with small files.

- What is a “small” file? A “large” file?  
Answer: “It depends.”
- Files do not stripe across FlexGroup member volumes.
- Large files and files that grow over time can potentially fill member volumes.
- FlexGroup members that fill up prematurely can create “out of space” issues.  
“Large” files are not necessarily a great fit, unless you size the FlexGroup volumes properly.



## Creating FlexGroup volumes

ONTAP System Manager

Search actions, objects, and pages

**ADD VOLUME**

NAME  
videolibrary

STORAGE VM  
c1\_svm3

Add as a cache for a remote volume (FlexCache)  
Simplifies file distribution, reduces WAN latency, and lowers WAN bandwidth costs.

**Storage and Optimization**

CAPACITY  
15 TiB

PERFORMANCE SERVICE LEVEL  
Performance

Not sure? [Get help selecting type](#)

OPTIMIZATION OPTIONS  
 Distribute volume data across the cluster (FlexGroup)

**Create a FlexGroup volume**


NetApp 40 © 2023 NetApp, Inc. All rights reserved.

To create a FlexGroup volume, follow the same procedure that you use to create a FlexVol volume. In the ONTAP System Manager Volumes page, click **Add Volume**. You can enter the volume name and size now or wait until after you click More Options. Scroll down in the Add Volume details page until you reach the storage and optimization section. Selecting the “Distribute volume data across the cluster” checkbox causes System Manager to create a FlexGroup volume rather than a FlexVol volume.

## Commonly used FlexGroup volume options

CLI: `volume create`

Volume option	What the volume option does
<code>-aggr-list</code>	The option specifies the names of aggregates that contain constituent volumes. Each entry in the list creates a constituent on the specified aggregate.
<code>-aggr-list-multiplier</code>	The option specifies the number of times to iterate over the aggregates that are listed with the <code>-aggr-list</code> parameter during the creation of a FlexGroup volume.
<code>-max-constituent-size</code>	The option specifies the maximum size of a constituent volume. The default value is determined by identifying the maximum FlexVol size setting on all nodes that the FlexGroup volume uses. The smallest value that is found is selected as the default for the maximum constituent size for the FlexGroup volume.

 41 © 2023 NetApp, Inc. All rights reserved.

Use the `volume create` command to create a FlexGroup volume. The options in the table are new options for the `volume create` command that are specific to FlexGroup creation. Use `-aggr-list` option to specify the aggregates that will host the FlexGroup constituent volumes. Use the `-nodes` option to use all of the aggregates on the specified nodes.


After you create a FlexGroup volume, to change the volume options or size, you must use the `volume modify` command.


Another command that has been added to ONTAP software for management of FlexGroup volumes is `volume expand`. The `volume expand` command enables you to add constituents to a FlexGroup volume. To add constituents, use the command with either the `-aggr-list` or `-aggr-list-multiplier` option. Simply specify the aggregates to which you want to add constituents and the number of constituents that you want to add to each aggregate. ONTAP software does the rest.

## Managing a FlexGroup volume

### Recommended practices


- To increase capacity, grow existing member volumes before adding new members.
  - FlexGroup volumes currently do not support shrinking or renaming volumes.
- Monitor free space and inode counts of member volumes.
  - If you have an 80% threshold, take action.
- Use nondisruptive volume move to relocate member volumes to newly added nodes.
  - Then expand the FlexGroup volume to add more members.
- Add new members in multiples. Adding single members can create hotspots.
- Consider disabling change or notify on CIFS shares if they are unneeded.

 42 © 2023 NetApp, Inc. All rights reserved.





# Addendum FlexCache volumes

 43 © 2023 NetApp, Inc. All rights reserved.

## FlexCache software

Topology	Configuration	Systems	Licensing
<ul style="list-style-type: none"><li>• Intracluster caching</li><li>• Within a storage VM</li><li>• Across storage VMs</li><li>• Cross-cluster caching</li></ul>	<ul style="list-style-type: none"><li>• Write-around caches</li><li>• Support for up to 10 caches per origin volume</li><li>• Protocol: NFSv3, NFSv4, and SMB</li></ul> <p>Cache volumes are FlexGroup volumes.</p>	<ul style="list-style-type: none"><li>• NetApp FAS</li><li>• NetApp AFF</li><li>• NetApp ONTAP Select</li><li>• NetApp Cloud Volumes ONTAP</li></ul>	<ul style="list-style-type: none"><li>• No-cost capacity-based licensing</li><li>• Based on cache-volume capacity</li><li>• Aggregated at a cluster level</li></ul>



## Creating FlexCache volumes

**ONTAP System Manager** Search actions, objects, and pages

**DASHBOARD**

**INSIGHTS**

**STORAGE** ^

Overview

Volumes

LUNs

Consistency Groups

Shares

Quotas

Storage VMs

Tiers

**NETWORK** v

**EVENTS & JOBS** v

**PROTECTION** v

**HOSTS** v

**CLUSTER** v

### Add Volume

NAME  
videolibrary\_mirror\_1

STORAGE VM  
c1\_svm3

Add as a cache for a remote volume (FlexCache)  
Simplifies file distribution, reduces WAN latency, and lowers WAN bandwidth costs.

CLUSTER  
cluster1 Refresh

STORAGE VM  
c1\_svm3

VOLUME  
videolibrary

FOLDERS TO PREPOPULATE ⓘ

+ Add

**Create a FlexCache volume**

NetApp 45 © 2023 NetApp, Inc. All rights reserved.


You can create FlexCache volumes from the Volumes page like you create FlexGroup volumes.

1. Click **Add**, and then in the Add Volume window, click **More options**.
2. Enter a volume name and select the storage VM that will own the FlexCache volume.
3. Select the Add as a cache for a remote volume checkbox.
4. Specify the size of the FlexCache volume. The default size is 10% of the origin volume size.
5. Click **Save**.

**Note:** Cluster and storage VM peer relationships must be established before you configure a destination FlexCache volume on a different storage VM or cluster than the origin.



## **Addendum Autobalance aggregate**

 46 © 2023 NetApp, Inc. All rights reserved.

## Autobalance aggregate syntax

Enable autobalancing and modify the thresholds with the following commands:

Enable the autobalance feature for the cluster:

```
::> autobalance aggregate config modify -is-enabled true
```

Modify the threshold when an aggregate is considered unbalanced:

```
::> autobalance aggregate config modify  
-aggregate-unbalanced-threshold-percent <integer>
```

Modify the threshold when an aggregate is considered balanced:


```
autobalance aggregate config modify -aggregate-available-threshold-percent
```

You can control the threshold at which point the aggregate is considered unbalanced.

The Auto Balance Aggregate feature attempts to move volumes from an unbalanced aggregate until it is lower than the percentage that is specified by the `aggregate-available-threshold-percent` option.

# Module 7

## Data access

 1 © 2023 NetApp, Inc. All rights reserved.

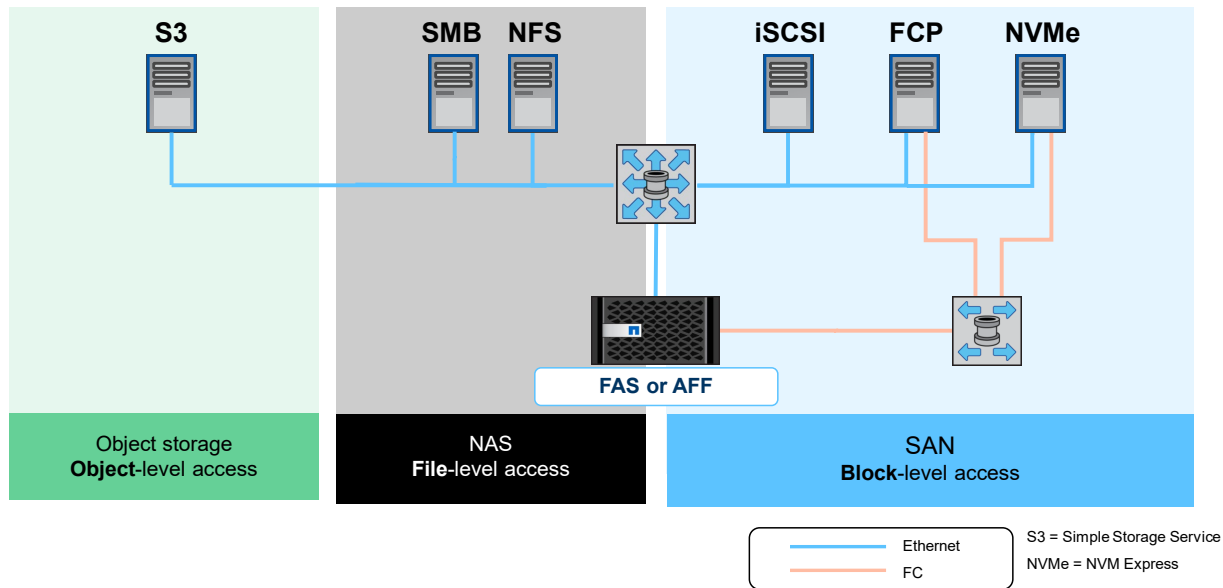
## About this module

This module focuses on enabling you to do the following:

- Use NAS protocols to access data
- Use SAN protocols to access data
- Use object protocols to access data

# Unified storage

Review




NetApp 3 © 2023 NetApp, Inc. All rights reserved.

NAS is a file-based storage system that uses NFS and SMB protocols to make data available over the network. CIFS is a dialect of SMB.


A SAN is a block-based storage system that uses the FC, iSCSI, or NVMe protocols to make data available over the network.

A storage system that can manage both NAS and SAN data is referred to as *unified storage*. NetApp ONTAP 9.7 software expanded the term *unified storage* to include object-based storage.

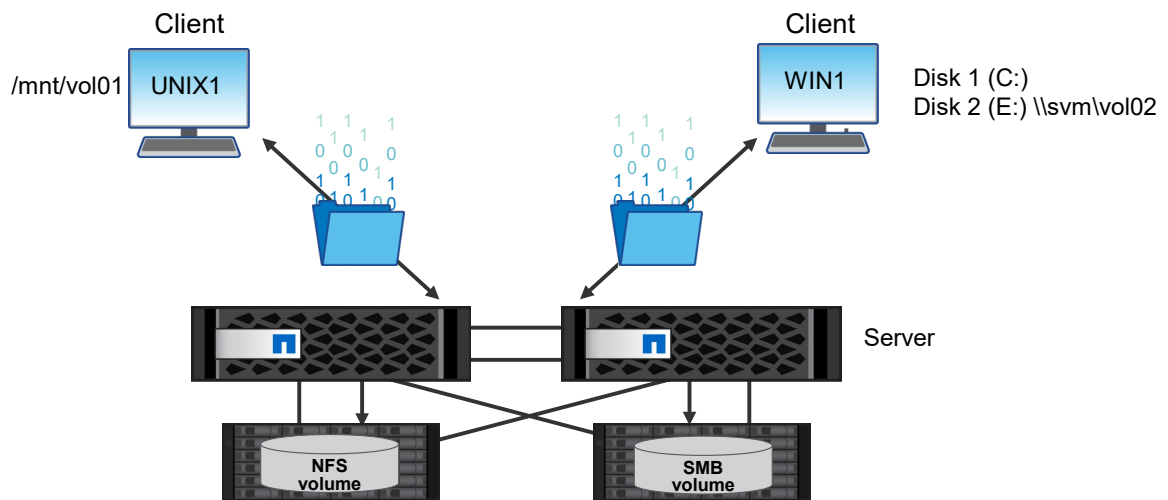


# Lesson 1

## Use NAS protocols to access data

 4 © 2023 NetApp, Inc. All rights reserved.

## NAS file system



NetApp 5 © 2023 NetApp, Inc. All rights reserved.

NAS is a distributed file system that enables users to access resources, such as volumes, on a remote storage system as if the resources were on a local computer system.

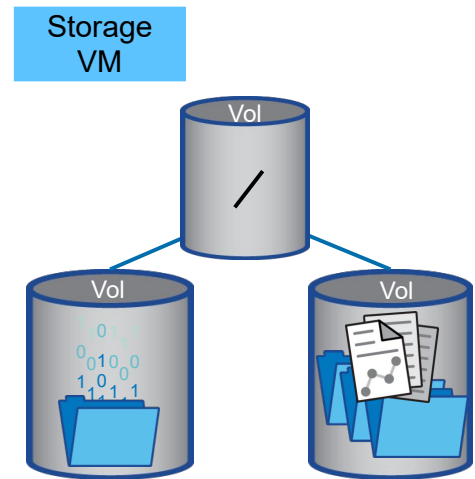
NAS provides services through a client-server relationship. Storage systems that make file systems and other resources available for remote access are called *servers*. The server is set up with a network address and provides file-based data storage to other computers, called *clients*, that use the server resources.

ONTAP software supports the NFS and SMB protocols.



## Storage system resources

- FlexVol volume
  - Data container to manage data in a storage VM (storage virtual machine, also known as SVM)
  - Exportable by mounting to a namespace junction
- Qtree
  - Volume partition that is created on a storage system
  - Exportable by mounting to a namespace junction
- Directory
  - Volume partition that is created on the NAS client
  - Not exportable



**NetApp** 6 © 2023 NetApp, Inc. All rights reserved.

With NAS protocols, you need to create file systems and other resources that are available to clients through either NFS or SMB.

FlexVol volumes are data containers that enable you to partition and manage your data. In a NAS environment, volumes contain file systems. The first resource to create is the volume.

In ONTAP software, the volume is associated with a storage VM (storage virtual machine, also known as SVM). The storage VM is a virtual management entity within which you create a namespace. Volumes are joined to the namespace through junctions. The junctions are exported.

Qtrees enable you to partition FlexVol volumes into smaller segments, which you can manage individually. ONTAP software creates a default qtree, called `qtree0`, for each volume. If you do not create and put data in another qtree, all the data resides in `qtree0`. Qtrees enable you to partition data without incurring the overhead that is associated with creating another FlexVol volume. Like volumes, qtrees can be exported and mounted. You might create qtrees to organize data or to manage one or more of the following factors: quota limitations, security style, or workload performance.

You can also create a directory or a file on the client in a FlexVol volume to use as a resource to export or share. A qtree is a partition on the storage system. A directory is a partition that the client creates within a FlexVol volume.

## Namespace and junction paths

- Create a projects volume under the storage VM root:

```
::> volume create -vserver svm4  
-aggregate sas_data_23 -volume projects  
-size 5GB -junction-path /projects
```

– THEN –


- Create a second volume named thesis:

```
::> volume create -vserver svm4  
-aggregate sas_data_18 -volume thesis  
-size 10GB
```

- Mount the second volume under /projects:

```
::> volume mount -vserver svm4 -volume thesis  
-junction-path /projects/thesis
```

It is best practice to mount data volumes directly to the root volume.



thesis All Volumes Edit More

Overview Snapshot Copies SnapMirror (Local or Remote)  
Back Up to Cloud File System Security Quota Reports

STATUS  
Online

STYLE  
FlexVol

MOUNT PATH  
/projects/thesis

TIERING POLICY  
None

STORAGE VM  
svm4

Capacity

USED	AVAILABLE	SIZE
260 KiB	150 MiB	158 MiB

260 KiB logical used

SNAPSHOT CAPACITY  
7.89 MiB Available 0 Bytes Used

NetApp 7 © 2023 NetApp, Inc. All rights reserved.

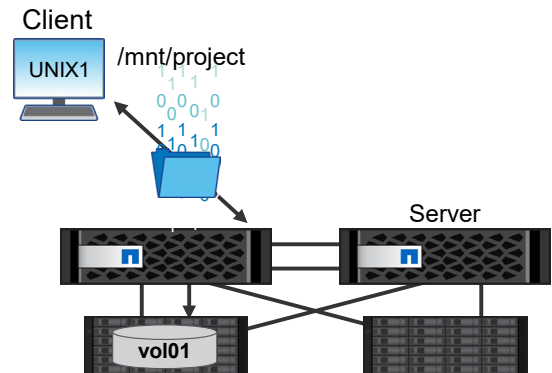
Volume junctions are a way to join individual volumes into one logical namespace. Volume junctions are transparent to CIFS and NFS clients. When NAS clients access data by traversing a junction, the junction appears to be an ordinary directory. A junction is formed when a volume is mounted to a mount point below the root and is used to create a file-system tree. The top of a file-system tree is always the root volume, which is represented by a slash mark (/). A junction points from a directory in one volume to the root directory of another volume.

You must mount a volume at a junction point in the namespace to enable NAS client access to the contained data. Specifying a junction point is optional when you create a volume. However, you cannot export data in the volume or create a share until you mount the volume to a junction point in the namespace. You can add new volumes to the namespace at any time by mounting the volumes to a junction point.

Nested mounts, like the ones in this example, are not recommended because they create a dependency. If the projects volume goes offline, the thesis volume becomes inaccessible. NetApp recommends that you mount all NAS volumes to the storage VM root volume to create a flat namespace without dependencies.

# NFS

- vol01 is *exported* to UNIX1 with read/write permission.
- UNIX1 *mounts* vol01 to /mnt/project with read/write permission.



NetApp 8 © 2023 NetApp, Inc. All rights reserved.

NFS provides services through a client-server relationship.

- NFS servers make file systems and other resources available for remote access by *exporting* them.
- NFS clients access an exported file system by *mounting* it to the local directory hierarchy.

When a client mounts a file system that a server exports, users on the client computer can view and interact with the mounted file systems on the server. This interaction must be within the limits of the granted permissions.

## NFS implementation steps

1. Verify or add the NFS protocol license.
2. Enable NFS functionality on the storage VM.
3. Create or identify the necessary resources.
4. Export the available resources.
5. Configure NFS authentication.
6. Authorize the user.
7. Mount the exported resources.



**NetApp** 9 © 2023 NetApp, Inc. All rights reserved.

The basic process for implementing the NFS protocol between a UNIX client and an ONTAP storage system involves several steps:

- License NFS, and then enable NFS functionality on the storage VM.
- You need resources to export, so create volumes, qtrees, and data LIFs.
- Determine which clients have which type of access to the resources. You need a way to authenticate client access and authorize users with appropriate permissions, including read-only or read/write.
- After the client has been granted access to the exported resource, the client mounts the resource and grants access to the users.

# Storage VM creation: NFS

## Storage VM basic details

The screenshot shows the NetApp ONTAP System Manager interface. On the left is a navigation sidebar with categories: DASHBOARD, INSIGHTS, STORAGE (expanded), NETWORK, EVENTS & JOBS, and PROTECTION. The 'STORAGE' section is active, showing a 'Storage VMs' table with columns for Name, State, Subtype, and Configured... Three storage VMs are listed: svm1, svm2, and svm3, all in a 'running' state with a 'default' subtype. A green arrow labeled 'Create' points to the '+ Add' button above the table. To the right, the 'Add Storage VM' dialog is open. It has a 'STORAGE VM NAME' field containing 'svm5'. Under the 'Access Protocol' section, three tabs are visible: 'SMB/CIFS, NFS, S3' (selected), 'iSCSI', and 'FC'. The 'SMB/CIFS, NFS, S3' tab contains several checkboxes: 'Enable SMB/CIFS' (unchecked), 'Enable NFS' (checked), 'Allow NFS client access' (unchecked), and 'Enable S3' (unchecked). A warning message is displayed below the 'Allow NFS client access' checkbox: 'Add at least one rule to allow NFS clients to access volumes in this storage VM.' The 'EXPORT POLICY' is set to 'Default'.

NetApp 10 © 2023 NetApp, Inc. All rights reserved.

Use this page in NetApp ONTAP System Manager (formerly OnCommand System Manager) to create a storage VM that uses the NFS protocol.

The feature licenses that are installed on the cluster determine which tabs (NAS or SAN) and protocol checkboxes you see in the Access Protocol section. This example shows an ONTAP system with only NAS protocol licenses installed.

## Storage VM creation: NFS

Enable NFS client access

The screenshot displays the NetApp storage VM configuration interface. On the left, the 'Access Protocol' section shows 'SMB/CIFS, NFS, S3' selected. Underneath, 'Enable NFS' is checked, and 'Allow NFS client access' is also checked. A warning icon indicates that at least one rule must be added to allow NFS clients to access volumes. Below this, the 'EXPORT POLICY' is set to 'Default', and the 'RULES' section is currently empty with a '+ Add' button. On the right, a 'New Rule' dialog box is open. It has a 'CLIENT SPECIFICATION' field with '192.168.0.0/16'. Under 'ACCESS PROTOCOLS', 'SMB/CIFS', 'FlexCache', 'NFS', 'NFSv3', and 'NFSv4' are all checked. The 'ACCESS DETAILS' table is as follows:

Type	Read-only Access	Read/Write Access	Superuser Access
All	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
All (As anonymous user)	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
UNIX	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Kerberos 5	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Kerberos 5i	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Kerberos 5p	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
NTLM	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

At the bottom of the dialog are 'Cancel' and 'Save' buttons. A green callout box on the left says 'Allow NFS client access.' and another says 'Add an export rule.'

Here you see the next set of steps to configure the NFS protocol within the storage VM. By default, access to the new storage VM is blocked. Select the **Allow NFS client access** checkbox to enable access.

A new default export policy is created for each storage VM. Initially, this policy does not contain any rules, so client hosts are denied access. Add rules to the policy to control access to the storage VM or assign a different export policy to the storage VM.

To add a new export policy rule, first enter the client specification to which the rule applies. You can select client hosts by name, address, netgroup, or domain. Then, choose the access protocols that clients can use to access the storage VM. You must choose which methods are attempted to authenticate the user and the level of access to grant if the authentication method succeeds. Finally, you must decide whether users with superuser privileges on the client hosts should be granted superuser privileges when accessing the storage VM.

NFS export policies can be applied to storage VMs, volumes, and qtrees. An export policy that is assigned to an individual volume overrides the export policy that is assigned to a storage VM.

## Storage VM creation: NFS

Configure network interfaces

The screenshot displays the ONTAP System Manager interface for configuring network interfaces. The left sidebar shows a navigation menu with 'STORAGE' selected, and 'Storage VMs' highlighted. The main content area is titled 'NETWORK INTERFACE' and includes the instruction: 'Use multiple network interfaces when client traffic is high.' Two interface configurations are shown:

- cluster2-01:** IP ADDRESS: 192.168.0.13, SUBNET MASK: 24, GATEWAY: Add optional gateway, BROADCAST DOMAIN AND PORT: Default.
- cluster2-02:** IP ADDRESS: 192.168.0.14, PORT: Automatical...

A checkbox is checked with the text: 'Use the same subnet mask and gateway for all of the following interfaces'.

NetApp 12 © 2023 NetApp, Inc. All rights reserved.

When enabling the NFS protocol within the storage VM, you need to configure some logical network interfaces. Enter the IP address and network mask for each logical network interface through which the storage VM is accessed by using NFS. Add LIFs to distribute the traffic across multiple network ports and cluster nodes.

## Storage VM creation: NFS

### Storage VM administrator details

Optionally, create a storage VM administrator account.

Optionally, create a storage VM management LIF.

NetApp 13 © 2023 NetApp, Inc. All rights reserved.

To create a storage VM administrator account, select the **Manage administrator account** checkbox, and then enter a password.

Optionally, you can create a LIF that is dedicated to the management of the storage VM. Generally, this step is necessary only for storage VMs that enable only SAN protocols. The system management functions of the storage VM can be accessed through any NAS data LIF.



## NFS client mounts

Use the UNIX `mount` command on the client to mount an exported NFS resource from the storage system.

```
unix1# mkdir /mnt/project1  
unix1# mount <SVM LIF IP>:/project/proj1 /mnt/project1
```

To enable an NFS client, mount a remote file system after NFS starts. For a temporary mount, use the `mount` command to attach the resource to a directory. For a persistent mount, add an entry to the file system table file.


Usually, only a privileged user can mount file systems with NFS. However, if the `user` option is set in `/etc/fstab`, you can enable users to mount and unmount selected file systems by using the `mount` and `umount` commands. The setting can reduce traffic by mounting file systems only when necessary. To enable user mounting, create an entry in `/etc/fstab` for each file system to be mounted.

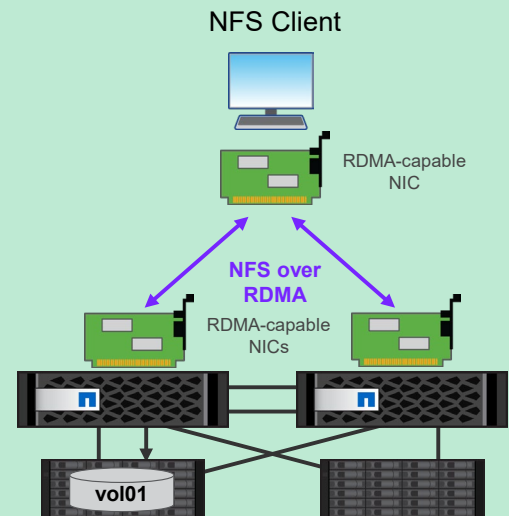
## NFS over RDMA

RoCE (RDMA over Converged Ethernet)

NFS over RDMA is for customers with latency sensitive or high-bandwidth workloads.

- Storage systems must be running ONTAP 9.10.1 or later software.
- Storage system controllers must have RDMA support (currently NetApp AFF A400, AFF A700, and AFF A800).
- Storage system controllers must include network interfaces that are RDMA capable (for example, Mellanox CX-5 or CX-6).
- Data LIFs must be configured to support RDMA.
- Clients must be using Mellanox RDMA-capable network interface cards (NICs) and Mellanox OFED (MOFED) network software.

 15 © 2023 NetApp, Inc. All rights reserved.



NFS over RDMA uses RDMA adapters, enabling data to be copied directly between storage system memory and host system memory, circumventing CPU interruptions and overhead. NFS over RDMA enables the use of NVIDIA GPUDirect Storage for GPU-accelerated workloads on hosts with supported NVIDIA GPUs.

NFS over RDMA configurations are designed for customers with latency sensitive or high-bandwidth workloads such as machine learning and analytics. NVIDIA has extended NFS over RDMA to enable GPU Direct Storage (GDS). GDS further accelerates GPU-enabled workloads by bypassing the CPU and main memory altogether, using RDMA to transfer data between the storage system and GPU memory directly. NFS client hosts must include the Mellanox OpenFabrics Enterprise Distribution (OFED) drivers for high-speed network cards to enable RDMA networking.

NFS over RDMA is supported beginning with ONTAP 9.10.1 software. NFS over RDMA configurations are only supported for the NFSv4.0 protocol when used with the Mellanox CX-5 or CX-6 adapter. This adapter provides support for RDMA using version 2 of the RoCE protocol. GDS is only supported using NVIDIA Tesla-family and Ampere-family GPUs with Mellanox NIC cards and MOFED software. NFS over RDMA support is limited to only node-local traffic. Standard FlexVol volumes or FlexGroup volumes in which all constituents are on the same node are supported and must be accessed from a LIF on the same node.

With ONTAP 9.12.1 software, System Manager supports network interface configuration for NFS over RDMA and identifies RoCE capable ports.



## Complete an exercise

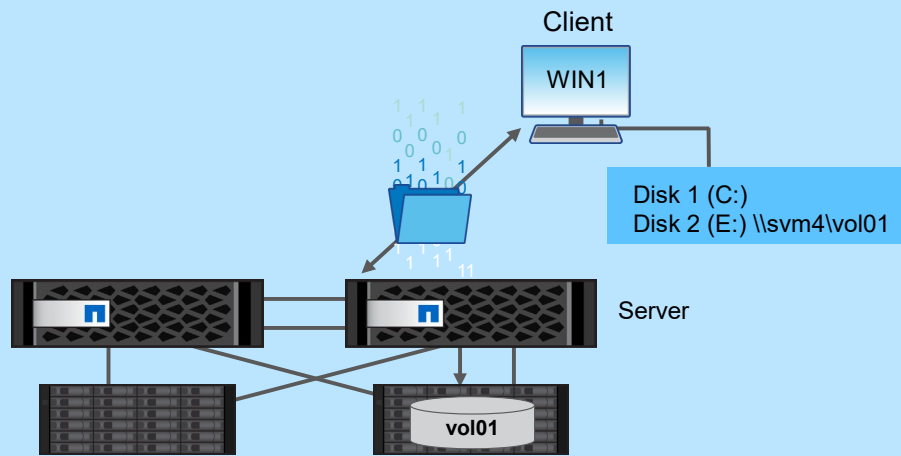
Module 7  
Data access

### Configuring the NFS Protocol in a Storage VM

- Access your lab equipment.
- Open your Exercise Guide to Module 7.
- Complete Exercise 1.
- Share your results.

This exercise requires approximately  
**20 minutes.**

## SMB



NetApp 17 © 2023 NetApp, Inc. All rights reserved.

SMB is an application-layer network file-sharing protocol that Microsoft Windows software uses. SMB enables users or applications to access, read, and write to files on remote computers as they would on a local computer. For the purposes of this course, the terms SMB and CIFS are used interchangeably (although the definitions of the two terms are not strictly the same).

A user or application can send network requests to read and write to files on remote computers.

SMB provides access to files and directories that are stored on the remote computer, through sharing resources. The rules of network protocols such as IPv4 and IPv6 control the network read and write process, which is also called network I/O.

## SMB implementation steps

1. Verify or add the CIFS protocol license.
2. Enable SMB functionality on the storage VM.
3. Create or identify the necessary resources.
4. Share the available resources.
5. Configure SMB authentication.
6. Authorize the user.
7. Map the shared resources.



**NetApp** 18 © 2023 NetApp, Inc. All rights reserved.

The basic process for implementing the SMB protocol between a Windows client and an ONTAP storage system involves several steps:

- License CIFS, and then enable SMB functionality on the storage VM.
- Create volumes, qtrees, and data LIFs.
- Determine which clients have which type of access to the resources. You need a way to authenticate client access and authorize users with appropriate permissions, including read-only or read/write.
- After the client is granted access to the shared resource, the client maps the resource and grants access to the users.

# Storage VM creation: SMB

## Storage VM basic details

The screenshot displays the ONTAP System Manager interface. On the left, a navigation sidebar includes sections for DASHBOARD, INSIGHTS, STORAGE (with sub-items like Overview, Volumes, LUNs, Consistency Groups, Shares, Buckets, Qtrees, Quotas, Storage VMs), NETWORK, and EVENTS & JOBS. The main area shows a table of Storage VMs with columns for Name, State, Subtype, Configured Pr..., and IPspace. A green arrow labeled 'Create' points to the '+ Add' button. A modal dialog titled 'Add Storage VM' is open, showing a text field for 'STORAGE VM NAME' containing 'svm\_SMB'. Under the 'Access Protocol' section, the 'SMB/CIFS, NFS, S3' tab is selected, and the 'Enable SMB/CIFS' checkbox is checked. A green arrow labeled 'Protocols' points to this section.

Name	State	Subtype	Configured Pr...	IPspace
svm1	running	default	NFS	Default
svm2	running	default	NFS, SMB/CIFS	Default
svm3	running	default	NFS, SMB/CIFS	Default
svm4	running	default	NFS	Default
svm5	running	default	S3	Default

You create the storage VM in System Manager in much the same way as for NFS. In this example, the major difference is the selection of the SMB protocol instead of NFS. In environments with both SMB and NFS clients, selecting both protocols is a typical way to enable clients to use their native protocol to access the resources.

## Storage VM creation: SMB

Configure the CIFS protocol

**Information to create a machine record in Active Directory**

**DNS and NTP settings from the Admin storage VM**

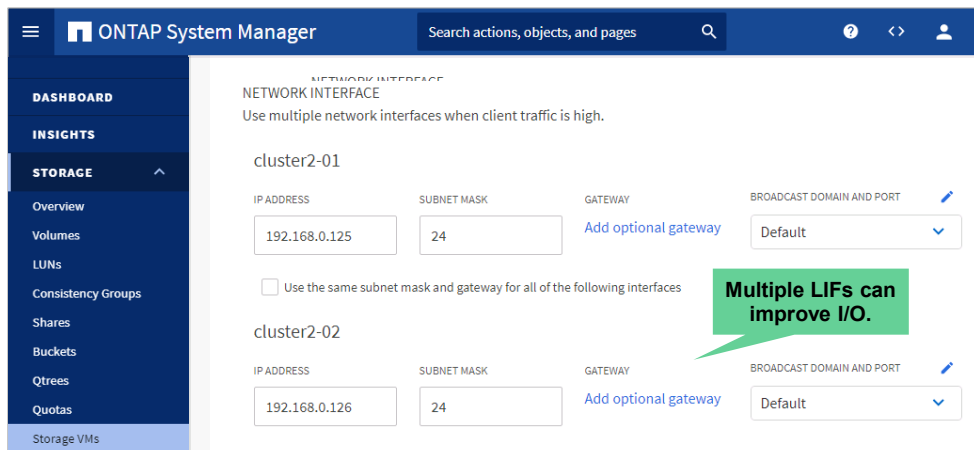
NetApp 20 © 2023 NetApp, Inc. All rights reserved.

When enabling the SBM protocol on the storage VM, you must also create a server entry in Active Directory. Provide the Active Directory domain name and login credentials of a user who can update the domain. Enter a unique SMB server name for the storage VM.

By default, a storage VM uses the same DNS domain and server as the ONTAP cluster. You can change the DNS information for the storage VM if necessary.

## Storage VM creation: SMB

Configure network interfaces



The screenshot shows the ONTAP System Manager interface for configuring network interfaces. The left sidebar contains navigation options: DASHBOARD, INSIGHTS, STORAGE (selected), Overview, Volumes, LUNs, Consistency Groups, Shares, Buckets, Qtrees, Quotas, and Storage VMs. The main content area is titled 'NETWORK INTERFACE' and includes the instruction 'Use multiple network interfaces when client traffic is high.' Below this, there are two sections for 'cluster2-01' and 'cluster2-02'. Each section has input fields for 'IP ADDRESS', 'SUBNET MASK', 'GATEWAY', and 'BROADCAST DOMAIN AND PORT'. The 'IP ADDRESS' field for cluster2-01 is '192.168.0.125' and for cluster2-02 is '192.168.0.126'. The 'SUBNET MASK' field for both is '24'. The 'GATEWAY' field has a link 'Add optional gateway'. The 'BROADCAST DOMAIN AND PORT' field has a dropdown menu set to 'Default'. A green callout box with a speech bubble points to the configuration area, containing the text 'Multiple LIFs can improve I/O.' There is also a checkbox labeled 'Use the same subnet mask and gateway for all of the following interfaces' which is currently unchecked.

NetApp 21 © 2023 NetApp, Inc. All rights reserved.

When enabling the SMB protocol, you need to configure some logical network interfaces. Enter the IP address and network mask for each logical network interface through which the storage VM is accessed by using SMB. Both SMB and NFS protocol traffic can share a data LIF.



## Storage VM creation: SMB

### Storage VM administrator details

Optionally, create a storage VM administrator account.

Optionally, create a storage VM management LIF.

In an exercise for this module, you create a storage VM to serve both NFS and SMB.

NetApp 22 © 2023 NetApp, Inc. All rights reserved.

If you plan to delegate management of the new storage VM, you can create a separate storage VM administrator account. The storage VM can be administered through any NAS data LIF that belongs to the storage VM. Therefore, you do not need to create a management LIF for the storage VM.

## Create an SMB share

The screenshot displays the NetApp ONTAP System Manager interface. On the left is a navigation sidebar with sections for DASHBOARD, INSIGHTS, STORAGE, and NETWORK. The STORAGE section is expanded, showing options like Overview, Volumes, LUNs, Consistency Groups, Shares, Buckets, Qtrees, Quotas, Storage VMs, and Tiers. The 'Shares' page is active, showing a table of existing shares. A green arrow points to the '+ Add' button, which is labeled 'Create'. The table lists shares with columns for Name, Storage VM, and Path. The 'Add Share' dialog box is open on the right, showing the following fields and options:

- SHARE NAME: Input field containing 'SMB\_share'.
- STORAGE VM: Dropdown menu showing 'svm\_SMB'.
- FOLDER NAME: Input field containing '/', with a 'Browse' button.
- DESCRIPTION: Text area containing 'Entire SVM namespace'.
- ACCESS PERMISSION: Table with columns 'User/Group', 'User Type', and 'Access Permission'. It shows 'Everyone' with 'Windows' user type and 'Full Control' permission.
- Buttons: '+ Add', 'Save', and 'Cancel'.
- Checkboxes: 'Enable Continuous Availability' (unchecked) and 'Encrypt data while accessing this share' (unchecked).

NetApp 23 © 2023 NetApp, Inc. All rights reserved.

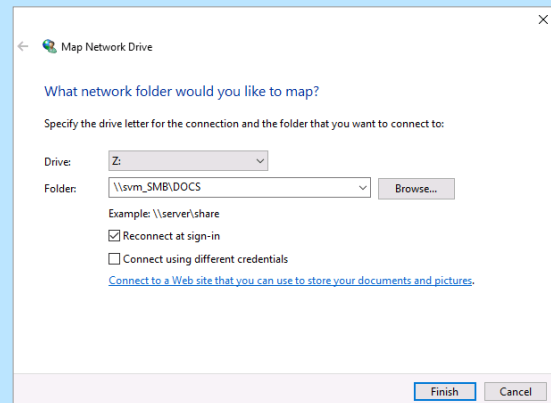
Grant access to your SMB users by creating a share.

1. Enter a share name.
2. Select an SMB-enabled storage VM to host the share.
3. Enter a slash (/) for the folder name to share the entire namespace of the storage VM or click **Browse** to select an individual volume or qtree.
4. Click **Add**.
5. Grant share-level access permissions.

## Mapping a share to a client

- CLI:

```
C:\> net view \\svm_SMB
C:\> net use e: \\svm_SMB\DOCS
/user:marketing\jdoe
```
- UI:
  - Use Microsoft Windows File Explorer.
  - Map a network drive: \\svm\_SMB\DOCS



NetApp 24 © 2023 NetApp, Inc. All rights reserved.

The `net view` command displays a list of computers with shared resources that are available on the specified computer.

To use the `net view` command, complete these steps:

1. Click the **Start** button.
2. Point to **Programs**.
3. Click the **MS-DOS** prompt.
4. At the command prompt, type `net view \\<computer_name>`, where `<computer_name>` is the name of a computer with resources that you want to view. If the computer does not have an address entry in DNS, enter its IP address instead.

You can connect or disconnect a computer from a shared resource, or you can display information about computer connections. The command also controls persistent net connections. Used without parameters, the `net use` command retrieves a list of network connections.

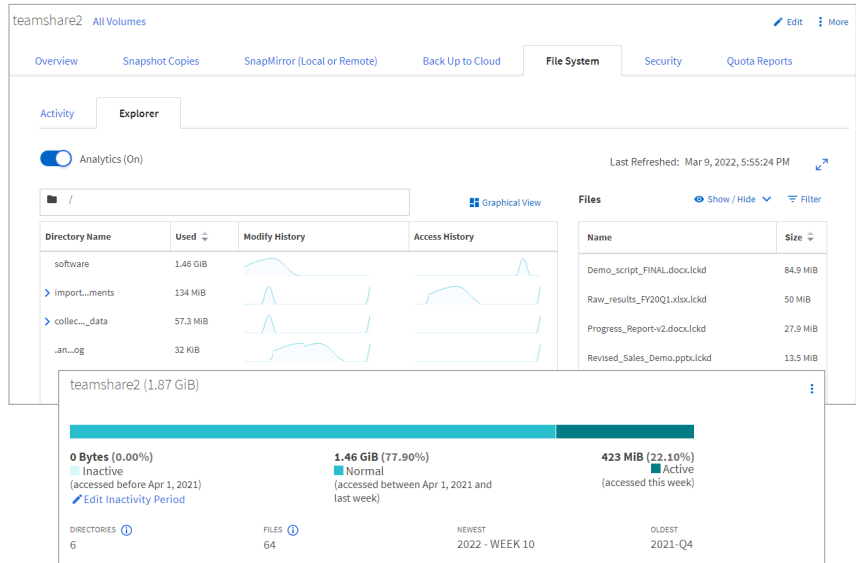
You can also use Windows to map a share to a client.

# File System Analytics

Explore a NAS volume

## Visualization through NetApp ONTAP System Manager

- Collects granular, hierarchical subdirectory metadata to provide visibility into the following:
- Capacity usage and trends
- File and directory counts
- File activity and age histogram
- APIs available for application integration



NetApp 25 © 2023 NetApp, Inc. All rights reserved.

One challenge that storage administrators face when managing NAS is data visibility. A volume can have thousands, millions, or billions of files and folders. Using standard NAS queries such as find, du, and ls to find useful information about your data, such as file size, age, and file count, can be a headache. With the ONTAP File System Analytics feature, you can have near-instant, up-to-date access to the file system information that you want, simply by opening System Manager.

You can enable this feature per volume. A low-impact background process runs on a schedule to keep information about the directory tree available. When you need to find the size of a file, go to System Manager and look in the directory tree. You can also quickly discover which data is cold and can benefit from tiering to cloud or S3.

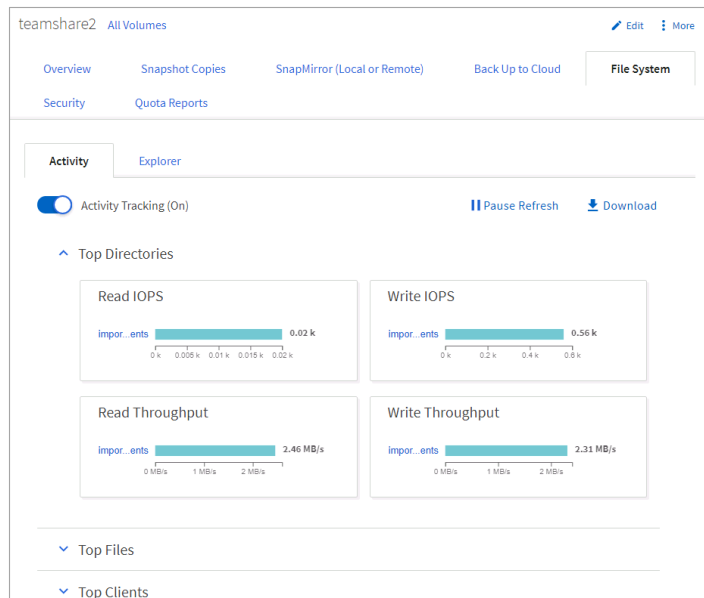
After you enable file analytics and after ONTAP scans the file system hierarchies, you can view a history of activity. View directory activity in list view, or switch to the graphical view to see which directories contain the most inactive data.

If you want to get this information without accessing the UI, you can use RESTful APIs to query ONTAP for the information.

# File System Analytics

## Activity Tracking

- Can be enabled on volumes to identify hotspots
- Monitors volume activity and reports the top 25 most active directories, files, storage VMs, client hosts, and users
- Displays performance metrics for the busiest objects in a volume
- Displays the largest directories in a volume



NetApp 26 © 2023 NetApp, Inc. All rights reserved.

Beginning with ONTAP 9.10.1 software, File System Analytics includes an Activity Tracking feature that enables you to identify hot objects in a FlexVol volume or FlexGroup volume.

File System Analytics monitors volume activity and identifies the busiest files, directories, storage VMs, client hosts, and users.

For each object, Activity Tracking displays read IOPs, write IOPs, read throughputs, and write throughputs.


Activity Tracking refreshes every 10 to 15 seconds to display hot spots that are seen in the system over the previous 5-second interval.

You can stop the refreshes with the Pause Refresh button to give yourself more time to interpret the information.

ONTAP 9.12.1 software extends File System Analytics to report the top directories by size.

Activity data can be downloaded in a CSV format that displays all the point-in-time data that is captured for the selected volume.

## Additional NAS learning

 27 © 2023 NetApp, Inc. All rights reserved.

- Where can you learn about advanced topics like the following?
  - Protocol versions and features
  - Export policies and rules
  - Shares
  - Authentication
  - Permissions
  - Using multiple protocols
  - Managing scalable NAS containers
- ONTAP NAS Fundamentals (online course)
- ONTAP NFS Administration (virtual/instructor-led course)
- ONTAP SMB Administration (virtual/instructor-led course)

- *ONTAP NAS Fundamentals* (online course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours00000000022332](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours00000000022332)
- *ONTAP NFS Administration* (virtual/instructor-led course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours00000000015071](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours00000000015071)
- *ONTAP SMB Administration* (virtual/instructor-led course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours00000000015070](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours00000000015070)

**Note:** Fundamentals and administration courses are updated at least once a year. The URLs might change, but the names rarely do. Use the search engine in the NetApp LearningCenter to find the most recent version of the courses and offerings.



## Complete an exercise

Module 7  
Data access

### Configuring the SMB Protocol in a Storage VM

- Access your lab equipment.
- Open your Exercise Guide to Module 7.
- Complete Exercise 2.
- Share your results.

This exercise requires approximately  
**15 minutes.**



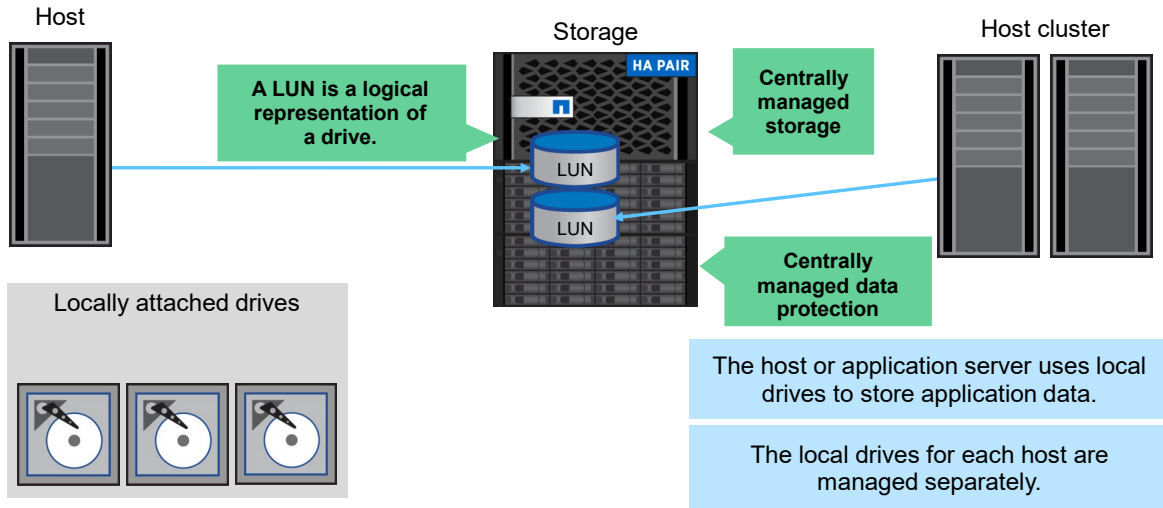
## Lesson 2

# Use SAN protocols to access data

 29 © 2023 NetApp, Inc. All rights reserved.



# SAN



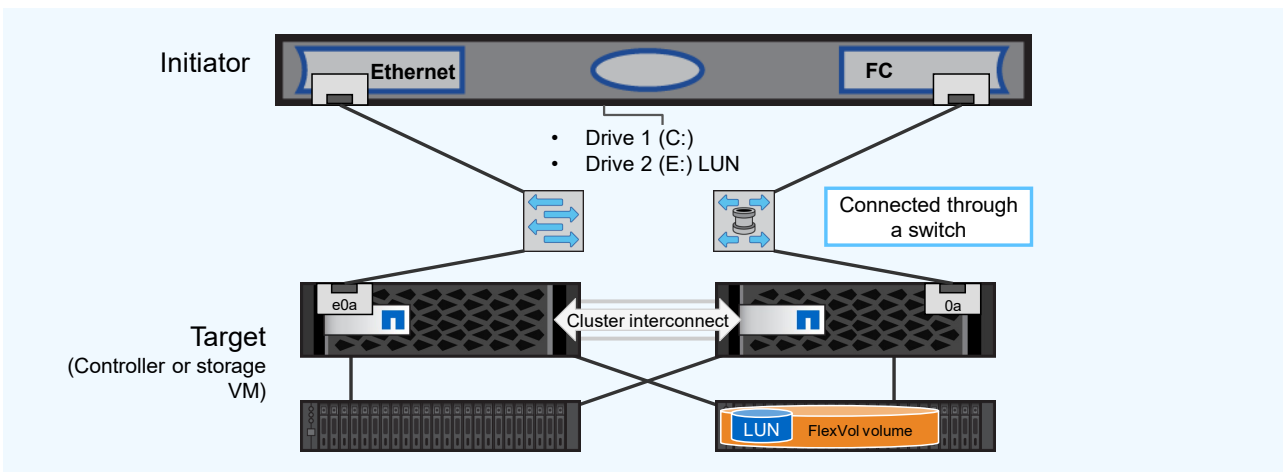
HA = high-availability

In an application server environment, locally attached drives, also called direct-attached storage (DAS), are separately managed resources. In an environment with more than one application server, each server storage resource also needs to be managed separately.

A SAN provides access to a LUN, which represents a SCSI-attached drive. The host operating system partitions, formats, writes to, and reads from the LUN as if the LUN were any other locally attached drive. The advantages of using SAN storage include support for clustered hosts, where shared drives are required, and centrally managed resources. In the example, if the administrator did not use a SAN, the administrator would need to manage separate resources for each application server and host cluster. As well as enabling centrally managed resources, SAN uses ONTAP Snapshot copy technology to enable centrally managed data protection and enterprise-grade data management.

## Connecting initiator to target

How can you connect an initiator to a target?



NetApp 31 © 2023 NetApp, Inc. All rights reserved.

ONTAP software supports the iSCSI, FC, FCoE, NVM Express over Fibre Channel (NVMe/FC), and NVM Express over TCP (NVMe/TCP) SAN protocols.

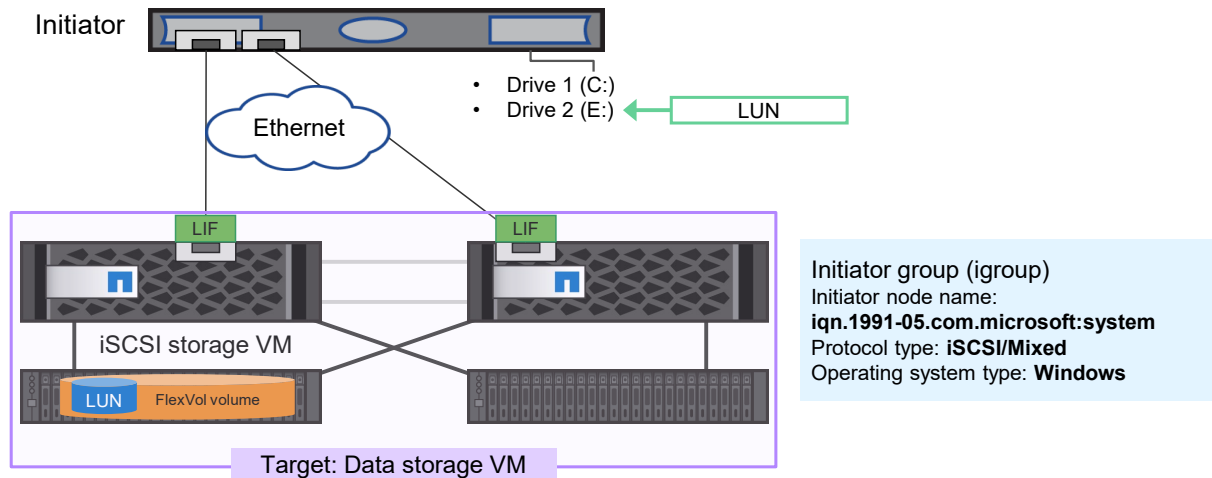
Data is communicated over ports and LIFs.

- In an Ethernet SAN, the data is communicated over Ethernet ports.
- In an FC SAN and NVMe/FC SAN, the data is communicated over FC ports.
- For FCoE, the initiator has a converged network adapter (CNA), and the target has a unified target adapter (UTA).
- SAN data LIFs do not migrate or fail over the way that NAS does. However, the LIFs can be moved to another node or port in the storage VM.

NetApp recommends the following practices:

- Use at least one LIF per node, per storage VM, per network.
- Use redundant connections to connect the initiator to the target.
- Use redundantly configured switched networks to provide resiliency if a cable, port, or switch fails.

## iSCSI architecture



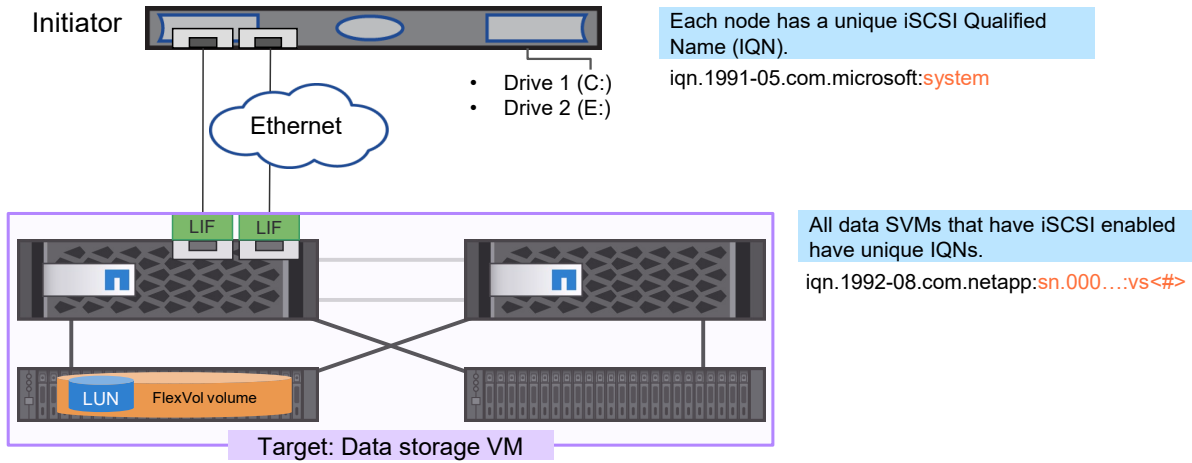
NetApp 32 © 2023 NetApp, Inc. All rights reserved.

Initiator groups (igroups) are tables of FC protocol host worldwide port names (WWPNs) or iSCSI host node names. You can define igroups and map the igroups to LUNs to control which initiators have access to LUNs. In the example, the initiator uses the iSCSI protocol to communicate with the target.

Typically, you want all the host initiator ports or software initiators to have access to a LUN. The example shows a single host. The iSCSI Software Initiator iSCSI Qualified Name (IQN) is used to identify the host.

An igroup can have multiple initiators, and multiple igroups can have the same initiator. However, you cannot map a LUN to multiple igroups that have the same initiator. An initiator cannot be a member of igroups of differing operating system types. In the example, the initiator runs Windows.

## iSCSI node names

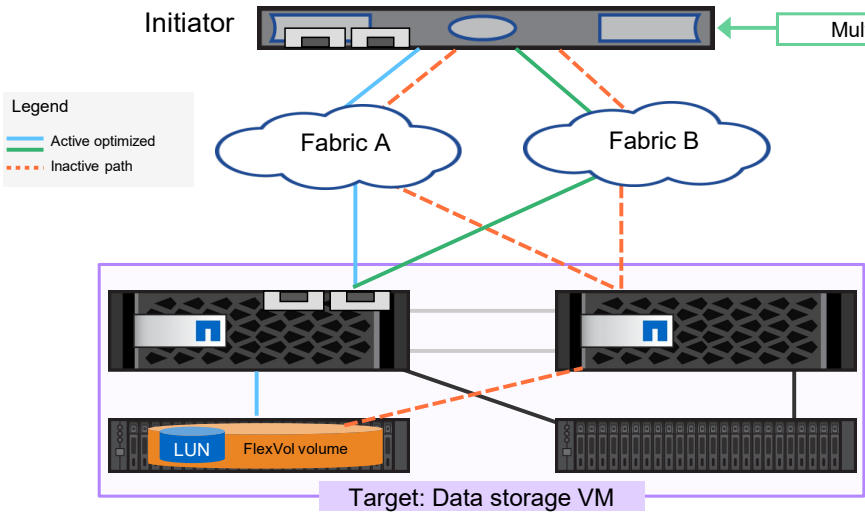


NetApp 33 © 2023 NetApp, Inc. All rights reserved.

Each storage VM is a separate target. Each storage VM is assigned a unique node name when a SAN protocol is enabled:

- iSCSI uses an IQN.
- FC, FCoE, and NVMe/FC use a worldwide node name (WWNN).
- NVMe/TCP uses an NQN.

## SAN multipath



- NetApp recommends multiple pathways from the initiator to the target LUN through separate SAN fabrics.
- ONTAP software uses asymmetric logical unit access (ALUA) to advertise the available and optimal paths.
- Multipathing software on the host identifies paths and manages path failure and recovery.
- NetApp Unified Host Utilities provide easy connection from a Windows or Linux host computer.

NetApp 34 © 2023 NetApp, Inc. All rights reserved.


To provide fault tolerance, an initiator has at least two connections to the target. Each connection sees the LUN independently and would cause confusion without the use of multipath drivers and software. The multipath driver can present the LUN as a single instance.

The figure illustrates a path through node 1 where the LUN is located. The second path is through node 2, which has a connection to the node 1 drive shelves. The multipath driver uses asymmetric logical unit access (ALUA) to identify the path to the node where the LUN is located as the active *direct data path*. The direct data path is sometimes called the *optimized path*. The active path to the node where the LUN is not located is called the *indirect data path*. The indirect data path is sometimes called the *nonoptimized path*. If all the optimized paths fail, the multipath software automatically switches to the nonoptimized paths, maintaining the host access to storage.

NetApp provides Unified Host Utilities software to enable you to easily connect a Windows or Linux host computer to NetApp storage systems. This software supports FC, FCoE, iSCSI, and SAS connections to the storage system.

## iSCSI implementation steps

1. Verify or add the iSCSI protocol license.
2. Enable iSCSI functionality on the storage VM.
3. Create or identify the necessary resources.
4. Map the LUN to the appropriate igroup.
5. Locate the LUN on the host computer and prepare the drive.



The basic process for implementing the iSCSI protocol between an initiator and an ONTAP storage system includes several steps:

- License iSCSI, and then enable iSCSI functionality on the storage VM. You must also identify the iSCSI Software Initiator node name.
- Create a volume, LUN, igroup, and data LIFs.
- Determine which hosts have access to the resources and map the hosts to the LUN.

The LUN is discovered on the host and prepared.

## Storage VM creation: iSCSI

### Storage VM basic details

The screenshot displays the ONTAP System Manager interface. On the left, a navigation menu includes Dashboard, Insights, Storage, Overview, Volumes, LUNs, Consistency Groups, Shares, Buckets, Qtrees, Quotas, Storage VMs, Tiers, Network, and Events & Jobs. The main area shows a table of Storage VMs:

Name	State	Subtype	Co
svm1	running	default	NFS
svm2	running	default	NFS
svm3	running	default	NFS
svm4	running	default	NFS
svm5	running	default	S3

An 'Add Storage VM' wizard is open, showing the following configuration:

- STORAGE VM NAME: svm\_iSCSI
- Access Protocol: iSCSI (selected), SMB/CIFS, NFS, S3, FC, NVMe
- Enable iSCSI:
- NETWORK INTERFACE: cluster2-01, cluster2-02
- For cluster2-01: IP ADDRESS: 192.168.0.181, SUBNET MASK: 24, GATEWAY: Add optional gateway, BROADCAST DOMAIN AND PORT: Default
- For cluster2-02: IP ADDRESS: 192.168.0.182, SUBNET MASK: 24, GATEWAY: Add optional gateway, BROADCAST DOMAIN AND PORT: Default

A green callout box with an arrow points to the wizard, containing the text: **Configure iSCSI LIFs.**

NetApp 36 © 2023 NetApp, Inc. All rights reserved.

View the steps to create a storage VM for an iSCSI environment.

You can enable the iSCSI protocol on an existing storage VM by using System Manager or the `vserver iscsi create -vserver <vserver_name>` command. Verify that the operational status of the iSCSI service on the specified storage VM is up and ready to serve data.

The storage VM creation wizard creates logical network interfaces for the new storage VM. Enter the IP address and subnetwork mask for each logical network interface. You can optionally provide the address of a gateway for each network interface.

To create an iSCSI LIF manually by using the CLI, you must specify the `-data-protocol` parameter as `iscsi`.

CLI LIF creation example:

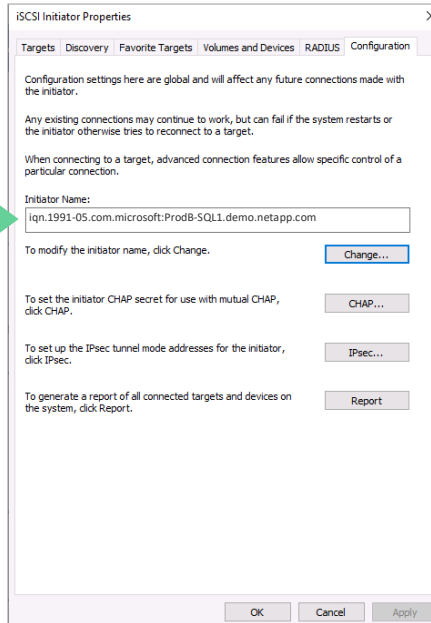
```
::> network interface create -vserver svm7-lif svm7_iscsi_lif1 -data-protocol iscsi  
-home-node rtp-nau-01 -home-port e0e -subnet-name snDefault
```

**Note:** You should create at least one LIF for each node and each network on all storage VMs that serve data with the iSCSI protocol. NetApp recommends having network redundancy through either multiple networks or link aggregation.

# Windows iSCSI implementation

Identify the iSCSI node name

iSCSI initiator name



After you enable a SAN protocol on a storage VM, the next step is to provision storage. Before you can grant the iSCSI client hosts access to the provisioned storage, you must learn the unique IQN of each client host.

The iSCSI Software Initiator creates the iSCSI connection on Windows hosts. The iSCSI Software Initiator is built into Windows Server.

If the system has not yet used an iSCSI Software Initiator, a dialog box appears and requests that you turn on the service. Click **Yes**. The iSCSI Initiator Properties dialog box then appears. You need to identify the iSCSI initiator name before you start the storage VM creation wizard.



## LUN creation

The screenshot shows the ONTAP System Manager interface. The left sidebar contains navigation options: DASHBOARD, INSIGHTS, STORAGE (selected), Overview, Volumes, LUNs, Consistency Groups, Shares, Buckets, Qtrees, Quotas, Storage VMs, Tiers, NETWORK, EVENTS & JOBS, and PROTECTION. The main content area is titled 'LUNs' and has a '+ Add' button with a green 'Create' callout. Below the table header, it says 'No data was found.' The 'Add LUNs' dialog box is open, showing the following fields:

- NAME PREFIX: SQL\_prodB\_Titan\_PartsDB
- NUMBER OF LUNs: 4
- CAPACITY PER LUN: 64 GIB
- HOST OPERATING SYSTEM: Windows
- LUN FORMAT: Windows 2008
- HOST INITIATORS: iqn.1991-05.com.microsoft:ProdB-SQL1.demo.netapp.com

Buttons at the bottom of the dialog are 'More Options', 'Cancel', and 'Save'. A footer at the bottom left reads 'NetApp 38 © 2023 NetApp, Inc. All rights reserved.'

Use the System Manager LUNs page to view and configure storage for SAN client hosts.

In the Add LUNs page, enter a name prefix for the LUNs to be created. This name is also applied to the volume that is created to contain the new LUNs.

Next, specify the number of LUNs to create and the usable size of each LUN.

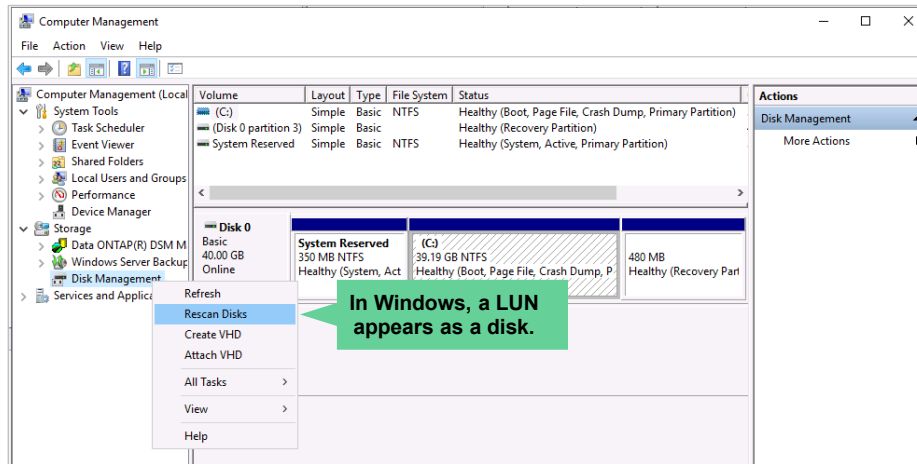
You must specify the operating system type of the client host or hosts that access this LUN. Selecting the incorrect type can result in I/O misalignment and poor performance.

Finally, enter the IQN of one or more client hosts to grant them access to the LUNs. A new host igroup is created automatically.

If you want to assign the LUNs to an existing host igroup, click the **More Options** button. You can also override the default performance service level from which the LUNs are provisioned and enable Snapshot copies and SnapMirror replication of the LUNs.

# Windows LUN implementation

Discover the LUN



To configure the LUN with NTFS, first discover the LUN by selecting **Disk Management > Rescan Disks**.

You can discover and prepare the LUN in Windows in many ways. Each version of Windows might have slightly different tools that you can use. This module illustrates the most often used method. In Windows, a LUN appears as a disk and is labeled as a disk.

1. Open Computer Management.
2. Select Disk Management.
3. If the LUN that you created is not displayed, rescan disks by right-clicking **Disk Management**, or, from the Action menu, select **Rescan Disks**.



## Complete an exercise

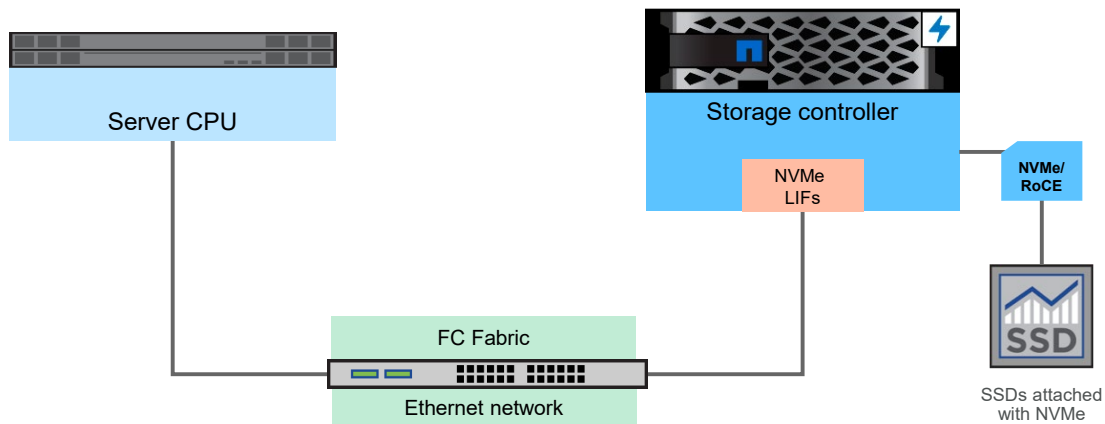
Module 7  
Data access

### Configuring iSCSI in a storage VM

- Access your lab equipment.
- Open your Exercise Guide to Module 7.
- Complete Exercise 3.
- Share your results.


This exercise requires approximately  
**20 minutes.**

## Nonvolatile memory express



RoCE = RDMA over Converged Ethernet

NetApp 41 © 2023 NetApp, Inc. All rights reserved.



Nonvolatile memory express (NVMe) is a block-access protocol that was first supported in ONTAP 9.4 software. The NVMe protocol accesses devices more efficiently than the old Small Computer Systems Interface (SCSI) protocol that was used with iSCSI and FC.


The NVMe over Fibre Channel (NVMe/FC) protocol can use an existing FC network to provide block access to LUNs that reside in the cluster. NVMe/FC was first supported in ONTAP 9.4 software.

The NVMe over TCP (NVMe/TCP) protocol, first available with ONTAP 9.10.1 software, can use an existing Ethernet network to provide block access to LUNs that reside in the cluster.

Some NetApp controllers support “end-to-end” NVMe access using SSD drive shelves connected by remote direct memory access (RDMA) over converged Ethernet (RoCE). Review the Hardware Universe to determine whether your storage controller has been added to the list of supported models.

## SAN terminology

FC	iSCSI	NVMe
Worldwide node name (WWNN)	iSCSI qualified name (IQN)	NVMe qualified name (NQN)
LUN	LUN	Namespace
Initiator group (igroup) LUN mapping/ LUN masking	Initiator group (igroup) LUN mapping/ LUN masking	Subsystem
ALUA	ALUA	Asymmetric namespace access (ANA)

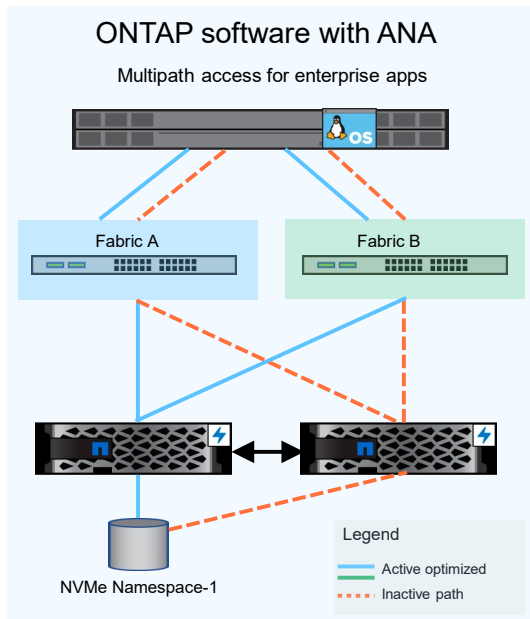
 42 © 2023 NetApp, Inc. All rights reserved.

NVMe adds some new names for some frequently seen structures. The table maps some structures that have different names from the names that are used in FC and iSCSI.

An NVMe qualified name (NQN) identifies an endpoint and is like an iSCSI IQN. A namespace is analogous to a LUN. Both represent an array of blocks that is presented to an initiator. A subsystem is analogous to an igroup and is used to mask an initiator so that it can see and mount namespaces. Asymmetric namespace access (ANA) enables monitoring and managing multiple paths between the initiator and target.

Starting with ONTAP 9.11 software, you can convert LUNs into NVMe namespaces and NVMe namespaces into LUNs. ONTAP 9.12 software expands where NVMe can be deployed to include 12-node ONTAP clusters and MetroCluster IP configurations. With ONTAP 9.12.1 software, you can create up to 8K NVMe subsystems.

## FC-NVMe with ANA



NetApp 43 © 2023 NetApp, Inc. All rights reserved.

### Supported NVMe clients

- Multipath (storage path) failover with *asymmetric namespace access*
- SUSE Enterprise Linux 15
- Red Hat Enterprise Linux 8.1+
- Oracle Linux 8.2
- Application-based high availability only (without ANA)
- Windows
- VMware ESXi

The ANA protocol defines how the NVMe subsystem communicates path and subsystem errors back to the host so that the host can manage paths and failover from one path to another. ANA performs a similar function for FC-NVMe as ALUA does for FC.

ANA categorizes paths as active or inactive. The host is connected to both the fabrics and has multipath access to the namespace through two LIFs (one from each fabric—blue in the figure). These paths are active optimized. The host is also connected to the namespace through an inactive path, as the dashed amber lines in the figure show. In a path failure or controller failover, partner takeover occurs, and the inactive path is turned to an active optimized path for the host to access the data from the namespace. For example, when there is a failure in the path or controller attached to Fabric A, the controller failover notification is sent, and controller B takes over. The inactive path from the host to Fabric A and Fabric B is turned into an active optimized state, and the host can access the data through the controller B.

The implementation of FC-NVMe evolves with each update to the ONTAP software. The expanded FC-NVMe ecosystem now includes VMware ESXi and Windows hosts, in addition to Red Hat, Oracle Linux, and SUSE Linux, with storage path resiliency. Organizations can experience NVMe/FC performance for most workloads.

The NVMe/TCP protocol is new and not yet as widely adopted. Supported operating systems include only Red Hat Enterprise Linux, Oracle Linux, and SUSE Linux so far.

## Storage VM creation: NVMe

### Storage VM basic details

1. Install the NVMe\_of license.
2. Enable the NVMe protocol on a storage VM.
3. Create NVMe protocol data LIFs.
4. Create one or more NVMe namespaces.
5. Grant client hosts access to the namespaces.

Configure NVMe LIFs.

STORAGE VM NAME  
svm\_NVMe

Access Protocol  
SMB/CIFS, NFS, S3  NVMe **Licensed protocols**

Enable NVMe/FC

CONFIGURE FC PORTS ?

Nodes	1a	1b
cluster1-01	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Nodes	1a	1b
cluster1-02	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Enable NVMe/TCP

NetApp 44 © 2023 NetApp, Inc. All rights reserved.

View the steps to create a storage VM for an NVMe environment. The NVMe protocol can now be used over Ethernet and Fibre Channel. The same NVMe\_of license works for both.

You can enable the NVMe protocol on an existing storage VM by using System Manager or the `vserver nvme create -vserver <vserver_name>` command.

The storage VM creation wizard creates logical network interfaces for the new storage VM. Either choose the FC ports to use or enter the IP address and subnetwork mask for each logical network interface. You can optionally provide the address of a gateway for each network interface.

To create an NVMe LIF manually by using the CLI, you must specify the `-data-protocol` parameter as `fc-nvme` or `nvme-tcp`.

CLI LIF creation example:

```
::> network interface create -vserver svm7-lif svm8_NVMe_lif1 -data-protocol nvme-tcp  
-home-node cluster1-02 -home-port e0f
```

**Note:** You should create at least one LIF for each node and each network on all storage VMs that serve data with the NVMe protocol. NetApp recommends having network redundancy through either multiple networks or link aggregation.

## Namespace creation

The screenshot shows the ONTAP System Manager interface with the 'Add NVMe Namespace' dialog box open. The dialog box has the following fields and options:

- NAME PREFIX:** Oracle\_DB7
- NUMBER OF NAMESPACES:** 5
- CAPACITY PER NAMESPACE:** 10 TiB
- HOST OPERATING SYSTEM:** Linux
- HOST NQN:** nqn.2014-08.com.redhat:oradb7

Three green callout boxes provide instructions:

- Create:** Points to the '+ Add' button in the 'NVMe Namespaces' table.
- Select number and size of namespaces.** Points to the 'NUMBER OF NAMESPACES' and 'CAPACITY PER NAMESPACE' fields.
- Select client operating system type.** Points to the 'HOST OPERATING SYSTEM' dropdown menu.
- Enter client host NQNs.** Points to the 'HOST NQN' text input field.

At the bottom of the dialog box, there are buttons for 'More Options', 'Cancel', and 'Save'. Below the dialog box, it says 'No data was found.'

Use the System Manager Namespaces page to view and configure storage for SAN client hosts.

In the Add NVMe Namespaces page, enter a name prefix for the namespaces to be created. This name is also applied to the volume that is created to contain the new namespaces.

Next, specify the number of namespaces to create and the usable size of each namespaces.


You must specify the operating system type of the client host or hosts that access this namespaces. Selecting the incorrect type can result in I/O misalignment and poor performance.

Finally, enter the NQN of one or more client hosts to grant them access to the namespaces. A new host subsystem is created automatically.

If you want to assign the namespaces to an existing subsystem, click the **More Options** button. You can also override the default performance service level from which the namespaces are provisioned and enable Snapshot copies and SnapMirror replication of the namespaces.



## Additional SAN learning

 46 © 2023 NetApp, Inc. All rights reserved.

Where can you learn about advanced topics like the following?

- FC and FCoE protocols
- Windows and Linux initiator implementation
- LUN management and mobility enhancements
- *ONTAP SAN Fundamentals* (online course)
- *ONTAP SAN Administration* (virtual instructor-led course and instructor-led course)
- *ONTAP SAN Implementation* (virtual instructor-led course and instructor-led course)

- *ONTAP SAN Fundamentals* (online course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours000000000018109](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours000000000018109)
- *ONTAP SAN Administration* (virtual instructor-led course and instructor-led course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours000000000029947](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours000000000029947)
- *ONTAP SAN Implementation* (virtual instructor-led course and instructor-led course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours000000000015072](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours000000000015072)

**Note:** Courses are regularly updated and revised. Check the LearningCenter for the latest version and offerings.



## Complete an exercise

Module 7  
Data access

### Configuring NVMe in a Storage VM


- Access your lab equipment.
- Open your Exercise Guide to Module 7.
- Complete Exercise 4.
- Share your results.

This exercise requires approximately  
**15 minutes.**



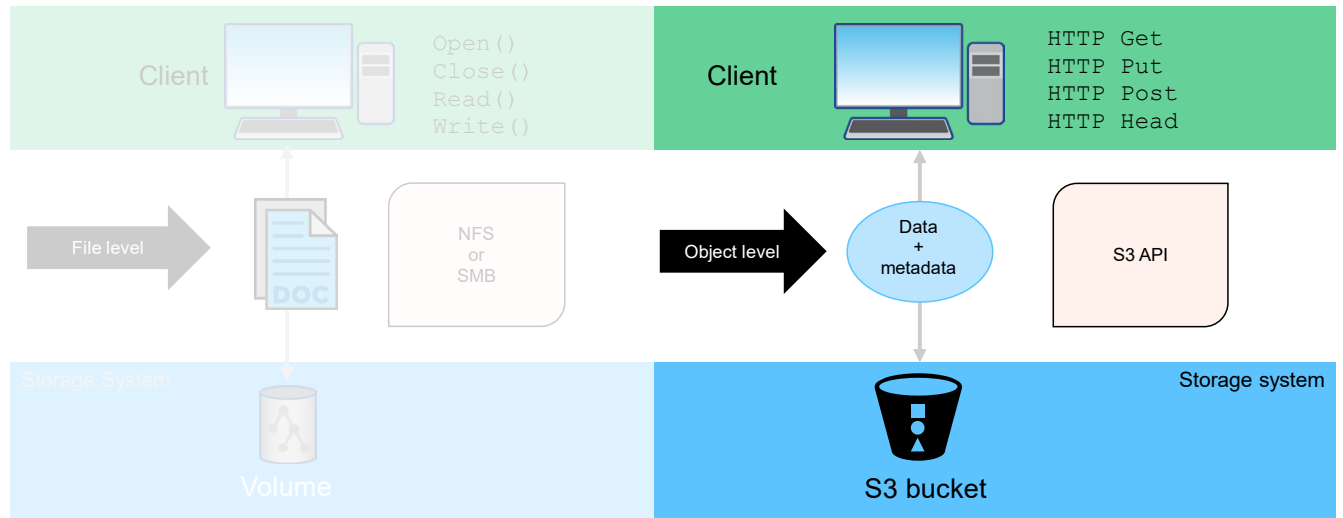
## **Lesson 3**

# **Use object protocols to access data**

 48 © 2023 NetApp, Inc. All rights reserved.

# NAS and object

## Overview



NetApp 49 © 2023 NetApp, Inc. All rights reserved.

Clients and applications can request data at the file level or the object level.


Traditionally, application data was stored in files within file systems. Applications access that data by first opening the file and reading the contents. Changed or new data is written to the file and the file is closed. NAS extended that paradigm by providing file-level access to data on a storage system that uses the NFS and SMB protocols. In a NAS environment, the storage system exports or shares access to a volume, which represents the file system.

Object storage provides object-level access to data on a storage system that uses S3. In an object storage environment, the storage system uses S3 buckets to store objects. Data that is stored by using object storage is typically accessed by a client application that uses an object protocol or API, rather than through file access at the user level.


In this lesson, you examine the ONTAP object-based storage solution.

## S3 objects and buckets

### Basics



- Objects can contain any type of data, like files can.
- Objects are stored in buckets.
- Multiple versions of an object can be kept.
  
- Buckets are a flat namespace with no hierarchy.
  - Objects are identified by a unique key ID.
  - Objects can be tagged with key=value pairs.
- ONTAP S3 buckets reside in FlexGroup volumes and can span multiple aggregates and storage nodes.

 50 © 2023 NetApp, Inc. All rights reserved.

Like a file, an object contains data. The type of data is unimportant. It can be a text document, a picture image, a movie snippet, or an entire system backup. Like a file, an object has information that describes it. A file has a path name that describes the file location in the file system hierarchy. An object has a unique identifier known as a search key, or key ID. Objects are kept inside a single flat container known as a bucket. The objects are not nested as files inside a directory inside other directories. Therefore, the key ID must be unique within the entire bucket. The key ID is alphanumeric, so you can choose a key that makes sense to you. Some applications that use an S3 object store to store files use the file path name for the key ID value. But more often, the object key ID is a machine-generated identifier.

In addition to the file path name, file systems keep information about files, such as the file owner and access control list. This information is called metadata. With S3, applications can provide metadata with the object data. This metadata comes in the form of key=value pairs. You can create your own key=value pairs as needed. The S3 protocol defines many key=value pairs that applications can use to control how objects are stored and managed.

In ONTAP software, the underlying architecture for a bucket is a FlexGroup volume. A FlexGroup volume is a single namespace that is made up of multiple constituent member volumes but is managed as a single volume. Individual objects in a bucket are allocated to individual member volumes and are not striped across volumes or nodes.


When used by buckets, FlexGroup volumes use elastic sizing. Buckets are limited by only the physical maximums of the underlying hardware and have been tested to 20PB and 400 billion files. Architectural maximums can be higher.

The following example creates a 1TB object store server bucket for storage VM vs1:

```
::> vservers object-store-server bucket create -vservers vs1 -bucket testbucket -aggr-list  
aggr1 -size 1TB
```

## FabricPool use of S3

- Use FabricPool technology to migrate cold data to a lower local storage tier.
- Use FabricPool technology to migrate cold data to another ONTAP storage system.



NetApp 51 © 2023 NetApp, Inc. All rights reserved.


There are three primary use cases for client access to NetApp ONTAP S3 services:

- **ONTAP systems that use ONTAP S3 as a local FabricPool tier.** The S3 server and the bucket that contain the capacity tier are on the same cluster as, but a different HA pair than, the performance tier.
- **ONTAP systems that use ONTAP S3 as a remote FabricPool capacity (cloud) tier.** The S3 server and the bucket that contain the capacity tier (for cold data) are on a different cluster than the performance tier (for hot data).
- **External S3 client applications.** ONTAP S3 serves S3 client applications that run on third-party systems.

## ONTAP S3 maturation

Each new version of ONTAP software adds new capabilities to the ONTAP S3 protocol

- ONTAP 9.8 software (initial ONTAP S3 release):
  - Object upload / download / list
  - Bucket creation / deletion
- ONTAP 9.9.1 software:
  - User-defined metadata and object tagging
- ONTAP 9.10.1 software:
  - S3 SnapMirror feature
  - S3 audit
- ONTAP 9.11.1 software:
  - Object versioning
- ONTAP 9.12.1 software:
  - Management of bucket policies
  - Partial object uploads, object copy

 52 © 2023 NetApp, Inc. All rights reserved.



Each new version of ONTAP adds new capabilities for the S3 protocol.

- The ONTAP S3 protocol was introduced with ONTAP 9.8 software and initially provided very basic functions. You could use the S3 protocol to create and delete buckets and upload and download objects.
- ONTAP 9.9.1 software added the ability to store object tags and user-defined object metadata.
- ONTAP 9.10.1 software added the ability to protect S3 buckets. S3 buckets can be backed up to NetApp and non-NetApp destinations or replicated to other ONTAP S3 configurations by using the S3 SnapMirror feature. Access to S3 buckets can be audited, similarly to how ONTAP audits access to NAS shares.
- Support for storing multiple versions of an object was added in ONTAP 9.11.1 software.
- ONTAP 9.12.1 software extends the ONTAP S3 API calls to include the management of bucket policies, the copying of objects, and the uploading of part of an object by copying it from an existing object.

## Comparing ONTAP S3 and StorageGRID solution



Or



### ONTAP S3

- Basic S3 protocol access
- Limited by ONTAP cluster size
- Suitable for FabricPool cloud tier and simple S3 client applications

### StorageGRID

- Globally dispersed object namespace design
- Full S3 command set
- Assumption of rich metadata
- Policy-engine driven data movement
- Integration with public cloud services

**StorageGRID remains the NetApp industry-leading solution for object storage.**

NetApp 53 © 2023 NetApp, Inc. All rights reserved.

The ONTAP S3 implementation offers only basic S3 protocol access. The implementation does not provide all the capabilities of a globally distributed, dedicated object storage solution like StorageGRID.

The StorageGRID solution is a feature-rich object store that is designed to manage objects through the entire lifecycle. A minimal StorageGRID deployment requires 4 storage nodes and 1 admin node. StorageGRID becomes a more cost-effective solution than ONTAP S3 or the public cloud when the total amount of data to be stored exceeds 300TB.



## S3 implementation steps

1. Install the S3 license.
2. Create an S3 storage VM.
3. Create an S3 user.
4. Obtain the user key and secret key.
5. Create a bucket.



## Storage VM creation: S3

Create an S3-enabled storage VM

The screenshot displays the ONTAP System Manager interface. On the left, a navigation menu includes Dashboard, Insights, Storage (Overview, Volumes, LUNs, Consistency Groups, Shares, Buckets, Qtrees, Quotas, Storage VMs), Tiers, Network, and Events & Jobs. The main area shows 'Storage VMs' with a list of existing VMs (svm1 to svm5) and an '+ Add' button. The 'Add Storage VM' dialog is open, showing the 'SMB/CIFS, NFS, S3' tab. The 'STORAGE VM NAME' is 'svm\_S3'. Under 'Access Protocol', 'Enable S3' is checked. The 'NETWORK INTERFACE' section shows two interfaces: 'cluster2-01' (IP: 192.168.0.195, Subnet Mask: 24) and 'cluster2-02' (IP: 192.168.0.196, Port: Automatical...). The 'Storage VM Administration' dialog is also open, with 'Manage administrator account' checked, and fields for 'USER NAME' (vsadmin), 'PASSWORD', and 'CONFIRM PASSWORD'. A 'Save' button is visible at the bottom right of the dialog.

NetApp 55 © 2023 NetApp, Inc. All rights reserved.

You create a storage VM for S3 in nearly the same way as you do for NFS and SMB.

To add a new storage VM, click **Storage > Storage VMs**, and then click **Add**.

If this system is a new system with no existing storage VMs, click **Dashboard > Configure Protocols**.

If you are adding an S3 server to an existing storage VM, click **Storage > Storage VMs**, select a storage VM, click **Settings**, and then click under **S3**.

As when configuring other data storage VMs, you need to enter the IP addresses of the data LIFs you want to create and decide whether to create an administrator account for the storage VM.

## Storage VM creation: S3

### Certificates

Access Protocol

SMB/CIFS, NFS, S3    iSCSI    FC

Enable SMB/CIFS

Enable NFS

Enable S3

S3 SERVER NAME

s3.company.com

Enable TLS

PORT

443

CERTIFICATE

Use system-generated certificate [?](#)

Use external-CA signed certificate

Use HTTP (Non-secure)

Added Storage VM

STORAGE VM: svm\_S3    S3 SERVER NAME: s3.company.com

User Details

USER NAME: sm\_s3\_user

⚠ The secret key will not be displayed again. Save this key for future use.

ACCESS KEY: IQJ0H121N2V3JJS4NAGO

SECRET KEY: [Show secret key](#)

Certificate

CERTIFICATE SERIAL NUMBER: 16D4F2D5ED718972    CERTIFICATE EXPIRATION DATE: Tuesday, Feb 13, 2024, 5:52 PM

CERTIFICATE DETAILS

```
-----BEGIN CERTIFICATE-----
MIIDZjCCAK6gAwIBAgIIftTy1e1xiIwDQYJKoZIhvcNAQELBQAwJTEWMBQGA1UE
AxQNU1ZNX1NZU01HUI9DQTELMakGA1UEBhMCVWwHhcNMjE4MTc1MjQ0WhcN
-----
```

[Download](#)    [Close](#)

NetApp 56 © 2023 NetApp, Inc. All rights reserved.

S3 security configuration differs from the other NAS protocols.

When you click **Enable S3**, you are prompted for the S3 Server Name. This name is the Fully Qualified Domain Name (FQDN) that clients use.

Select the certificate type. Whether you select a system-generated certificate or one of your own, a certificate is required for client access. If you select the system-generated certificate, you see the certificate information when the new storage VM creation is confirmed. Click **Download** and save the certificate for client access. If you need the certificate information again, click **Storage > Storage VMs**, select the storage VM, and then click **Settings**.

The storage VM creation confirmation window also includes the S3 user access key and secret access key. The secret key is not displayed again, so be sure to download and save the key for use when configuring client access.

## Create S3 user accounts

The screenshot displays the ONTAP System Manager interface. The main window shows the 'Storage VMs' page with a table listing storage VMs (svm5, svm\_S3) and a 'Protocols' panel for S3 showing 'STATUS Enabled', 'TLS Enabled', and 'HTTP Disabled'. An 'Add User' dialog box is open, showing fields for 'NAME' (image\_archiver), 'ACCESS KEY' (ESMFV12214HT17T75E53), and 'SECRET KEY' (P693Q6j4c\_\_bi6ppM14VdCUatX\_zH7h\_ZZ2yZ2Ny). A 'Save' button is highlighted in red. The dialog box also includes a 'Download' button and a 'Close' button. A warning message states: 'The secret key will not be displayed again. Save this key for future use.'

You must edit the storage VM settings to add S3 users, and to add users to groups. User authorization is required on all ONTAP object stores to restrict connectivity to authorized clients.

To edit the storage VM, click **Storage > storage VMs**, click the storage VM, click the **Settings** tab, and then click the edit button (pencil icon) in the S3 panel.

To add a user, click **Users**, and then click **Add**. Enter a name, and then click **Save**. Be sure to save the access key and secret key, both of which are required for access from S3 clients.

If desired, to add a group, click **Groups**, and then click **Add**. Enter a group name, and then select from a list of users. You can select an existing group policy, add a policy now, or add a policy later.

The `vserver object-store-server user create` command creates an object store user. This command generates an access key and a secret key to use for authentication.

```
> vserver object-store-server user create -vserver <svm_name> -user <user_name>
```

The `vserver object-store-server user show` command displays information about the object store user. You must be in the advanced privilege level to view the secret access key.

```
> set adv
```

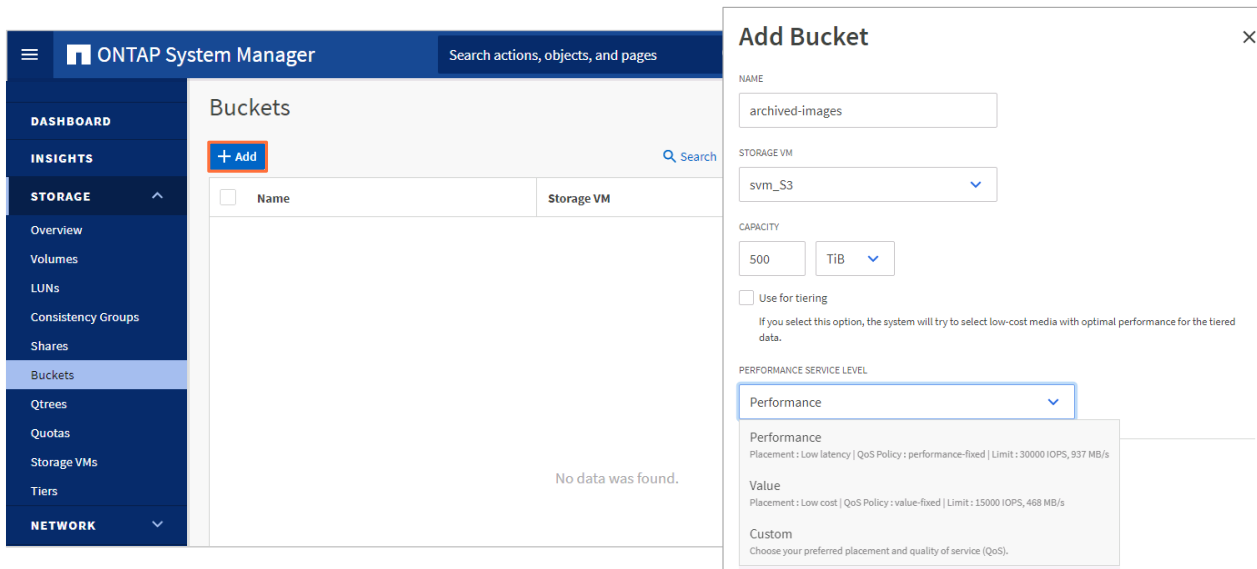
```
*> object-store-server user show <username>
```

The `vserver object-store-server user regenerate-keys` command regenerates a new access key and secret key for an object store user.

```
*> vserver object-store-server user regenerate-keys -vserver vs1 -user user1
```

## Create an S3 bucket

### Performance service levels



NetApp 58 © 2023 NetApp, Inc. All rights reserved.

For remote client access, you must configure buckets in an S3-enabled storage VM. If you create a bucket in a storage VM that is not S3-enabled, the bucket is available for only local tiering.

Beginning in ONTAP 9.8 software, when you provision storage, quality of service (QoS) is enabled by default. You can disable QoS or choose a custom QoS policy during the provisioning process or later.

To add a new bucket on an S3-enabled storage VM, complete the following steps:

1. Click **Storage > Buckets**, and then click **Add**.
2. Enter a name, select the storage VM, and then enter a size.

If you click **Save** at this point, a bucket is created with a QoS (performance) level that is the highest available for your system.

You can click **More Options** to configure user permissions and performance level when you configure the bucket, or you can modify these settings later.

If you intend to use the S3 object store for FabricPool tiering, consider selecting **Use for tiering** rather than a performance service level.

## Create an S3 bucket

### Controlling access

Permissions

Copy access permissions from an existing bucket

Principal	Effect	Actions	Resources	Conditions
No data				

[+ Add](#)

[Save](#) [Cancel](#)

New Permission

PRINCIPAL

image\_archiver X

EFFECT

Allow

ACTIONS

GetObject X PutObject X  
DeleteObject X ListBucket X

RESOURCES

archived-images,archived-images/\*

Conditions

[+ Add](#)

[Cancel](#) [Save](#)

NetApp 59 © 2023 NetApp, Inc. All rights reserved.

No users are granted access to the bucket unless group policies are already in effect.

You should not use the S3 root user to manage ONTAP object storage and share its permissions because this user has unlimited access to the object store. Instead, create a user or group with administrative privileges that you assign.

You must create users and groups before configuring their permissions.

Choose the following:

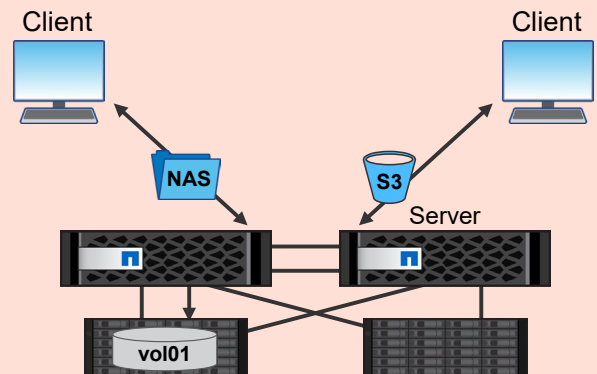
- Who can access the bucket (the principals)
- Whether the access rule allows or prevents access (the effect)
- Which actions the user can perform
- The resources on which the user can perform those actions

After the bucket has been created, verify access to the new bucket by entering the following into an S3 client application or into another ONTAP system:

- The S3 server certificate authority (CA) certificate
- The user access key and secret key
- The S3 server FQDN name and bucket name

## S3 access to NAS shares

- S3 users can create, read, delete, and list files in NAS shares.
- A NAS directory is assigned a "NAS bucket."
- ONTAP software translates S3 object identifiers into NAS pathnames.
  - NFS: /mountpoint/folder1/folder2/filename
  - SMB: \\servername\folder1\folder2\filename
  - S3: "folder1/folder2/filename"
- The bucket policy controls access by S3 users.
  - By default, access is denied to all S3 users.
  - S3 user names must map into local NFS or SMB users.
  - S3 user names can map to a default local user.



NetApp 60 © 2023 NetApp, Inc. All rights reserved.

Beginning with ONTAP 9.12.1 software, you can enable clients running the S3 protocol to access the same data that is being served to clients that use the NFS and SMB protocols.

ONTAP creates "S3 NAS buckets" that enable S3 clients to create, read, delete, and list files in NAS storage by using S3 object requests. A NAS bucket can be associated with a single WAFL directory, a Qtree, or an entire FlexVol volume or FlexGroup volume. NAS buckets are different from ONTAP S3 buckets and are not managed in the same way. S3 NAS buckets cannot be used as a capacity tier for FabricPool technology.

ONTAP translates NAS pathnames to and from S3 object identifiers. The S3 object identifier consists of the file or folder pathname, relative to the directory associated with the NAS bucket.

Before enabling S3 access to NAS shares, you must enable the SVM to speak the S3 protocol and accept S3 protocol traffic on a data LIF. Next, the SVM object store service must be started. Only then can you create a NAS bucket.

The bucket policy determines which S3 users can perform which actions on which resources. The S3 user names must map into an ONTAP local user name or a default user name. Lightweight Directory Access Protocol (LDAP) is not supported. Mapping an S3 user name to a corresponding LINUX/UNIX or Windows user enables authorization checks on the NAS files to be honored when those files are accessed by S3 clients.

## S3 bucket access policies

Policy statements use the following structure to specify permissions:

Grant **<Effect>** to allow/deny  
**<Principal>** to perform  
**<Action>** on  
**<Resource>** when  
**<Condition>** applies.

The screenshot shows a 'New Permission' dialog box with the following configuration:

- PRINCIPAL:** All users of this stor...
- EFFECT:** Allow
- ACTIONS:** ListBucket, GetObject, PutObject, DeleteObject
- RESOURCES:** nas-svm1-data1,nas-svm1-data1/\*
- Conditions:**

KEY	OPERATOR	VALUE
source_ips	ip_address	192.168.0.5,192.168.0.11

**NetApp** 61 © 2022 NetApp, Inc. All rights reserved.

Create a bucket policy to grant S3 users access to the NAS share. A S3 bucket policy consists of one or more policy statements, each which contains the following elements:


- **Principal:** To whom the policy applies  
Provide S3 user names or use the default (all users).
- **Effect:** Select whether the selected actions are allowed or denied.
- **Actions:** A list of actions that are allowed or disallowed  
ONTAP software currently supports the following actions on S3 NAS buckets: GetObject, PutObject, DeleteObject, ListBucket, GetBucketAcl, GetObjectAcl, GetObjectTagging, PutObjectTagging, DeleteObjectTagging, GetBucketLocation, GetBucketVersioning, PutBucketVersioning, and ListBucketVersions. Wildcards are accepted for this parameter.
- **Resources:** The folder or file paths in which the actions are allowed or denied or the default, which is the root directory of the bucket
- **Conditions:** One or more conditions can be applied to the principle to limit the scope of the policy statement. In this example, the statement applies to all users accessing the storage system from client hosts with either IP address 192.168.0.5 or 192.168.0.11.

S3 bucket policies can also be managed from the command line by using the `vserver object-store-server bucket policy` commands.



# Knowledge check

Module 7: Data access

 62 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

A volume called svm1\_vol2 is created on the aggr2 aggregate and mounted to the junction path /svm1/vol2. An administrator moves the volume to the aggr1 aggregate.

After the move, what is the path to the volume?

- a. /aggr1/svm1/svm1\_vol2
- b. /svm1/vol2
- c. /vol/svm1\_vol2
- d. /aggr1/svm1\_vol2

## Knowledge check

**When you create a storage VM to support SAN protocols, which configuration steps need to be made? (Choose three)**

- a. Configure a SAN LIF on each cluster node.
- b. Choose an IPspace for the storage VM.
- c. Create an interface group (ifgroup).
- d. Create a storage VM management LIF.

## References

- ONTAP 9 Documentation Center:  
<http://docs.netapp.com/ontap-9/index.jsp>
- TR-4080 Best Practices for Modern SAN
- TR-4684: Implementing and Configuring Modern SANs with NVMe/FC

*NVMe organization website*  
<https://nvmexpress.org/>

## Module summary


This module focused on enabling you to do the following:

- Use NAS protocols to access data
- Use SAN protocols to access data
- Use object protocols to access data



## Complete exercises

Module 7  
Data access

 67 © 2023 NetApp, Inc. All rights reserved.

### Configuring the S3 Protocol in a Storage VM

#### Managing NAS Storage VMs

- Access your lab equipment.
- Open your Exercise Guide to Module 7.
- Complete Exercises 5 and 6.
- Share your results.

This exercise requires approximately  
**25 minutes.**


See the instructions in your Exercise Guide.



## Share your experiences

Roundtable discussion


- Were you able to use both the SMB and NFS protocols to access the same volume in the namespace?
- How does partitioning and formatting a LUN from the Windows host differ from partitioning and formatting a physical disk in Windows?
- Why do you need FlexVol volumes?

 68 © 2023 NetApp, Inc. All rights reserved.

Have a roundtable discussion with the class to answer these questions. You should also add any comments about experiences or lessons learned during the exercises that others might find helpful.

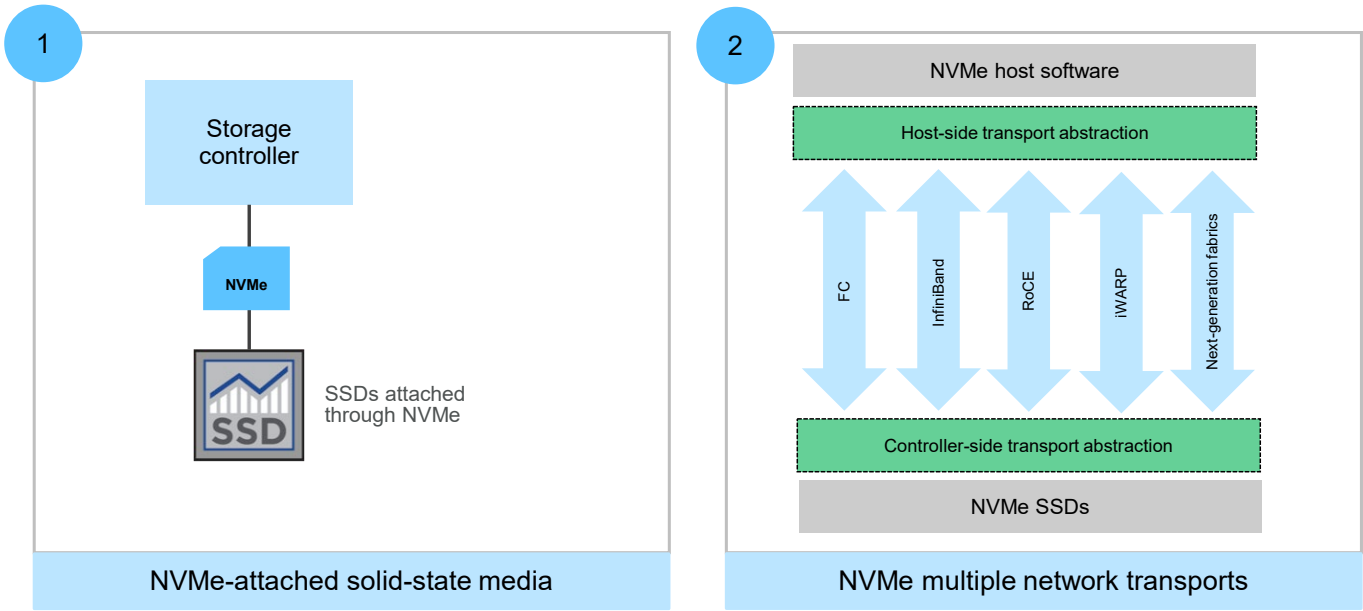


# Addendum NVMe and NVMe/FC

 69 © 2023 NetApp, Inc. All rights reserved.



# NVMe and modern SAN



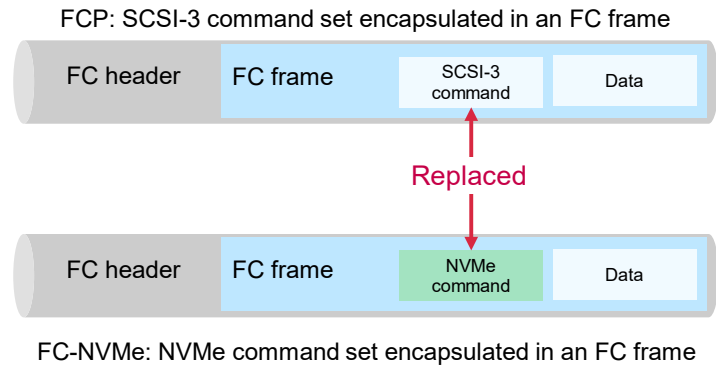
NetApp 70 © 2023 NetApp, Inc. All rights reserved.

NVMe is most often used to attach disks and disk shelves. Implementing end-to-end NVMe requires NVMe-attached solid-state media and NVMe transport from the storage controller to the host server. NVMe-oF adds NVMe as a new block storage protocol type. NVMe-oF defines and creates specifications for how to transport NVMe over various network storage transports such as FC and InfiniBand.

## FC-NVMe and FCP frames

- Share hardware and fabric components
- Can coexist on the same optical fibers, ports, switches, and storage controllers

NVMe/FC and FCP look similar.




FCP and NVMe share hardware and fabric components and can coexist on the same FC infrastructure, including host bus adapters (HBAs), switches, zones, targets, and cabling.

See the NetApp Interoperability Matrix Tool (IMT) to verify the latest supported solution stack for ONTAP software.

FC-NVMe and FCP look similar. FCP encapsulates the SCSI-3 command descriptor block inside FC frames. FC-NVMe swaps out the SCSI-3 command descriptor block for the new NVMe command set, thus offering substantial improvements to throughput and latency.




# Addendum S3 deployments

 72 © 2023 NetApp, Inc. All rights reserved.

## S3 protocol in ONTAP 9.12 software

ONTAP 9.8 and later software supports S3 object storage in production environments.

Focus	Supported	Not supported
Data protection	<ul style="list-style-type: none"> <li>NetApp Cloud Sync service</li> <li>System-scheduled NetApp Snapshot copies</li> <li>NetApp Volume Encryption (NVE)</li> <li>NetApp Storage Encryption (NSE)</li> </ul>	<ul style="list-style-type: none"> <li>Erasure coding</li> <li>MetroCluster software</li> <li>NDMP</li> <li>SnapLock compliance software</li> <li>SyncMirror feature</li> <li>Storage-Level Access Guard</li> <li>SMTape</li> <li>Storage VM disaster recovery</li> <li>Transport Layer Security (TLS)</li> <li>User-created Snapshot copies</li> <li>WORM</li> </ul>
Storage efficiency	<ul style="list-style-type: none"> <li>Inline deduplication</li> <li>Inline compression</li> <li>Compaction</li> </ul>	<ul style="list-style-type: none"> <li>Aggregate-level efficiencies</li> </ul>
Additional features	<ul style="list-style-type: none"> <li>Quality of service (QoS) maximums (ceiling)</li> <li>QoS minimums (floors)</li> </ul>	<ul style="list-style-type: none"> <li>FabricPool technology</li> <li>NetApp FPolicy software</li> <li>Qtrees</li> <li>Quotas</li> </ul>

 73 © 2023 NetApp, Inc. All rights reserved.

In ONTAP 9.7 software, S3 object storage was introduced as a product preview. That version was not intended for production environments. Only ONTAP 9.8 and later releases support S3 object storage in production environments.

The ONTAP S3 protocol is appropriate if you want S3 capabilities on existing clusters without additional hardware and management. For deployments larger than 300TB, NetApp StorageGRID software remains the NetApp flagship solution for object storage. For more information, see the [StorageGRID documentation](#).

In this ONTAP release, some standard features of FlexGroup volumes are not supported when used for S3 object storage:


- FlexCache volumes
- MetroCluster software
- NDMP
- SnapLock software
- SnapMirror Cloud
- The SyncMirror feature
- SMTape
- Storage VM disaster recovery
- Volume clone of the FlexGroup volume containing ONTAP S3 buckets
- User-created Snapshot copies

Unsupported S3 object storage functionality includes:

- Erasure coding

# Module 8

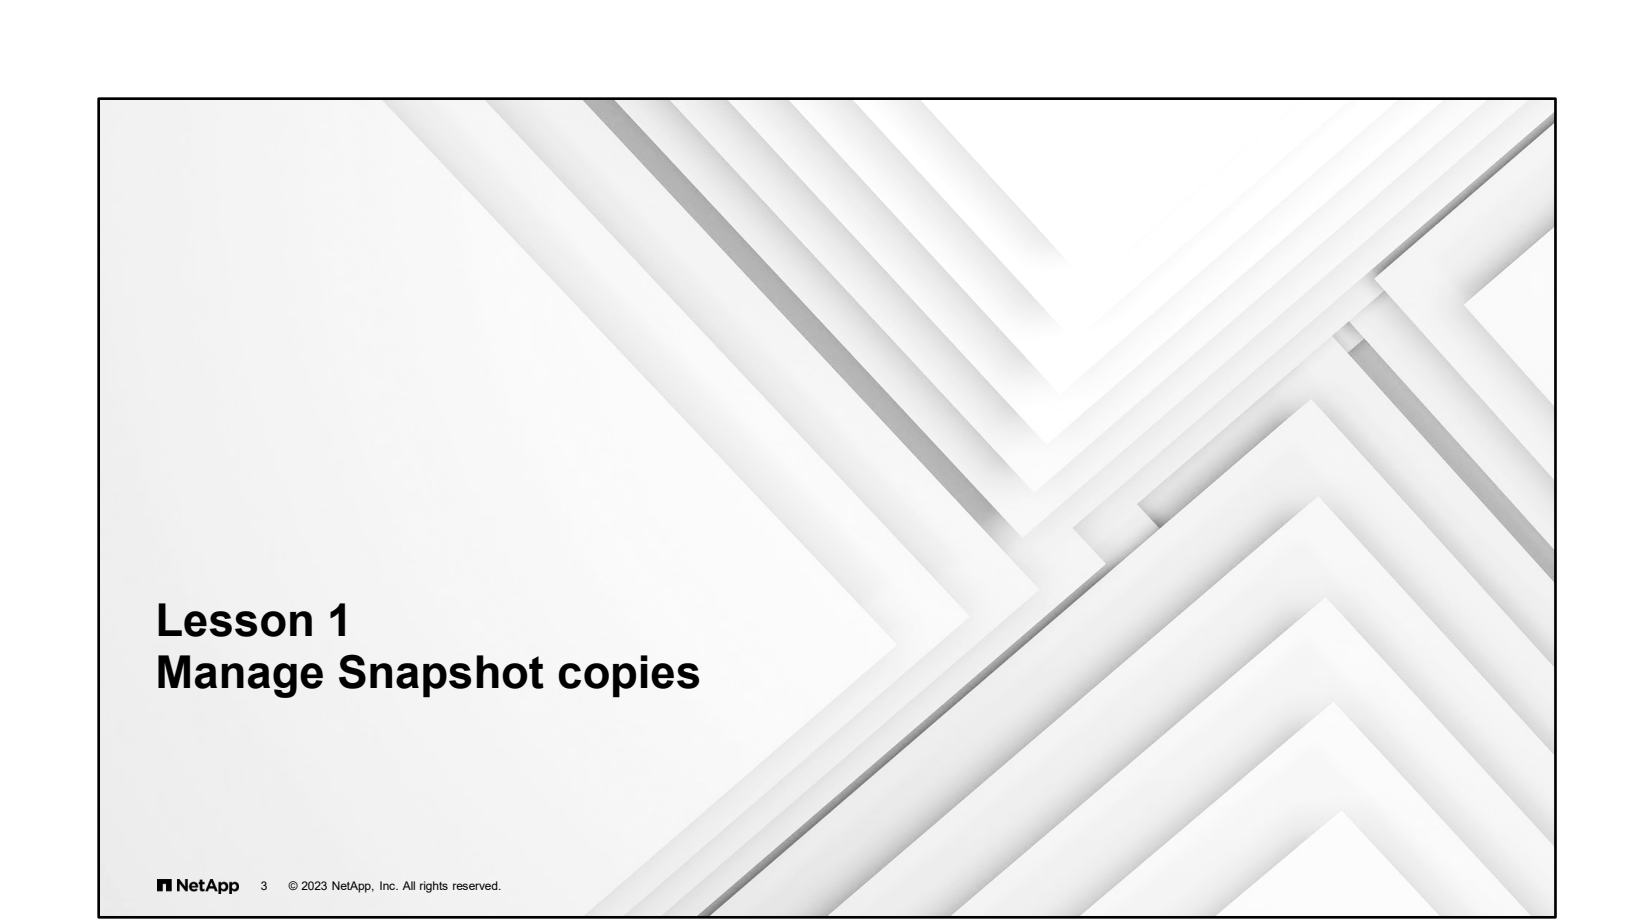
## Data protection

 1 © 2023 NetApp, Inc. All rights reserved.

## About this module


This module focuses on enabling you to do the following:

- Manage Snapshot copies
- Restore data from Snapshot copies
- Back up and replicate data
- Use encryption to prevent unauthorized access to data



# Lesson 1

## Manage Snapshot copies

 3 © 2023 NetApp, Inc. All rights reserved.

## Snapshot copies

- A Snapshot copy is a read-only, point-in-time image of a FlexVol volume.
- The copy consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the most recent Snapshot copy was made.
- Snapshot copies owe their efficiency to the NetApp WAFL file system, which uses metadata to point to blocks on disk and writes to a new block rather than overwrite existing blocks.
- Instead of moving old blocks to a pool of space for Snapshot copies, old blocks remain in place. Only the pointers move from the active file system to the Snapshot copies.

Understanding the technology that is used to create a Snapshot copy helps you to understand how space is used. Furthermore, understanding the technology helps you to understand features such as FlexClone technology, deduplication, and compression.

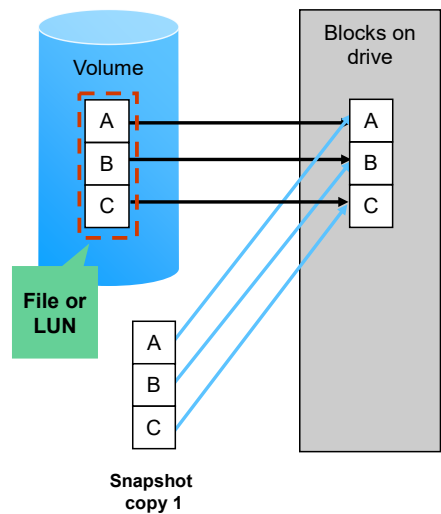
A Snapshot copy is a local, read-only, point-in-time image of data. Snapshot copy technology is a built-in feature of NetApp WAFL storage virtualization technology and provides easy access to old versions of files and LUNs.

Snapshot technology is highly scalable. A Snapshot copy can be created in a few seconds, regardless of the size of the volume or the level of activity on the NetApp storage system. After the copy is created, changes to data objects are reflected in updates to the current version of the objects, as if the copy did not exist. Meanwhile, the Snapshot copy of the data remains stable. A Snapshot copy incurs no performance overhead. Users can store as many as 1023 Snapshot copies per volume. All the Snapshot copies are accessible as read-only and online versions of the data.



## Snapshot copy technology

### Create Snapshot copy 1



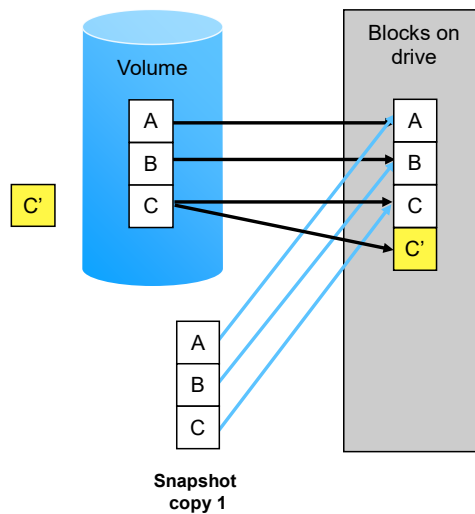
#### 1. Create Snapshot copy 1:

- Pointers are copied.
- No data is moved.

NetApp ONTAP software uses pointers to map logical FlexVol volume blocks to locations on physical drives. These pointers are kept in the volume inode map. When ONTAP software creates a Snapshot copy, ONTAP software preserves the inode map at that point in time and then continues to change the inode map on the active file system. ONTAP software then retains the old version of the inode map. No data is moved when the Snapshot copy is created.

## Snapshot copy technology

Continue writing data



1. Create Snapshot copy 1.
2. Continue writing data:
  - Data is written to a new location on the disk.
  - Pointers are updated.

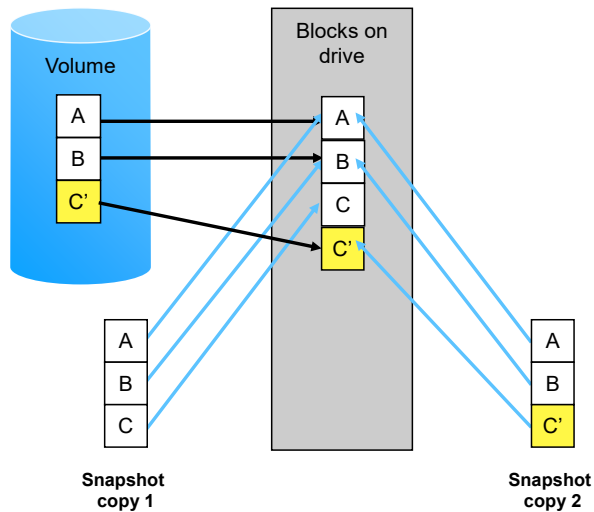
NetApp 6 © 2023 NetApp, Inc. All rights reserved.

When ONTAP software writes changes to disk, the changed version of block C is written to a new location. In the example, C' is the new location. ONTAP software changes the pointers rather than moving data.

The file system avoids the parity update changes that are required if new data is written to the original location. If the WAFL file system updated the same block, the system would need to perform multiple parity reads to update both parity disks. The WAFL file system writes the changed block to a new location, writing in complete stripes and without moving or changing the original data blocks.

## Snapshot copy technology

### Create Snapshot copy 2



1. Create Snapshot copy 1.
2. Continue writing data.
3. Create Snapshot copy 2:
  - Pointers are copied.
  - No data is moved.
  - Block C consumes Snapshot space because the active file system no longer references Block C.

NetApp 7 © 2023 NetApp, Inc. All rights reserved.

When ONTAP software creates another Snapshot copy, the new Snapshot copy points only to the unchanged blocks A and B and to block C'. Block C' is the new location for the changed contents of block C. ONTAP software does not move any data. The system keeps building on the original active file system. The method is simple and good for disk use. Only new and updated blocks use additional block space.

When Snapshot copy 1 is created, the copy consumes no space because the copy holds only pointers to blocks on disk. When C' and Snapshot copy 2 are created, the primary pointer from block C changes from the active file system to Snapshot copy 1. Snapshot copy 1 now owns the block and the space that the block consumes. If Snapshot copy 1 is deleted, the C block has no more pointers referencing it. The block is returned to the available free space.

## Create a Snapshot copy

The screenshot shows the NetApp ONTAP System Manager interface. The left sidebar contains navigation options: DASHBOARD, INSIGHTS, STORAGE (expanded), Overview, Volumes (selected), LUNs, Consistency Groups, Shares, Buckets, Qtrees, Quotas, and Storage VMs. The main content area displays the 'Volumes' page for 'c1\_svm1\_vol1'. A table lists several volumes, with 'c1\_svm1\_vol1' selected. A modal window titled 'Add Snapshot Copy' is open, showing a text input field for 'SNAPSHOT COPY NAME' containing 'snap.2022-02-18\_194404'. Below the input field are 'Cancel' and 'Add' buttons. An orange arrow points from the '+ Add' button in the volume table to the modal. At the bottom of the screenshot, a terminal window shows the following command:

```
cluster1::> snapshot create -vserver svm4 -volume svm4_vol1002  
-snapshot vol2-pre-app-upgrade
```

NetApp 8 © 2023 NetApp, Inc. All rights reserved.

You can use NetApp ONTAP System Manager (formerly OnCommand System Manager) or clustershell to create, schedule, and maintain Snapshot copies for volumes and aggregates.

## Snapshot copy design

Snapshot copies are the first line of defense against accidental data loss or inconsistency.

- Do not create more Snapshot copies than necessary.
- Check and adjust the volume Snapshot copy reserve defaults.
- To control storage consumption, configure Snapshot copy automatic deletion and volume automatic increase.
- See TR-4678 for guidance on planning Snapshot copies of NetApp FlexGroup volumes.



To provide efficient use of drive space, deploy only the required number of Snapshot copies on each volume. If you deploy more Snapshot copies than are required, the copies consume more drive space than necessary.

You might need to adjust default settings for the Snapshot copy reserve for volumes:

- The Snapshot copy reserve guarantees that you can create Snapshot copies until the reserved space is filled.
- When Snapshot copies fill the reserved space, the Snapshot blocks compete for space with the active file system.

NetApp ONTAP FlexGroup volumes have special considerations for taking a Snapshot copy. All FlexGroup volumes must temporarily halt data access to help ensure a crash-consistent state. If the Snapshot copy does not finish in 10 seconds, the copy fails. Technical report TR-4678 covers the process of configuring FlexGroup Snapshot copies for use by ONTAP Snap and Flex features. See the References pages for the URL link to the technical report.

## Naming conventions for Snapshot copies

- A Snapshot copy name can have a prefix or schedule name, timestamp, comment, and label:

vserver	volume	snapshot
-----	-----	-----
svm4	svm4_vol1002	2HourSnapshot.2023-03-11_1030
		

- Snapshot copy names cannot be longer than 255 characters.

Administrators can use the Snapshot copy prefix, timestamp, and comment features to quickly determine why a Snapshot copy was created.


The Snapshot copy name is made up of the *prefix* and *timestamp*. The prefix is an optional string of characters that you can specify when defining a Snapshot policy. If no prefix is specified, the Snapshot schedule name is used for the prefix. For manually triggered Snapshot copies, the prefix is set to “snap.”

The Snapshot name, including the prefix and timestamp, is limited to 255 characters.

When you run the `volume snapshot create` command from the CLI, you can use the `comment` option to annotate a Snapshot copy. You can view the comment by using the `instance` or `fields` option of the `volume snapshot show` command.

## The Snapshot copy reserve

- The Snapshot reserve is a storage space set aside inside a volume.
  - Often depicted as a partition
  - Actually a soft quota
- The reserve holds blocks that are no longer in the active file system but are still referenced by Snapshot copies.
- The reserve is not used for file system writes.
- The reserve can be increased or decreased.



NetApp 11 © 2023 NetApp, Inc. All rights reserved.

The `volume show -fields percent-snapshot-space` command displays the percentage of storage space that has been set aside for Snapshot copies.

Use the `volume modify` command to change the percentage of storage space that is set aside for the Snapshot copies of a volume. For example, to increase the Snapshot copy reserve from 5% to 10% for the volume named *engineering*, enter the following command:

```
volume modify -vserver svm4 -volume engineering -percent-snapshot-space 10
```

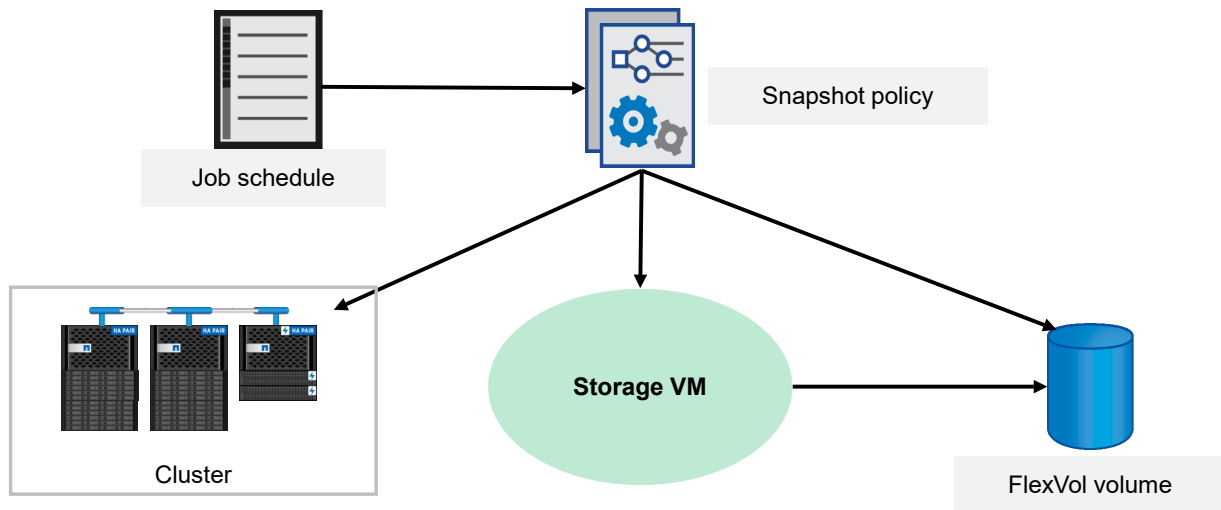
By default, volume Snapshot copies are stored in the Snapshot copy reserve storage space. The Snapshot copy reserve space is not counted as part of the volume disk space that is allocated for the active file system. (For example, when you enter the `volume show` command for a volume, the amount of available disk space that is shown does not include the amount of disk space that the `snap reserve` command reserves.)

When a Snapshot copy is first created, none of the Snapshot copy reserve is consumed. The Snapshot copy protects the active file system at a point in time when the Snapshot copy was created. As the Snapshot copy ages and the active file system changes, the Snapshot copy begins to own the data blocks that the current active file system deleted or changed. The Snapshot copy begins to consume the Snapshot copy reserve space. The amount of disk space that Snapshot copies consume can grow, depending on the length of time that a Snapshot copy is retained and the rate of change of the volume.

Sometimes, if the Snapshot copy is retained for a long period and the active file system has a high rate of change, the Snapshot copy can consume 100% of the Snapshot copy reserve. If the Snapshot copy is not deleted, the copy can consume a portion of the drive space that is intended for the active file system. Monitor and manage Snapshot copies so that drive space is properly managed.

**Note:** Even if the Snapshot copy reserve is set to 0%, you can still create Snapshot copies. If no Snapshot copy reserve exists, Snapshot copies consume blocks from the active file system over time.

## The Snapshot policy



NetApp 12 © 2023 NetApp, Inc. All rights reserved.

Storage VM = storage virtual machine, also known as SVM

A Snapshot policy enables you to configure the frequency and maximum number of Snapshot copies that are created automatically.

When you create a storage VM (storage virtual machine, also known as SVM), you can specify a Snapshot policy that becomes the default for all FlexVol volumes that are created for the storage VM. When you create a FlexVol volume, you can specify which Snapshot policy you want to use.

The default Snapshot policy might meet your needs. The default Snapshot copy policy is useful if users rarely lose files.

The default Snapshot policy specifies the following:

- Weekly schedule to keep two weekly Snapshot copies
- Daily schedule to keep two daily Snapshot copies
- Hourly schedule to keep six hourly Snapshot copies

However, if users often lose files, you should adjust the default policy to keep Snapshot copies longer:

- Weekly schedule to keep two weekly Snapshot copies
- Daily schedule to keep six daily Snapshot copies
- Hourly schedule to keep eight hourly Snapshot copies

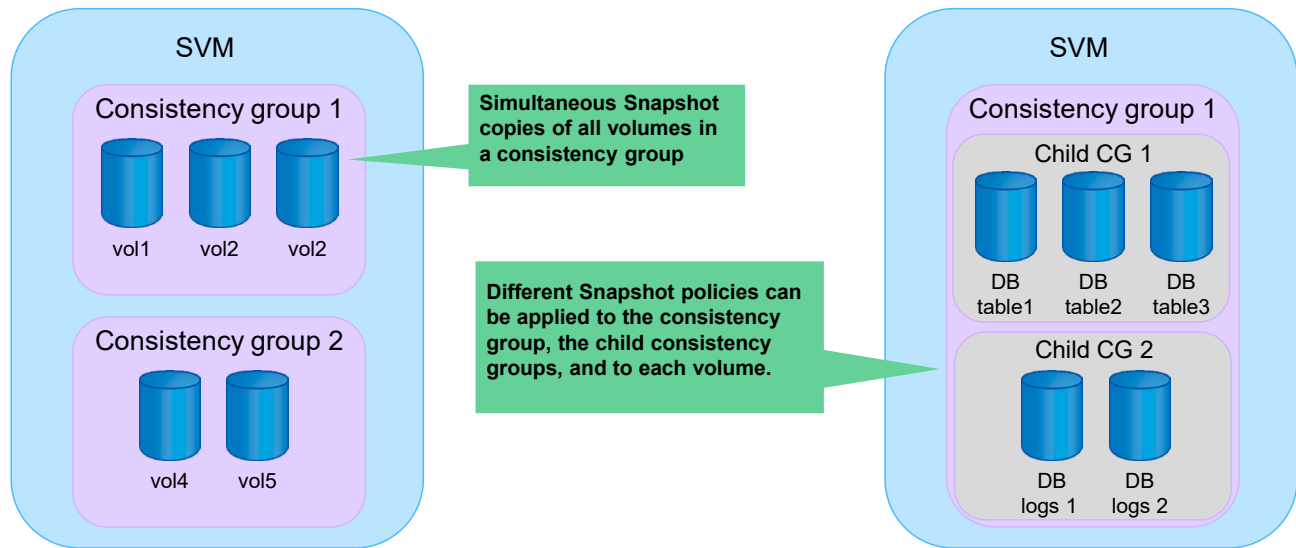
For typical systems, only 5% to 10% of the data changes each week. Six daily and two weekly Snapshot copies consume 10% to 20% of disk space. Adjust the Snapshot copy reserve for the appropriate amount of disk space for Snapshot copies.

Each volume on a storage VM can use a different Snapshot copy policy. For active volumes, create a Snapshot schedule that creates Snapshot copies every hour and keeps the copies for just a few hours, or turn off the Snapshot copy feature.

You back up Snapshot copies to the vault destination. If an empty label ("") is specified, the existing label is deleted.



## Snapshot consistency groups



NetApp 13 © 2023 NetApp, Inc. All rights reserved.

A consistency group is a collection of volumes that provides a write-order consistency guarantee for an application workload spanning multiple volumes. Consistency groups provide simultaneous crash-consistent or application-consistent Snapshot copies of a collection of volumes at a point in time. Consistency groups support any FlexVol volume regardless of protocol (NAS, SAN, or NVMe) and can be managed through the ONTAP REST APIs or in ONTAP System Manager under the **Storage > Consistency Groups** menu item.

Consistency groups can exist on their own or in a hierarchical relationship. An individual consistency group is a collection of volumes. Volumes can have individual local Snapshot policies in addition to inherited Snapshot policies from a consistency group.

Larger application workloads might require multiple consistency groups. In these situations, multiple consistency groups can be placed together in a hierarchical relationship. In this configuration, single consistency groups become the child components of a parent consistency group. The parent consistency group can include up to five child consistency groups.

Consistency groups offer local protection through Snapshot policies and remote protection through SnapMirror Business Continuity (SM-BC). Beginning with ONTAP 9.8 software, you can use SM-BC to protect applications with LUNs, enabling applications to fail over transparently, ensuring business continuity in case of a disaster.

Beginning in version 9.11.1, ONTAP software supports two-phase commits for consistency group (CG) Snapshot creation. A two-phase CG Snapshot creation breaks the Snapshot creation process into two phases. The first phase executes prechecks, triggers Snapshot creation, and starts a timer for a designated interval. If phase one completes successfully, the second phase must be invoked within the designated interval to commit to the consistency group Snapshot creation. This feature is only available with the ONTAP REST API.

ONTAP 9.12.1 enhances manageability through consistency group cloning, and the ability to add volumes to or remove volumes from a consistency group.

## Typical workflow

1. Create a job schedule or use the default.
2. Create a Snapshot policy, and then specify the job schedule.
3. Assign the Snapshot policy to a FlexVol volume or storage VM.



## Create a job schedule

The screenshot shows the ONTAP System Manager interface. On the left is a navigation menu with categories: DASHBOARD, INSIGHTS, STORAGE, NETWORK, EVENTS & JOBS, and PROTECTION. The PROTECTION section is expanded, showing Overview, Relationships, HOSTS, and CLUSTER. The main content area is titled 'Local Policy Settings' and contains a 'Schedules' panel. An arrow points to the 'Add' button in the top right of the Schedules panel. The 'Add Schedule' dialog box is open, showing a 'Name' field with '5min', a 'SCHEDULE NAME' field with 'Every 4 hour at 20 past on weekdays', and a 'SCHEDULE TYPE' section with 'Cron' selected. The 'CRON SCHEDULE' field contains '20 4,8,12,16,20 \* \* 1,2,3,4,5'. Below this, it lists the corresponding times: 'At 04:20 AM, 08:20 AM, 12:20 PM, 04:20 PM and 08:20 PM, only on Monday, Tuesday, Wednesday, Thursday, and Friday'. There are 'Cancel' and 'Save' buttons at the bottom right of the dialog.

```
::> job schedule cron create -name 4hrs -hour 4,8,12,16,20 -minute 20  
-dayofweek Monday,Tuesday,Wednesday,Thursday,Friday
```

NetApp 15 © 2023 NetApp, Inc. All rights reserved.

To create a job schedule by using ONTAP System Manager, first navigate to the Protection Overview page. Expand the Local Policy Settings panel, and then click the arrow in the upper-right corner of the Schedules panel. On the Schedules page, click **Add**.

Provide the new job schedule with a unique name and specify whether the scheduled type is interval or cron format.

If the schedule type is interval, specify the frequency with which the job is run in minutes, hours, or days. If the type is cron, specify when the job should run. Use a combination of minutes past the hour, hours of the day, days of the month, months of the year, and days of the week in numeric form. An asterisk represents all occurrences.

## Create a Snapshot policy

The screenshot shows the ONTAP System Manager interface. On the left is a navigation menu with sections: DASHBOARD, INSIGHTS, STORAGE, NETWORK, EVENTS & JOBS, PROTECTION (expanded), HOSTS, and CLUSTER. The 'PROTECTION' section is active, showing 'Local Policy Settings' and 'Snapshot Policies'. An arrow points to the 'Snapshot Policies' panel. In the center, the 'Policies' page is visible, with an '+ Add' button highlighted. A modal window titled 'Add Snapshot Policy' is open, showing the following details:

- POLICY NAME:** 4\_hours\_weekday
- POLICY SCOPE:** Cluster (unselected), Storage VM (selected)
- STORAGE VM:** svm1
- Schedules:**

Schedule Name	Maximum Snapshot Copies	SnapMirror Label
Every 4 hours at 20 past on week...	9	weekday_4h
daily	8	daily
weekly	13	weekly

At the bottom of the modal, there is a '+ Add' button.

```
::> volume snapshot policy create -vserver svm4 -policy 4_hours_weekday  
-schedule1 4hrs -count1 5 -prefix1 weekday_4h -enabled true
```

NetApp 16 © 2023 NetApp, Inc. All rights reserved.

You can use ONTAP System Manager to include the new job schedule in a Snapshot policy. On the Protection Overview page, click the arrow in the upper-right corner of the Snapshot Policies panel. On the Policies page, click **Add**.

To create a Snapshot policy, enter a unique policy name and specify whether the policy applies to cluster and node storage VMs or to data storage VMs. If you select Storage VM, which is the typical selection, you need to select the data storage VM name from the list. Then, click the **Add** button to add one or more job schedules to the policy. You can include as many as five schedules. For each schedule, select the schedule name and the number of Snapshot copies to retain. You can also specify a SnapMirror label to use if these Snapshot copies are replicated. The vaulting subsystem uses the SnapMirror label when you back up Snapshot copies to the vault destination. If an empty label ("" ) is specified, the existing label is deleted.

To specify the prefix name to use on Snapshot copies, you must use the CLI.

To add job schedules to an existing Snapshot policy, use the `volume snapshot policy add-schedule` command:

```
volume snapshot policy add-schedule -policy <snapshot_policy> -schedule <schedule_name> -  
count <integer> [-prefix <text>]
```

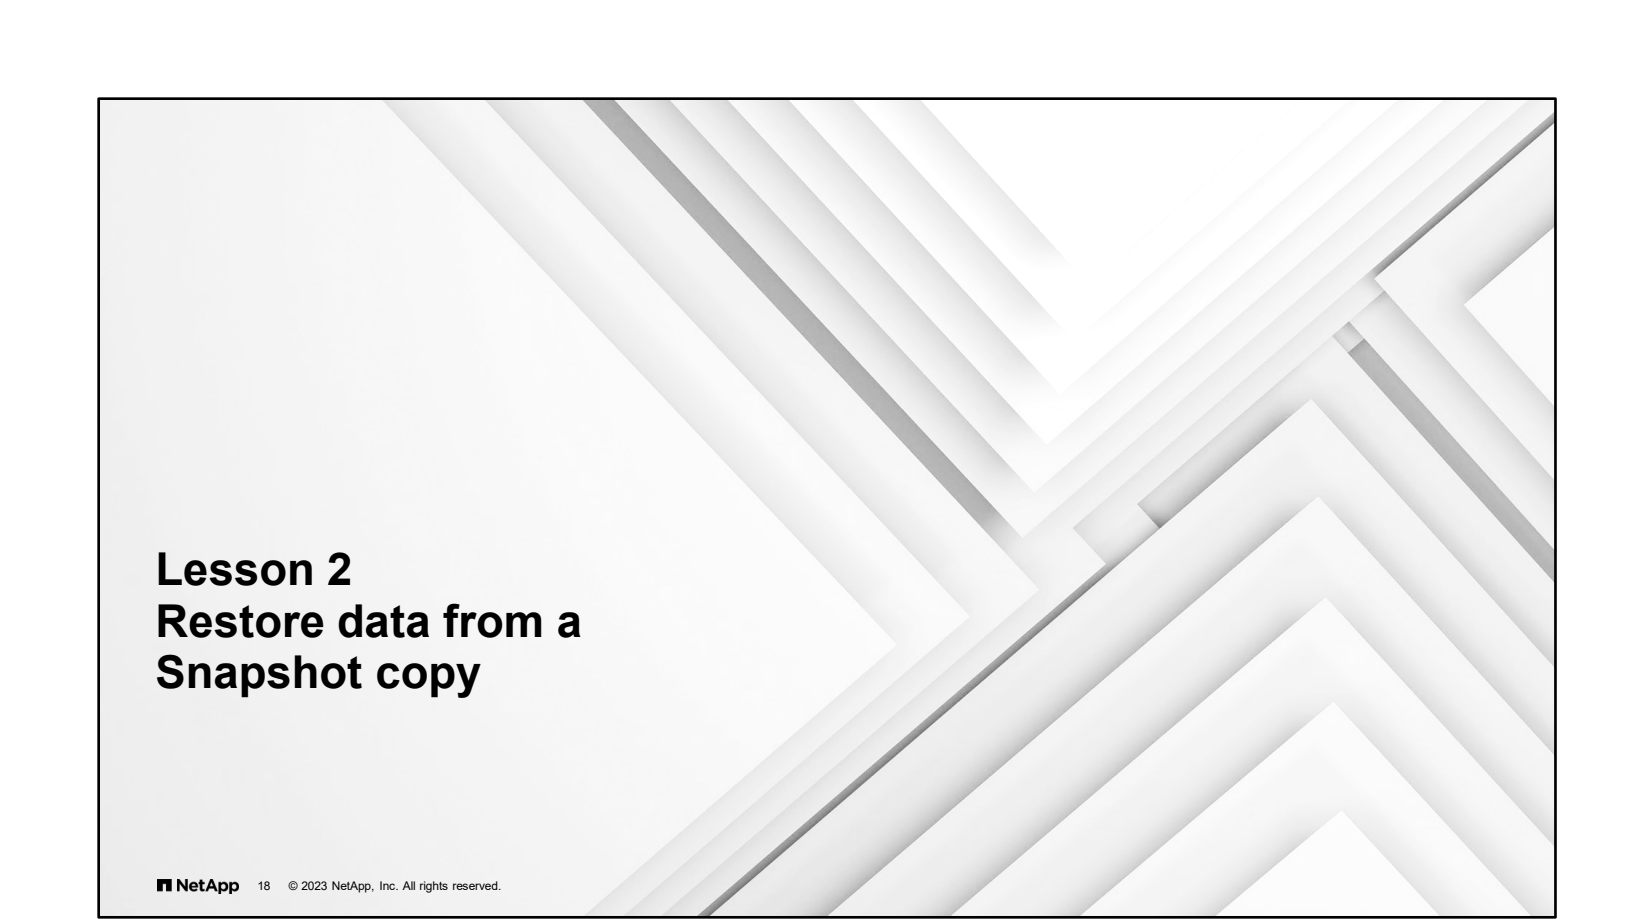
## Apply a Snapshot policy to a volume

The screenshot displays the ONTAP System Manager interface. On the left, the 'Volumes' table lists several volumes, with a context menu open for 'c1\_svm1\_vol1' showing the 'Edit' option highlighted. On the right, the 'Snapshot Copies (Local) Settings' panel is visible, showing the 'SNAPSHOT POLICY' dropdown set to '4\_hours\_weekday'. Below this, a table lists various snapshot policies with their schedules and labels. At the bottom, a terminal window shows the command: `::> vol modify -vserver svm4 -volume svm4_vol1002 -snapshot-policy 4_hours_weekday`.

Schedule Name	Maximum ...	Schedule	SnapMirror Label
Every 4 hours at 20 past on weekdays	0	At 04:20 AM, 08:20 AM, 04:20 PM and 08:20 PM, only on Monday, Tuesday, Wednesday, Thursday, and Friday, every month	-
daily	0	At 12:10 AM, every day	-
weekly	0	At 12:15 AM, only on Sunday, every month	-


NetApp 17 © 2023 NetApp, Inc. All rights reserved.

After you define the Snapshot policy, you can assign the policy to a storage VM or volume. On the System Manager Volumes page, click the **More** icon to the right of the volume name to display the menu, and then select **Edit**. On the Edit Volume page, scroll to the Snapshot Copies (Local) Settings section. From the list, select the Snapshot policy. The description is updated to reflect your choice. You can also enable automatic deletion of older Snapshot copies.




## **Lesson 2**

# **Restore data from a Snapshot copy**

 18 © 2023 NetApp, Inc. All rights reserved.

## Recovering data

Recover Snapshot data	Copy from a Snapshot copy	Use SnapRestore technology
<ul style="list-style-type: none"><li>• Copy data from Snapshot data.</li><li>• Use SnapRestore data recovery software.</li><li>• Use the Windows Previous Versions feature.</li></ul>	<ul style="list-style-type: none"><li>• Locate the Snapshot copy.</li><li>• Copy the file to the original location.</li><li>• Copy the file to a new location.</li></ul>	<ul style="list-style-type: none"><li>• Restore an entire volume.</li><li>• Quickly restore large files.</li></ul>

 19 © 2023 NetApp, Inc. All rights reserved.

You can use Snapshot copies to recover data in three ways:

- **Copy a file from a Snapshot directory:** To copy a lost or corrupted file from a Snapshot copy, navigate to the Snapshot directory on the client host. Locate the Snapshot copy that contains the correct version of the file. You can copy the file to the original location and overwrite existing data or copy the file to a new location.
- **Use the SnapRestore feature to recover data:** To revert a volume or file from a Snapshot copy, you need the SnapRestore license. You can revert a volume or file from the storage CLI or from the System Manager interface.
- **Use the Windows Previous Version feature:** Microsoft Windows users can use the Previous Versions feature to recover a previous version of a file from a Snapshot copy directly from the client host.

## Snapshot visibility to clients

- Snapshot directories are visible to NAS clients by default.
- Use commands to hide the Snapshot directories at the volume level or share level.
- Disable visibility of Snapshot directories for the volume:

```
::> vol modify -vserver svm4 -volume svm4_vol2 -snapdir-access false
```

- Disable visibility of Snapshot directories for an SMB share:

```
::> vserver cifs share properties remove -vserver svm4  
-share-name svm4_vol2 -share-properties showsnapshot
```

The `.snapshot` and `~snapshot` directories are visible to NFS and SMB clients by default.

**Note:** Show Hidden Files and Folders must be enabled on your Windows system.

CLI commands are available to control visibility from NAS clients of Snapshot directories on a volume.


You control access to `.snapshot` and `~snapshot` at the volume level by setting the `-snapdir-access` switch. You can also control access to `~snapshot` from SMB clients at the share level by using the `showsnapshot` share property.

NFS version 4 does not display the volume `.snapshot` directory. You can navigate through the `.snapshot` directory and recover previous versions of files, but the `.snapshot` directory will not appear in a listing of the directory contents.



## Snapshot view from a UNIX client

```
# ls /system/vol01/.snapshot
weekly.2023-01-15_0015  daily.2023-01-18_0010
daily.2023-01-19_0010  hourly.2023-01-19_0605
hourly.2023-01-19_0705  hourly.2023-01-19_0805
hourly.2023-01-19_0905  hourly.2023-01-19_1005
hourly.2023-01-19_1105  hourly.2023-01-19_1205
snapmirror.3_2147484677.2023-01-19_114126
```

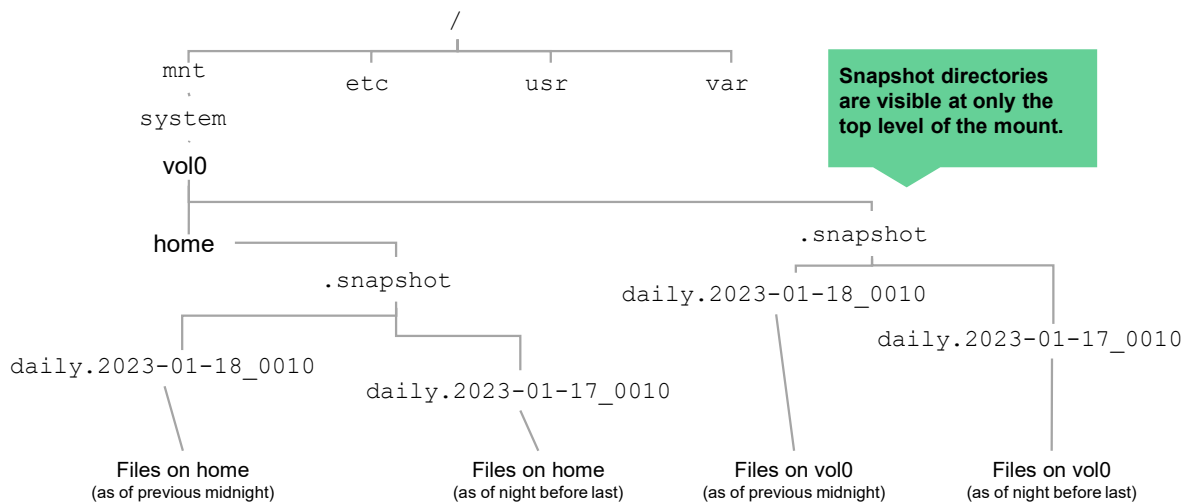
 21 © 2023 NetApp, Inc. All rights reserved.

Every volume in your file system contains a special Snapshot subdirectory. This subdirectory enables users to access earlier versions of the file system to recover lost or damaged files.

The Snapshot directory appears to NFS clients as `.snapshot`. The `.snapshot` directory is usually hidden. The directory is not displayed in directory listings, unless you use the `ls` command with the `-a` option.

When client Snapshot directories are listed, the timestamp is usually the same for all directories. To find the actual date and time of each Snapshot copy, use the `snap list` command on the storage system.

## Recovering files from the .snapshot directory of a UNIX host



NetApp 22 © 2023 NetApp, Inc. All rights reserved.

The `.snapshot` directory is at the root of a storage system volume.

In the example, the directory structure is shown for an NFS client that mounted `vol0` of a storage system to the mount point `/mnt/system` on the UNIX host.

The `home` directory and the `.snapshot` directory are visible at the root of the `vol0` mount.


You can open the `.snapshot` directory and access the files in the two Snapshot copies that are subdirectories of the `.snapshot` directory.

To restore a file from the `.snapshot` directory, rename or move the original file, and then copy the file from the `.snapshot` directory to the original directory.

## Recovering files from the ~snapshot directory

Snapshot copies are visible to Windows clients that have File Explorer configured to display hidden files.

Name	Date modified	Type	Size
~snapshot	10/6/2016 12:33 PM	File folder	
Financial	10/28/2016 3:38 PM	File folder	
Misc	10/28/2016 3:37 PM	File folder	
Projects	10/28/2016 3:37 PM	File folder	
7_JulyData	10/28/2016 3:44 PM	Text Document	0 KB
8_AugustData	10/28/2016 3:44 PM	Text Document	0 KB
file1	10/6/2016 12:33 PM	Text Document	0 KB
file2	10/28/2016 3:41 PM	Text Document	1 KB

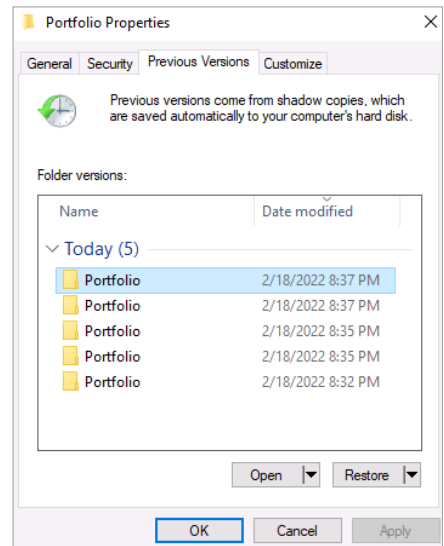
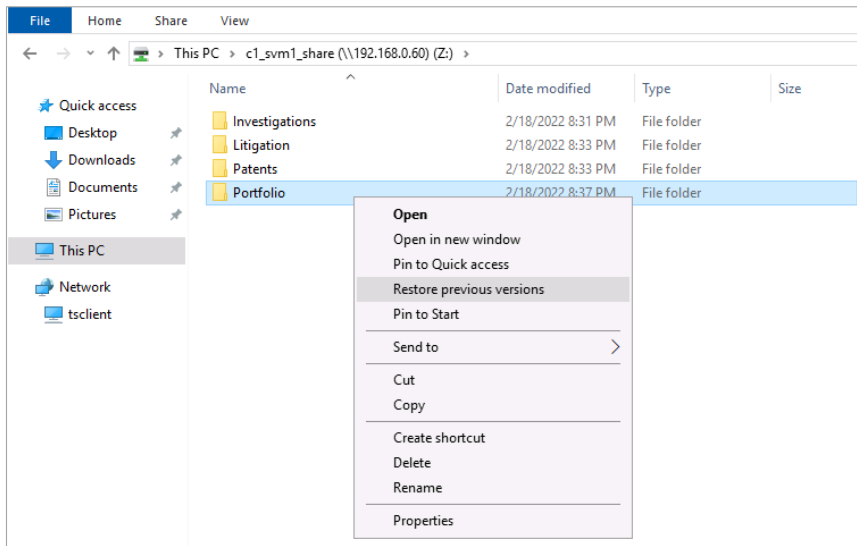
 23 © 2023 NetApp, Inc. All rights reserved.

Snapshot directories are hidden on Windows clients. To view them from clients that use SMB version 1, you must first configure File Explorer to display hidden files. Then navigate to the root of the CIFS share and find the directory folder. The subdirectory for Snapshot copies appears to CIFS clients as ~snapshot. Both automatic and manually created Snapshot copies are listed.

A user on a Windows client that uses SMB 2.x can, after connecting to a share, access the hidden ~snapshot directory by manually appending `\~snapshot` to the end of the share path. The hidden ~snapshot directory is accessible from two entry points: at the root of the share or at any junction point in the share space. The hidden ~snapshot directory is not accessible from non-junction subdirectories within the share.

To restore a file from the ~snapshot directory, rename or move the original file, and then copy the file from the ~snapshot directory to the original directory.


## Restoring previous versions in Windows



In Windows, right-click the file, and then select **Restore previous versions**.

## Snapshot copy technology

Restore from a Snapshot copy



- To restore a volume from Snapshot copy 1, use SnapRestore data recovery software.
- Snapshot copies that were created after Snapshot copy 1 are deleted.
- Unused blocks on drives are made available as free space.

NetApp 25 © 2023 NetApp, Inc. All rights reserved.

Suppose that after you create the Snapshot copy, the file or LUN becomes corrupted, which affects logical block C'. If the block is physically bad, RAID can manage the issue without recourse to the Snapshot copies. In the example, block C' becomes corrupted because part of the file is accidentally deleted. You want to restore the file.

To easily restore data from a Snapshot copy, use the SnapRestore feature. SnapRestore technology does not copy files. SnapRestore technology moves pointers from files in the good Snapshot copy to the active file system. The pointers from the good Snapshot copy are promoted to become the active file system pointers. When a Snapshot copy is restored, all Snapshot copies that were created after that Snapshot copy are destroyed. The system tracks links to blocks on the WAFL system. When no more links to a block exist, the block is available for overwrite and is considered free space.

Because a SnapRestore operation affects only pointers, the operation is quick. No data is updated, nothing is moved, and the file system frees any blocks that were used after the selected Snapshot copy was made. SnapRestore operations generally require less than 1 second. To recover a single file, the SnapRestore feature might require a few seconds or a few minutes.

## Reverting and restoring a volume

The screenshot shows the 'Volumes' page in NetApp System Manager. The volume 'c1\_svm1\_vol1' is selected. The 'Snapshot Copies' tab is active, displaying a table of snapshot copies. A context menu is open over the 'snap.2022-02-18\_203813' entry, with the 'Restore' option highlighted.

Name	Snapshot Copy Creation Time	Snapshot Restore Size
snap.2022-02-18_203813	Feb/18/2022 8:37 PM	86.7 MiB
snap.2022-02-18_203813	Feb/18/2022 8:37 PM	86.7 MiB
snap.2022-02-18_203813	Feb/18/2022 8:35 PM	86.6 MiB
snap.2022-02-18_203813	Feb/18/2022 8:35 PM	86.6 MiB

```
::> volume snapshot restore -vserver svm4 -volume svm4_vol2  
-snapshot svm4_vol2 snap
```

NetApp 26 © 2023 NetApp, Inc. All rights reserved.

Restoring a volume replaces the working copy of the volume with the Snapshot copy. The action results in a loss of all changes that were made since the Snapshot copy was created.

To view the Snapshot copies, click the volume name on the System Manager Storage Volumes page. Choose the Snapshot copy, and then select **Restore** from the menu. You are prompted for confirmation before the restore begins.

## Reverting and restoring a file

1. Verify that the volume is online and writable.
2. List the Snapshot copies in the volume.

```
::> volume snapshot show -vserver svm4 -volume svm4_vol2
```

3. Notify network users about the reversion.
4. Identify the names of the Snapshot copy and the file to restore and initiate the reversion.

```
::> volume snapshot restore-file -vserver svm4 -volume svm4_vol2  
-snapshot svm4_vol2_snap -path /svm4_vol2/myfile.txt
```

 27 © 2023 NetApp, Inc. All rights reserved.


After you complete the steps to revert a file, ONTAP software displays a warning message and prompts you to confirm your decision to revert the file. Press **Y** to confirm that you want to revert the file. If you do not want to proceed, press **Ctrl+C** or press **N**.

If you confirm that you want to revert a file in the active file system, the version in the Snapshot copy overwrites the file.

## SnapRestore technology versus copying

If a file is large (such as a database), you should use SnapRestore technology to revert instead of copying the file:

- Copying requires double the storage and time.
- Reverting saves time and reinstates the data.
- For reliability, NetApp recommends SnapRestore technology rather than alternative technologies.



Whether you restore by copying files from a Snapshot directory or from tape, copying large quantities of data can be time consuming. Instead, use the SnapRestore function to restore by reverting the volume or file.



## Snapshot automatic delete

- Use the `volume snapshot autodelete modify` command to modify the autodelete policy settings.
- Enable automatic Snapshot copy deletion on a volume:

```
::> volume snapshot autodelete modify -vserver svm4 -volume svm4_vol2  
-enabled true
```

- Trigger automatic deletion of the oldest unlocked Snapshot copies when the volume threshold is exceeded:

```
::> volume snapshot autodelete modify -vserver svm4 -volume svm4_vol2  
-trigger volume -commitment try -delete-order oldest_first
```

Snapshot automatic deletion determines when or whether Snapshot copies are automatically deleted. The autodelete policy settings determine when automatic deletion of volume and LUN Snapshot copies, NVMe Express (NVMe) namespace, or file clones is triggered. Automatic deletion of Snapshot copies and LUN, NVMe namespace, or file clones is useful when you want to automatically reclaim space in the volume.

The `volume snapshot autodelete modify` command has many options that you can use to modify the autodelete policy settings. Some of the most used options are the following:

- `enabled [true|false]`: Enable or disable automatic deletion.
- `commitment [try|disrupt|destroy]`: Determine which types of Snapshot copies are selected for deletion.
  - When set to `try`, the Snapshot copies that are not locked by any application are deleted.
  - When set to `disrupt`, the Snapshot copies that are not locked by clones are deleted (unless they are marked as preserved). In the `disrupt` mode, the Snapshot copies that are locked by SnapMirror software or Volume Move can be deleted, causing the transfer to be aborted.
  - When set to `destroy`, all the Snapshot copies are deleted.
- `defer-delete [scheduled|user_created|prefix|none]`: Determine which types of Snapshot copies are deleted last.
- `delete-order [newest_first|oldest_first]`: Delete oldest or newest Snapshot copies first.
- `trigger [volume|snap_reserve]`: Select the condition that starts automatic deletion.
  - Setting this option to `volume` triggers automatic deletion when the volume reaches the threshold capacity, and the volume Snapshot reserve is exceeded. The threshold capacity ranges between 85 and 98 percent, depending on the volume size.
  - Setting the option to `snap_reserve` triggers automatic deletion when the capacity of the Snapshot reserve is reached.



## Complete an exercise

Module 8  
Data protection

### Managing Snapshot copies

- Access your lab equipment.
- Open your Exercise Guide to Module 8.
- Complete Exercise 1.
- Share your results.


This exercise requires approximately  
**15 minutes.**

See the instructions in your Exercise Guide.



## Lesson 3

### Back up and replicate data

 31 © 2023 NetApp, Inc. All rights reserved.

## Disaster recovery and business continuance

There are two reasons for backing up and replicating data:

- **Disaster Recovery:** The ability to recover data that has been deleted, corrupted, infected by a virus, or physically lost because of a disaster
  - NDMP backup
  - SnapMirror vaults
- **Business Continuance:** Using up-to-date replicas of data to keep a business operating despite a disaster
  - SnapMirror software
  - MetroCluster cluster configuration

Both reasons are constrained by recovery time objectives (RTOs) and recovery point objectives (RPOs).

Regardless of how resilient a storage system is, some events can result in the corruption or loss of data. To reduce the effect of these events, make backup copies and replicas of the data. Whether you do both or only one is determined by your business needs. The two primary business needs are disaster recovery and business continuance.

For disaster recovery, the primary goal is the ability to restore the data. The amount of time required to recover the data is secondary. Disaster recovery is the least expensive option, so companies with limited budgets or less reliance on data often choose this option. The two primary ONTAP features that are used to create disaster recovery backups are NDMP backups and SnapMirror vaults.

For business continuance, the primary goal is for the company to continue doing business while recovery from a disaster. Business continuance is expensive because it generally requires the duplication of the production compute, storage, and network infrastructure. SnapMirror software is the primary ONTAP feature that is used to accomplish business continuance. MetroCluster configurations are hardware and software solutions to provide business continuance.

### RTO

The Recovery Time Objective (RTO) is the amount of time within which the service, data, or process must be made available again to avoid undesirable outcomes. Essentially, the RTO specifies how long the business can tolerate an outage.

### RPO

The Recovery Point Objective (RPO) is a point to which data must be restored or recovered to be acceptable to the acceptable data loss policy of the organization. Essentially, the RPO specifies how much data the business can tolerate to lose. Data backups incur a cost. The more frequent the backups, the more cost is incurred. If a system is scheduled to perform a backup once every 24 hours and the disaster occurs before the next backup finishes, 24 hours of new and changed data was not backed up and might not be recoverable.

## NDMP and SMTape backups

**NDMP** is the industry standard protocol that third-party backup applications use to back up data to physical or virtual tape devices.

- Backup applications can use NDMP to perform a Snapshot copy-based backup of a volume, directory tree, or file.
- NDMP supports baseline, differential, and incremental backups.

**SMTape** is a Snapshot copy-based solution that backs up volumes to tape.

- SMTape backs up and restores only entire volumes.
- SMTape is used primarily to back up Snapshot copies and to seed SnapMirror destination volumes.

NDMP is an industry standard protocol for communication between storage devices and backup devices, such as tape drives. Third-party backup applications can use the NDMP protocol that is embedded in ONTAP software to back up and restore the contents of FlexVol volumes.


The NDMP protocol is file based, so you can back up or restore an entire volume, a specific directory tree, or an individual file. NDMP understands file-modification times and can therefore perform both baseline and incremental backups.

SMTape uses the SnapMirror engine, which is discussed later, to back up blocks of data rather than files. (Think of SMTape as a SAN protocol and NDMP as a NAS protocol.) Although you can use SMTape for daily backup, the feature is most frequently used to seed a remote SnapMirror destination for large volumes. Rather than send a large amount of data over the network to the destination, SMTape creates a set of tapes that is shipped to the destination and recovered locally onto the destination volumes. SnapMirror replication is then initiated, and only the blocks that are new or changed since the tapes were created are transferred over the network.

## SnapMirror vault relationships

Previously known as SnapVault

- SnapMirror vault relationships create read-only backup copies on a destination volume.
- SnapMirror vault relationships are frequently used to back up multiple production clusters to a few remote high-capacity disaster recovery clusters.
- Use SnapMirror software instead of NDMP-based backups for the following reasons:
  - NDMP incrementals back up changed files. SnapMirror software backs up only changed blocks.
  - SnapMirror vaulting can store hundreds of daily backups, often for lower costs than removable media.
  - SnapMirror vaulting provides efficient use of WAN resources for off-site backups.



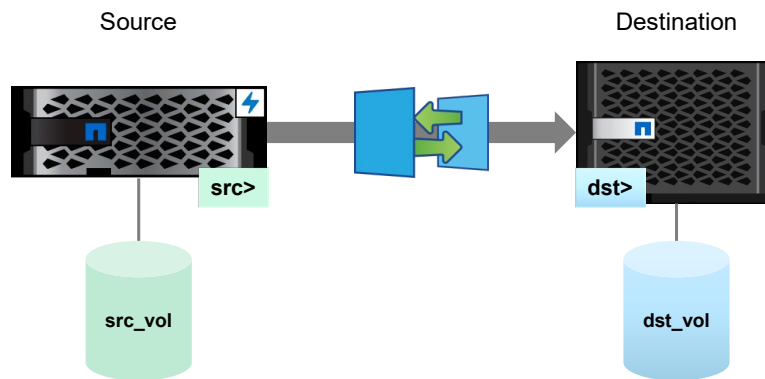
**NetApp** 34 © 2023 NetApp, Inc. All rights reserved.

Although a SnapMirror vault relationship is like NDMP backups to tape, the relationship has several significant advantages. The first advantage is that the backed-up data is always online and available. The second advantage is that after the initial baseline transfer of the entire volume contents, subsequent SnapMirror backups contain only the data blocks that have changed since the previous backup was performed. The size of a SnapMirror backup is considerably smaller than an NDMP backup of the same content, resulting in significant savings of resources. SnapMirror preserves storage efficiencies by transferring deduplicated compressed data.

The economies of scale often make it less expensive than using tapes or other removable media. Tapes have the cost of media, administrative overhead to load and remove them from tape libraries, and ongoing expenses for the physical transport and storage costs at archive facilities like Iron Mountain.

## SnapMirror technology

SnapMirror technology enables the mirroring of volumes to other local or geographically remote systems.



**NetApp** 35 © 2023 NetApp, Inc. All rights reserved.

SnapMirror technology is an ONTAP feature that enables you to replicate data for business continuance. SnapMirror technology enables you to replicate data from specified source volumes to specified destination volumes.


You can use SnapMirror technology to replicate data within the same storage system or between storage systems.

After the data is replicated to the destination storage system, you can access the data on the destination to perform the following actions:

- Provide users immediate access to mirrored data if the source fails
- Restore the data to the source to recover from disaster
- Balance resource loads
- Back up or distribute the data to remote sites

## SnapMirror features and benefits

- The SnapMirror destination is a replica. Changes to the source are mirrored to the destination.
- Updates to the destination can be made frequently because only new and changed data blocks, rather than entire files, are sent.
- SnapMirror technology uses deduplication and compression and supports dual paths to keep latency low and network bandwidth needs to a minimum.




**NetApp** 36 © 2023 NetApp, Inc. All rights reserved.

Unlike an NDMP backup, which requires a backup window, SnapMirror replicas can be created and updated as frequently as every 5 minutes. Only the new or changed blocks in a file, rather than the entire file, are sent to the destination. Deduplication and compression provide further efficiencies.



## SnapMirror foundational technology

- SnapMirror Asynchronous
- SnapMirror Synchronous
- SnapMirror vault relationships
- SVM DR
- SnapMirror Cloud
- SnapMirror Business Continuity
- S3 SnapMirror

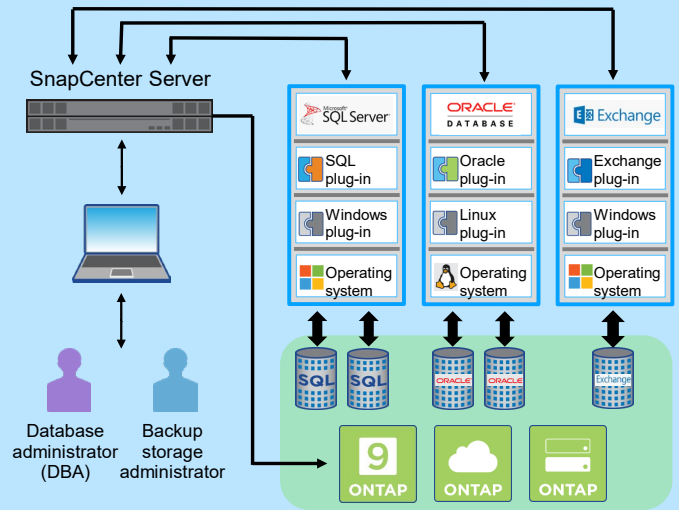


SnapMirror is a key NetApp technology. You can learn more by enrolling in the *ONTAP Data Protection* courses.

## NetApp SnapCenter software

- Application-consistent data protection for applications, databases, host file systems, and VMs running on NetApp ONTAP based systems anywhere in the hybrid cloud
- Application plug-ins for:
  - Microsoft Exchange Server
  - Microsoft SQL Server
  - Oracle databases on Linux
  - SAP HANA databases
  - VMware vSphere
- Available downloadable samples of plug-ins for other applications and databases

NetApp 38 © 2023 NetApp, Inc. All rights reserved.



Applications like databases cannot be easily recovered by simply copying their files from a backup. The data must be *quiescent* during the backup, and the state of the application must be preserved to create what is known as an *application-consistent* backup. NetApp SnapCenter software automates the work of creating and recovering application-consistent backups or replicas.

SnapCenter software is a simple, centralized, scalable platform that provides application-consistent data protection for applications, databases, host file systems, and VMs that run on ONTAP systems anywhere in the hybrid cloud. You can use SnapCenter software to manage data protection operations on multiple ONTAP storage systems. These systems can be on-premises FAS and AFF systems, off-premises systems, NetApp Cloud Volumes ONTAP systems, or ONTAP Select software. These systems are seamless to SnapCenter software and work the same, regardless of where the data resides.

Data protection is supported for Microsoft Exchange Server, Microsoft SQL Server, Oracle databases on Linux, and SAP HANA databases running on ONTAP systems. The SnapCenter Plug-in for VMware vSphere, which enables backup and recovery of VMs and datastores, is also available. SnapCenter software also supports a custom plug-in capability. You can use this capability to write your own plug-in for any application or database that SnapCenter software does not support. As examples, the plug-ins for DB2, MySQL, and MongoDB are available for you to download from the NetApp Automation Store. These plug-ins are not part of the SnapCenter software installation and are community supported.

Backup storage and database administrators all access SnapCenter software through a shared UI. When a manual or scheduled backup is triggered, SnapCenter Server contacts the appropriate plug-in that is installed on the application server or servers. The plug-in pauses the application I/O, while SnapCenter Server initiates the creation of a Snapshot copy on the ONTAP system. The ONTAP system creates a Snapshot copy of the volumes where the application data resides. The ONTAP system notifies SnapCenter Server when the process is complete, which typically takes less than one second. SnapCenter Server then directs the application plug-in to resume the application I/O. The Snapshot copies can then be automatically replicated to another ONTAP system, using SnapMirror technology.

## MetroCluster configuration


- A MetroCluster configuration geographically separates the partners in high-availability (HA) pairs.
- If a disaster damages the physical storage hardware or network access to the hardware, the cluster can continue to serve data.
- MetroCluster configuration is not an add-on feature or upgrade. Clusters must be physically installed and configured in a MetroCluster configuration.



Earlier, the course discussed cluster configurations and MetroCluster cluster configurations. In a standard cluster configuration, the partners in a high-availability (HA) pair are often physically located in the same cabinet. If the cabinet is destroyed by a fire or earthquake, the entire HA pair is lost. In a MetroCluster cluster configuration, the HA partners are geographically separated to reduce the likelihood of a disaster affecting both partners. Because of the complexity and physical requirements of a MetroCluster configuration, this type of cluster configuration must be decided when the cluster is purchased. MetroCluster configurations are popular in regions where geography, politics, and national borders make the use of remote disaster-recovery locations difficult.




# Lesson 4 Compliance

 40 © 2023 NetApp, Inc. All rights reserved.

## SnapLock

- SnapLock software is a high-performance compliance solution for organizations that use WORM storage to retain files in unmodified form for regulatory and governance purposes.
- Files are locked from modification for an administrator-defined length of time.
  - Files in Snapshot copies and SnapMirror destinations are also locked until the time limit for all files expires.
  - With SnapLock Enterprise, administrators can delete files before the time limit expires.
  - With SnapLock Compliance, administrators cannot delete files until after the time limit expires.
- SnapLock software requires a license.
- SnapLock software can be used in conjunction with storage encryption.



Recommended practice: Learn and practice using SnapLock software on a simulator before implementation, because some mistakes are irreversible.

**NetApp** 41 © 2023 NetApp, Inc. All rights reserved.

SnapLock is a licensed feature that enables you to lock files so that they cannot be altered in any way for a predetermined amount of time. Companies that process insurance, mortgages, and other legal and financial documentation use SnapLock technology to ensure that digital files cannot be altered and therefore become legally indefensible.

With ONTAP 9.9.1 software, the storage efficiency capabilities for SnapLock volumes and aggregates were extended to include data compaction, cross-volume-deduplication, adaptive compression, and Temperature Sensitive Storage Efficiency (TSSE), enabling even greater space savings.


With ONTAP 9.10.1 software, SnapLock volumes can co-exist with standard FlexVol volumes in the same data aggregate. This alleviates the need to dedicate drives to SnapLock aggregates making SnapLock compliance more attractive on smaller storage systems.

With ONTAP 9.11.1 software, SnapLock can now be enabled on FlexGroup volumes as well as FlexVol volumes.

The SnapLock feature has a learning curve, and mistakes can result in files or entire aggregates that cannot be deleted until the lock expires. Locks are often set for many years (up to 100 years)! If you plan to implement SnapLock software or take over administration of a cluster that uses SnapLock software, practice on a simulation before making any changes to the production cluster.

## Tamperproof Snapshot copies

- Snapshot copies can be locked even on volumes that do not have SnapLock enabled.
- Snapshot copies cannot be deleted until the expiration time is reached.
- To protect volume Snapshot copies from tampering:
  - The SnapLock license must be installed on the cluster.
  - The SnapLock cluster compliance clock must be initialized.
  - Snapshot locking must be enabled on the volume.
  - A retention time must be set on the snapshot.



**NetApp** 42 © 2023 NetApp, Inc. All rights reserved.


Beginning with ONTAP version 9.12.1, you can now lock a Snapshot copy on a non-SnapLock volume to provide protection from ransomware attacks. Locking Snapshot copies helps ensure that they are not deleted accidentally or maliciously.

Use the `–snapshot-locking-enabled true` option of the `volume create` or `volume modify` commands to enable tamperproof snapshots on new or existing volumes. Then specify the Snapshot expiration period with the `–expiry-time` option when creating a volume snapshot manually or include the Snapshot retention period in the volume snapshot policy.


The retention period for locked Snapshot copies takes precedence over the Snapshot copy keep count, which means the keep count limit is not honored if the Snapshot copy retention period for locked Snapshot copies has not expired.

In a SnapMirror relationship, you can set a retention period on a mirror-vault policy rule, and the retention period is applied for Snapshot copies replicated to the destination if the destination volume has Snapshot copy locking enabled. The retention period takes precedence over keep count; for example, Snapshot copies that have not passed their expiry will be retained even if the keep count is exceeded.

You can restore a locked Snapshot copy with the `volume snapshot restore` command only if the locked Snapshot copy is the most recent. If there are any unexpired Snapshot copies later than the one being restored, the Snapshot copy restore operation fails.




## Lesson 5 Storage encryption

 43 © 2023 NetApp, Inc. All rights reserved.

## What is NetApp Storage Encryption?

- NetApp Storage Encryption (NSE) is an ONTAP feature that provides support for self-encrypting drives (SEDs).
- SEDs protect data at rest (when the drive is powered off).
- NSE manages the authorization process with a key management server to grant storage controllers access to the encrypted data on the drives.
- The encryption process is transparent to end users and has a minimal effect on performance.



**NetApp** 44 © 2023 NetApp, Inc. All rights reserved.

A standard drive stores data unencrypted. If the drive is lost or stolen, the data is vulnerable to unauthorized access. Theft of storage devices is a real threat for financial, healthcare, and government institutions. To solve this issue, drive manufacturers create drives with built-in encryption, called self-encrypting drives (SEDs). All data that is written to a SED is encrypted and can be read only by systems that have successfully completed an authentication process with a key management server. NetApp Storage Encryption (NSE) is an ONTAP feature that supports SED use.

After you enable NSE and create an authorization key, a FAS or AFF system must send a password to the key server to enable access to the encrypted drives. When the storage system is running and the drives are powered on, the process is transparent to end users and performance is barely affected. Only when the drives are offline or not connected to an authenticated system is the data indecipherable. See the Interoperability Matrix Tool (IMT) for a list of supported key servers.


One caveat when using NSE: All drives that are attached to a standalone system or an HA pair of systems must be SEDs. Mixing encrypting and non-encrypting drives is not supported. Multinode clusters do support mixing HA pairs that use SEDs and HA pairs that use standard drives.

All NVMe SSDs are self-encrypting and can use the same onboard and Key Management Interoperability Protocol (KMIP) key servers as NSE drives. However, NVMe SSD encryption is not Federal Information Processing Standard (FIPS)-140-2 level 2 compliant, and NVMe SSDs cannot be mixed with NSE drives in the same HA pair.



## Onboard key management

- Onboard Key Manager is a less-expensive alternative to external KMIP servers. With onboard key management, the storage servers manage their own authentication to the NSE drives.
- Do not use Onboard Key Manager if any of the following conditions are true:
  - Your storage systems must comply with Federal Information Processing Standard (FIPS) 140-2 or the OASIS KMIP standard.
  - You need a centralized, multicluster solution. (OKM works only for the cluster that hosts the keys.)
  - Your business requires the added security of storing authentication keys separately from encrypted data.





A low-cost alternative to key management servers is to enable Onboard Key Manager. With Onboard Key Manager, the storage systems manage the authentication keys that unlock the NSE drives. This approach helps to ensure that encryption protects data at rest. However, Onboard Key Manager is not compliant with FIPS and does not work with more than one cluster. OKM also has a physical security flaw: Unauthorized users with physical access to the storage systems and disks can access the encrypted data.

You are prompted to enable Onboard Key Manager when you first provision storage by using the ONTAP System Manager.

## NetApp Volume Encryption

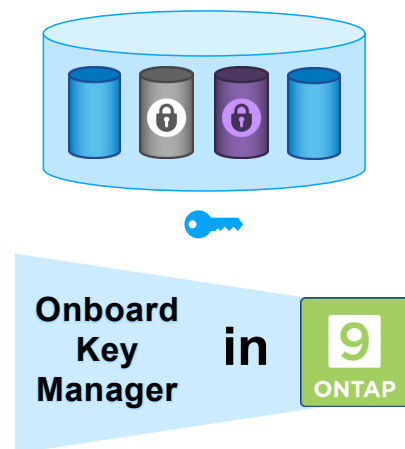
- NetApp Volume Encryption (NVE) is a software-based, data-at-rest encryption solution:
  - Encrypts sensitive data without relying on NSE drives
  - Uses Advanced Encryption Standard (AES)-256 encryption
  - Requires a license
- Each data volume has a unique encryption key:  
Decide which volumes to encrypt and which to leave unencrypted.
- Encryption requires zero management:
  - Snapshot copies and FlexClone volumes are also encrypted.
  - If you use a KMIP server, ONTAP software automatically uploads the encryption key to the server when you encrypt a volume.



NSE has some caveats that can be problematic in environments where not all data needs to be encrypted. NetApp Volume Encryption (NVE) provides a flexible solution. Using onboard key management, you can select a volume to encrypt and assign the volume a unique encryption key. Because the data blocks in the volume are encrypted, the encryption follows the blocks into Snapshot copies and FlexClone volumes.

## NetApp Aggregate Encryption

- You can use NetApp Aggregate Encryption (NAE) to assign encryption keys to the containing aggregate.
- Volumes that you create in the aggregate are encrypted by default, using the aggregate encryption keys.
- You can override the default encryption keys or disable encryption when you create a volume.



NetApp 47 © 2023 NetApp, Inc. All rights reserved.

Typically, every encrypted volume is assigned a unique key. Starting with ONTAP 9.6 software, you can use NetApp Aggregate Encryption (NAE) to assign keys to the aggregate that contain the volumes to encrypt.

You must use aggregate-level encryption if you plan to perform inline or background aggregate-level deduplication. Aggregate-level deduplication is otherwise unsupported by NVE.

NVE and NAE volumes can coexist on the same aggregate.

You can use the `volume move` command to convert an NVE volume to an NAE volume and vice versa. You can also replicate an NAE volume to an NVE volume.

## Additional storage security features



### Use data encryption by default.

Data is encrypted automatically when key management is configured.




### Honor “right to be forgotten.”

Manage new data-compliance regulations better with crypto-shredding of data through secure purge.



### Protect systems in transit.

Protected controller reboot and secure Unified Extensible Firmware Interface (UEFI) boot prevent unwanted access of systems outside the data center.



### Worry less about cloud security.

NVE support for NetApp Cloud Volumes ONTAP provides FIPS 140-2 certified encryption in the cloud.



More info in Addendum

**NetApp** 48 © 2023 NetApp, Inc. All rights reserved.

Starting with ONTAP 9.7 software, newly created aggregates and volumes are encrypted by default when the NVE license is installed, and onboard or external key management is configured.

Secure purge shreds data in the volume and Snapshot copies to meet data-compliance regulations.

For systems that move between data centers, protected controller reboot prevents unauthorized access if the storage hardware is stolen.

NVE also works to protect data in the cloud.

Learn more about secure purge and secure boot in the module addendum.




## Complete an exercise

Module 8  
Data protection


### Encrypting a Volume

- Access your lab equipment.
- Open your Exercise Guide to Module 8.
- Complete Exercise 2.
- Share your results.

This exercise requires approximately  
**15 minutes.**


 49 © 2023 NetApp, Inc. All rights reserved.

See the instructions in your Exercise Guide.



## Lesson 6

# Ransomware protection

 50 © 2023 NetApp, Inc. All rights reserved.

## Ransomware attacks

- There are two types of ransomware attacks:
  1. Denial of service to files by encrypting data  
The attacker withholds access to this data unless a ransom is paid.
  2. Theft of sensitive proprietary data  
The attacker threatens to release this data to the public domain unless a ransom is paid.
- Implement ONTAP ransomware protection in concert with your other antivirus defenses.



There are two forms of ransomware commonly in use today, data exfiltration and denial of service through encryption.

With data exfiltration, an attacker gains access to a company's IT systems, moving sensitive data to an unknown location outside of the company, and then threatens to publicly release that data unless a ransom is paid.

The more common version of ransomware is called denial of service ransomware. In this ransomware strategy, an attacker gets you to inadvertently download an encryption program (malware). After it is installed, that malware encrypts all the local client files and every single file that it can on NFS or SMB shares on the corporate network. After the files are encrypted, the original files are deleted and there is no longer any way for you to access the data in the files. You can still see the files because they are still on your network, but you cannot access them because the attacker has encrypted them. The attacker then demands that you pay a ransom to obtain the decryption key so that you can regain access to your data.

ONTAP ransomware protection addresses denial of service type attacks, with an anti-ransomware detection mechanism that is based on identification of the incoming data as encrypted.

The Cloud Secure feature of NetApp Cloud Insights performs user behavior analytics to identify suspect user activity. Cloud Secure can be used to identify intellectual property theft.


ONTAP ransomware protection and Cloud Secure should be implemented as part of a larger data security defense strategy that includes your firewalls and identity certificates.

## ONTAP Autonomous Ransomware Protection

### Last line of defense

- ONTAP uses workload analysis in NFS and SMB environments to proactively detect abnormal activity.
- Anti-ransomware can be enabled on selective volumes.
- When an attack is suspected, anti-ransomware does the following:
  - Creates new Snapshot backups
  - Alerts you to the attack



 52 © 2023 NetApp, Inc. All rights reserved.

The ONTAP anti-ransomware feature uses workload analysis in NAS (NFS and SMB) environments to proactively detect and warn about abnormal activity that might indicate a ransomware attack. When an attack is suspected, anti-ransomware also creates new Snapshot backups, in addition to existing protection from scheduled Snapshot copies.

The anti-ransomware feature starts in learning mode. NetApp recommends a period of at least 30 days so that the machine learning has a chance to understand the typical workloads on the NAS volumes. When anti-ransomware is put into active mode, it starts looking for abnormal volume activity that might potentially be ransomware.

If abnormal activity is detected, an automatic Snapshot copy is immediately taken, which provides a restoration point as close as possible to the file infection. Simultaneously, an automatic alert is generated that enables administrators to see the abnormal file activity so that they can determine whether the activity is indeed malicious and take appropriate action. If the activity was an expected workload, they can easily mark it as a false positive. The anti-ransomware machine learning notes the change in workload and no longer flags it as a potential attack. In addition, the feature does not disrupt I/O in any way. Instead, it provides administrators with native analytics, insights, and data recovery capabilities for unprecedented on-box ransomware detection.

No ransomware detection or prevention system can completely guarantee safety from a ransomware attack. While it is possible that an attack might go undetected, NetApp ransomware protection acts as an important additional layer of defense if antivirus software has failed to detect an intrusion. NetApp anti-ransomware can detect the spread of most ransomware attacks after only a small number of files are encrypted, act automatically to protect data, and alert you that a suspected attack is happening.



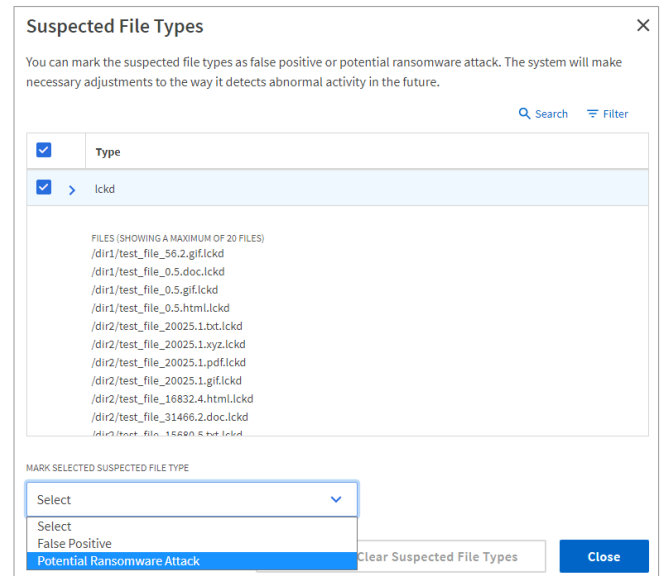
## Respond to abnormal activity

When the ONTAP anti-ransomware detects abnormal volume activity, it issues a warning notice with a list of suspected files.

You can respond to the warning in two ways:

1. **False positive**  
The identified file type is expected in your workload and can be ignored.
2. **Potential ransomware attack**  
The identified file type is unexpected in your workload and should be treated as a potential attack.

Normal monitoring resumes after the notices have been cleared.



**NetApp** 53 © 2023 NetApp, Inc. All rights reserved.

When the ONTAP anti-ransomware detects any combination of high data entropy, abnormal volume activity with data encryption, and unusual file extensions, it issues a warning with a list of suspected files. By examining the actual file contents, ONTAP can detect ransomware within the data instead of only relying on user file access patterns or extension names.

When the warning is issued, you can respond by marking the file activity in one of two ways:


1. **False positive**  
The identified file type is expected in your workload and can be ignored.
2. **Potential ransomware attack**  
The identified file type is unexpected in your workload and should be treated as a potential attack.

In both cases, normal monitoring resumes after updating and clearing the notices. ONTAP anti-ransomware records your evaluation, updates the logs with the new file types, and uses them for future analysis.

However, in the case of a suspected attack, you must determine whether it is an attack, respond to it if it is, and restore protected data before clearing the notices.

## NetApp Cloud Insights with Cloud Secure

- Use NetApp Cloud Insights to monitor your hybrid multi-cloud IT infrastructure.
- Cloud Secure helps you in the following ways:
  - Ensures corporate compliance by monitoring user data access patterns
  - Detects and stops ransomware attacks
  - Protects intellectual property from theft by malicious users
- ONTAP anti-ransomware activity can be monitored in Cloud Secure for a single unified view.



Cloud Insights is a SaaS infrastructure and service monitoring solution that works for on-premises, private cloud, and public cloud environments including AWS, Azure, and Google Cloud. Cloud Secure, a feature of NetApp Cloud Insights, analyzes data access patterns to identify risks from ransomware attacks.

Cloud Secure helps to protect your data with actionable intelligence on insider threats. It provides centralized visibility and control of all corporate data access across hybrid cloud environments to ensure that security and compliance goals are met.

1. **Visibility:** Gain centralized visibility and control of user access to critical corporate data stored on premises or in the cloud. Cloud Secure uniquely operates on both cloud and on-premises storage systems to give you real-time alerts of malicious user behavior.
2. **Protection:** Protect organizational data from being misused by malicious or compromised users through advanced machine learning and anomaly detection. Cloud Secure alerts you to any abnormal data access through advanced machine learning and anomaly detection of user behavior.
3. **Compliance:** Ensure corporate compliance by auditing user data access to critical corporate data stored on premises or in the cloud.

Like ONTAP ransomware protection, Cloud Secure can initiate a Snapshot backup of ONTAP volumes whenever an attack is detected.

## Recover from a ransomware attack

Perform the following steps when a ransomware attack is detected:

1. **Contain and isolate**  
Prevent the ransomware from spreading and re-infecting by isolating the infected clients from the network.
2. **Repair and patch**  
Use antivirus and anti-malware software to clean the infected client hosts.  
Install software security patches to prevent re-infection.
3. **Restore and recover**  
Identify backups that are not compromised.  
Restore damaged files from backup.



Your first instinct after a ransomware attack might be to instantly recover your data. You can certainly do this, but if you do not take other steps to make sure that the ransomware does not come back, you are likely to end up being reinfected and the effort will waste valuable time.

There are three key steps to remediate your environment properly and holistically from ransomware infection.

### Step 1: Contain and isolate

Ransomware infections typically start at a client host. To contain the outbreak, you must identify and isolate the infected clients by disconnecting them from the network.

Identification of infected client hosts can be accomplished by whatever means you typically use to monitor file share access. For ONTAP based systems with SMB shares, there are two CLI options: using `vserver cifs session show` and `vserver locks show`. For ONTAP based systems with NFS exports, the `vserver locks show` can be used for monitoring access to an NFS share.

### Step 2: Repair and patch

When you have identified the infected clients, it is highly recommended that you use local antivirus and anti-malware software to clean them.

After a machine has been isolated and cleaned of the infection, then it is time to determine how they were infected and install any recent software patches. Many ransomware attacks exploit known security flaws for which fixes are already available.

### Step 3: Recover and restore

Identifying good backups that are not infected with ransomware is a crucial part of this step. If you are using NetApp ONTAP Snapshot copies, they are read-only and immutable from ransomware. Snapshot copies named "Anti\_ransomware\_backup" are created when anti-ransomware detects a potential attack. You can restore data from these anti-ransomware copies or other Snapshot copies. If you are using backups other than NetApp Snapshot copies, make sure that they have not been infected with the ransomware before the restore.


The larger the amount of data that has been encrypted by the ransomware attack, the longer the restore process takes before you can regain access to all your data. NetApp ONTAP Snapshot copies provide nearly instant recovery.

With the knowledge of the affected files and the time of attack, it is possible to selectively recover the affected files from various Snapshot copies, rather than simply reverting the whole volume to one of the Snapshot copies.



# Knowledge check

Module 8: Data protection

 56 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

**True or false: Data can be written to a Snapshot copy.**

- a. true
- b. false

## Module summary

This module focused on enabling you to do the following:

- Manage Snapshot copies
- Restore data from Snapshot copies
- Back up and replicate data
- Use encryption to prevent unauthorized access to data

## Additional data protection learning

Learn about advanced topics like configuration of intercluster replication, fan-in and fan-out strategies, and NetApp data-protection interfaces.

- **ONTAP Data Protection Fundamentals**  
(online course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours000000000024323](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours000000000024323)
- **ONTAP Data Protection Administration**  
(2-day instructor-led course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours000000000022724](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours000000000022724)
- **ONTAP Security and Compliance Solutions Administration**  
(2-day instructor-led course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours000000000024832](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours000000000024832)
- **ONTAP MetroCluster Installation**  
(2-day instructor-led course)  
[https://netapp.sabacloud.com/Saba/Web\\_spf/NA1PRD0047/common/ledetail/cours000000000022663](https://netapp.sabacloud.com/Saba/Web_spf/NA1PRD0047/common/ledetail/cours000000000022663)

Courses are regularly revised and updated. See the NetApp LearningCenter for the latest versions of the courses and offerings.

## References Documentation

- ONTAP 9 Documentation Center:  
<http://docs.netapp.com/ontap-9/index.jsp>
- *[TR-4015 SnapMirror Configuration and Best Practices Guide](#)*
- *[TR-4678 Data Protection and Backup: NetApp FlexGroup Volumes](#)*
- *[TR-4569 Security Hardening Guide](#)*





## Complete an exercise

Module 8  
Data protection

### Enabling Anti-Ransomware Protection

Access your lab equipment.

- Open your Exercise Guide to Module 8.
- Complete Exercise 3.
- Share your results.


This exercise requires approximately  
**5 minutes.**

See the instructions in your Exercise Guide.



# Addendum

## Secure purge and secure boot

 62 © 2023 NetApp, Inc. All rights reserved.

## Secure purge

NVE for General Data Protection Regulation (GDPR) and US public sector

Manage data spillage and right to erasure

Cryptographically shred the contents of old data blocks belonging to deleted files.

Use when there is an immediate need to destroy data:

- When data with different classification levels accidentally ends up in the same volume
- To delete user data cryptographically to satisfy GDPR requirements



**NetApp** 63 © 2023 NetApp, Inc. All rights reserved.

To meet security and compliance requirements, secure purge enables storage administrators to nondisruptively scrub old data from NVE-enabled volumes. Scrubbing data on an encrypted volume prevents the data from being recovered from the physical media. Sample uses include cases of “spillage,” where data traces may have been left behind when blocks were overwritten, or to securely delete the data of a vacating tenant.

Secure purge works only for previously deleted files on NVE-enabled volumes. You cannot scrub an unencrypted volume. You must use KMIP servers to serve keys, not the onboard key manager.

Secure purge functions differently depending upon your ONTAP version. In ONTAP 9.7 and earlier:


- The secure purge feature triggers a volume move that re-encrypts the active filesystem data with a new key. The moved volume remains on the current aggregate. The old key is automatically destroyed so that purged data cannot be recovered from the storage media.
- Data blocks are not be purged if they are in use by a Snapshot copy.
- If the volume being purged is the source of a SnapMirror relationship, you must break the SnapMirror relationship before you can purge the volume.  
If there are busy Snapshot copies in the volume, you must release the Snapshot copies before you can purge the volume. For example, you might need to split a FlexClone volume from its parent.

In ONTAP 9.8 and later:

- The re-encryption method differs for volumes that do and do not use SnapMirror data protection.
  - Volumes that use SnapMirror data protection (DP) mode re-encrypt data by using the volume move re-encryption method.
  - Volumes that do not use SnapMirror data protection or volumes that use SnapMirror extended data protection (XDP) mode use the in-place re-encryption method.
- You can change these defaults by using the `volume encryption secure-purge -re-encryption-method` command. The `volume-move` method is faster compared to the `in-place-rekey` method. The volume-move method requires additional space in the same aggregate and involves a cutover during which client I/O operations are temporarily blocked.
- By default, all Snapshot copies in FlexVol volumes are automatically deleted during the secure purge operation. Snapshot copies in FlexGroup volumes and volumes that use SnapMirror data protection are not automatically deleted.

## Protect systems in transit

Protected controller reboot



**Passphrase required  
after reboot**

**NetApp** 64 © 2023 NetApp, Inc. All rights reserved.



**Secure transport**



**Equipment return**



**Mission-forward  
deployments**

For customers like the military, which has clusters on mobile platforms (trucks, ships, aircraft), the data on the system must not be accessible if unauthorized users gain access to the storage system. NSE disk and volume encryption serve as a first line of defense. However, that defense can be overcome by physically hacking the storage. Protected controller reboot renders the hardware inoperable until the correct passphrase is supplied.

## Secure boot

### Unified Extensible Firmware Interface


- Verifies that software is genuine ONTAP software during boot
- Prevents hacked or prerelease versions of ONTAP software anytime that the system boots
- Verifies signed ONTAP images by the boot loader



Secure boot is another security feature that is designed to protect all new AFF and FAS systems from use or exploitation through hacked or prerelease versions of ONTAP software. The feature protects customers from purchasing gray-market or stolen hardware.

# Module 9

## Storage efficiency

 1 © 2023 NetApp, Inc. All rights reserved.

## About this module


This module focuses on enabling you to do the following:

- Implement storage-efficiency features
- Use FlexClone software volumes



# Lesson 1

## Thin provisioning

 3 © 2023 NetApp, Inc. All rights reserved.



## Thick and thin provisioning of volumes

### Thick provisioning of volumes (space-guarantee = volume)

- Requires reserved space within the aggregate for volume creation
- Helps to prevent overcommitment of an aggregate
- Simplifies storage management

### Thin provisioning of volumes (space-guarantee = none)

- Does not require reserved space within the aggregate for volume creation
- Enables more aggressive allocation
- Does not prevent overcommitment of an aggregate
- Requires more complex storage management

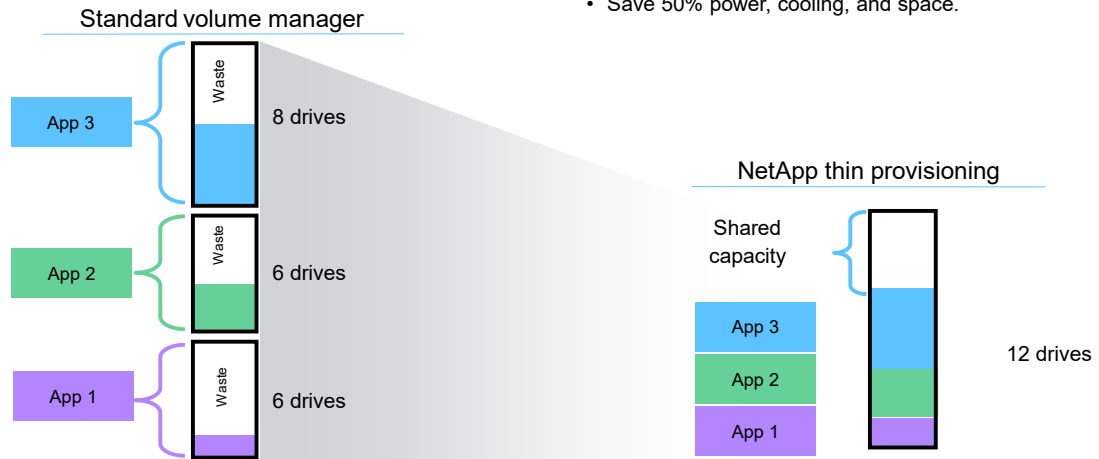
Administrators can manage storage systems by allocating volumes in one of two ways:

- Thick provisioning of volumes uses a space guarantee for a volume. A volume guarantee requires reserved space in the aggregate when the volume is created. Thick provisioning is a conservative approach that prevents administrators from overcommitting space to an aggregate. Thick provisioning simplifies storage management at the risk of wasting unused space.
- Thin provisioning of volumes uses a space guarantee of none, meaning that no space within the aggregate is reserved for the volume when the volume is created.

## Thin provisioning

- Typical: Only 40% of provisioned storage is used.

- NetApp: More than 70% of provisioned storage is used.
  - Buy 50% less storage.
  - Save 50% power, cooling, and space.



NetApp 5 © 2023 NetApp, Inc. All rights reserved.

When you compare the NetApp storage-use approach to competing approaches, one feature stands out. Flexible dynamic provisioning with FlexVol technology provides high storage-use rates and enables customers to increase capacity without physically repositioning or repurposing storage devices. NetApp thin provisioning enables users to overcommit data volumes, resulting in high-use models. You can think of the approach as “just-in-time” storage.


To manage thin provisioning on a cluster, use the volume command.

## Enable thin provisioning


The screenshot displays the NetApp ONTAP System Manager interface. On the left, a navigation sidebar includes sections for DASHBOARD, INSIGHTS, STORAGE (with sub-items like Overview, Volumes, LUNs, etc.), and NETWORK. The main content area is titled 'Volumes' and shows a list of volumes. The volume 'c1\_svm1\_vol1' is selected, and its details are shown in a right-hand pane titled 'Edit Volume'. In the 'Storage and Optimization' section, the 'Enable thin provisioning' checkbox is checked. Below the interface, a terminal window shows the command: `::> volume modify -vserver svm4 -volume svm4_vol_002 -space-guarantee none`. A green callout box with a white arrow points to the 'Enable thin provisioning' checkbox, containing the text 'Volume space-reservation setting'.

NetApp 6 © 2023 NetApp, Inc. All rights reserved.

Enabling thin provisioning is as simple as selecting a checkbox in NetApp ONTAP System Manager or using the `volume modify` command to set the space guarantee to `none`.



## Lesson 2 Deduplication and compression

 7 © 2023 NetApp, Inc. All rights reserved.

## Volume efficiency



### Deduplication

- Elimination of duplicate data blocks
- Inline or postprocess options
- Inline deduplication for AFF systems and Flash Pool systems, reducing the number of writes to SSDs



### Data compression

- Compression of data within a file
- Inline or postprocess options
- Two compression methods:
  - **Secondary:** 32KB compression groups
  - **Adaptive:** 8KB compression groups, which improve read performance

NetApp ONTAP software provides two features that can increase volume efficiency: deduplication and data compression. You can run deduplication and data compression together or independently on a FlexVol volume to reduce the amount of physical storage that the volume requires.

To reduce the amount of required physical storage, deduplication eliminates duplicate data blocks and data compression shrinks the data within a file. Depending on the version of ONTAP software and the type of drives that are used for the aggregate, the volume efficiency features can be run inline or postprocess.

Inline deduplication can reduce writes to SSDs. Inline deduplication is enabled by default on all new volumes that are created on AFF systems. Inline deduplication can also be enabled on new and existing Flash Pool volumes.

Data compression combines multiple 4KB NetApp WAFL blocks into compression groups before compression. Two data compression methods are available: secondary and adaptive.

## Enabling deduplication

The screenshot displays the ONTAP System Manager interface. On the left is a navigation sidebar with categories like DASHBOARD, INSIGHTS, STORAGE, and NETWORK. The main area shows a list of volumes under 'Volumes', with 'c1\_svm1\_vol1' selected. An 'Edit Volume' modal is open for this volume, showing the 'Storage Efficiency' section. In this section, 'Enable background deduplication' is checked, and 'Enable background compression' is also checked. A dropdown menu for 'STORAGE EFFICIENCY POLICY' is set to 'Default'. A green callout box with a speech bubble points to the 'Storage Efficiency' section, containing the text 'Deduplication and compression setting'. Below the interface, a terminal window shows the command: `::> volume efficiency on -vserver svm4 -volume svm4_vol_002`. The footer of the page includes the NetApp logo and copyright information: '© 2023 NetApp, Inc. All rights reserved.'

Deduplication improves the efficiency of physical storage space by eliminating redundant data blocks within a FlexVol volume. Deduplication works at the block level on an active file system and uses the WAFL block-sharing mechanism. Each block of data has a digital signature that is compared with all the other blocks in the data volume. If an exact match is identified, the duplicate block is discarded. A data pointer is modified so that the storage system references the copy of the data object that is stored on disk. The deduplication feature works well with datasets that have large quantities of duplicated data or white space. You can configure deduplication operations to run automatically or according to a schedule. You can run deduplication on new or existing data on any FlexVol volume.

The default storage efficiency policy uses a schedule to run the storage efficiency processes once daily. Use the `volume efficiency policy create` command to define a new storage efficiency policy.

## Configuring data compression

The screenshot shows the ONTAP System Manager interface. On the left is a navigation menu with sections for DASHBOARD, INSIGHTS, STORAGE, and NETWORK. The STORAGE section is expanded to show Volumes. A table lists several volumes, with 'c1\_svm1\_vol1' selected. To the right, the 'Edit Volume' panel is open for 'c1\_svm1\_vol1'. It shows the volume name, status (Online), style (FlexVol), and storage efficiency (Disabled). The 'Storage Efficiency' section is expanded, showing options to 'Enable background deduplication', 'Enable background compression', and 'Enable inline compression'. The 'Enable background compression' and 'Enable inline compression' options are checked. The 'Storage Efficiency Policy' is set to 'Default'.

```
::> volume efficiency modify -vserver svm4 -volume svm4_vol1002  
-compression true -inline-compression true
```

NetApp 10 © 2023 NetApp, Inc. All rights reserved.

Data compression enables you to reduce the physical capacity that is required to store data on a cluster by compressing data blocks within a FlexVol volume. Data compression optimizes the storage space and bandwidth that are required to replicate data during volume operations, such as moving volumes and performing SnapMirror transfers. You can compress standard data files, virtual disks, and LUNs. You cannot compress file system internal files, alternate data streams, or metadata.

To manage compression on a cluster, use the `volume efficiency` command.

## Characteristics of data compression

### Inline compression

- Data is compressed in memory before being written to the drives.
- Storage consumption and write operations are reduced.
- Throughput increases because of fewer I/O operations.

### Postprocess compression

- Uncompressed data is compressed during idle time.
- Only previously uncompressed blocks are compressed.
- Compression occurs before deduplication.
- NetApp ONTAP software can detect incompressible data before wasting effort.

For more information, see the *ONTAP 9 Logical Storage Management Guide*.

There are two types of data compression: inline and postprocess.

During inline data compression, the data is compressed in memory before being written to the disk, reducing the amount of write I/O. This option can affect your write performance and so should not be used for performance-sensitive environments on HDD configurations without proper testing. The AFF systems and the Flash Pool systems are exceptions. Inline compression can be used for primary workloads on AFF systems and Flash Pool systems.

ONTAP software employs two compression methods: adaptive and secondary. Adaptive compression combines fewer blocks into a compression group. The smaller compression group takes less time to decompress, which can improve read performance, compared to secondary compression. Adaptive compression provides higher performance but fewer space savings and is better suited for random workloads. Secondary compression combines more data blocks into a 32KB compression group. Secondary compression provides more space savings and is better suited for sequential workloads.

Beginning with version 9.8, ONTAP software uses temperature-sensitive storage efficiency. Cold data blocks are compressed by using secondary compression, and hot data blocks are compressed by using adaptive compression.

ONTAP software can run postprocess compression and inline compression. When enabled, postprocess compression and deduplication are run automatically when sufficient data has changed, or they can be scheduled to run at specific times. Like inline compression, postprocess compression compresses data only if the compression savings are >50% (adaptive compression) or >25% (secondary compression). Postprocess compression is not supported on AFF configurations.


For more information, see the *ONTAP Logical Storage Management Guide*.





## Lesson 3

### Flash efficiency

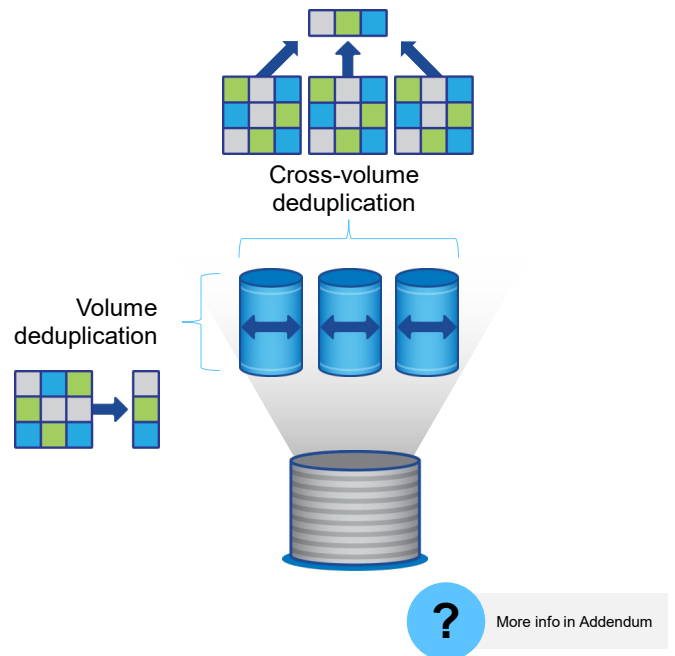
 12 © 2023 NetApp, Inc. All rights reserved.

## Aggregate inline deduplication

### Overview

Aggregate inline deduplication enables block sharing across multiple volumes within an aggregate:

- Is available on only AFF and All SAN Array (ASA) systems
- Uses the volume efficiency parameter:  
`-cross-volume-inline-dedupe`



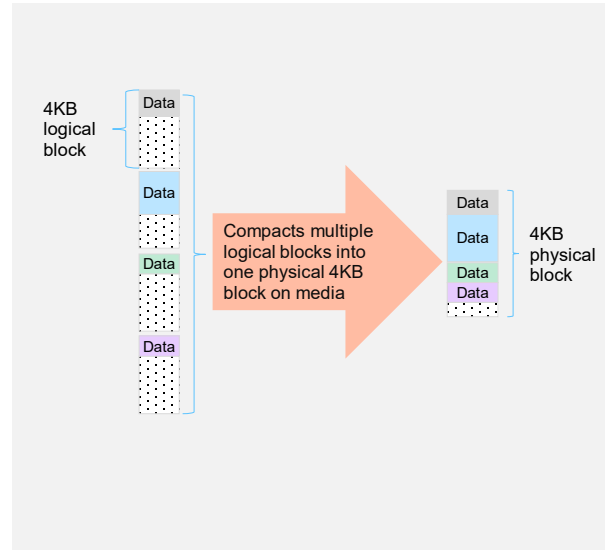
**NetApp** 13 © 2023 NetApp, Inc. All rights reserved.

Aggregate inline deduplication is available only on AFF and All SAN Array (or ASA) systems and is enabled by default. You can enable and disable the feature by using the `volume efficiency` command parameter `-cross-volume-inline-dedupe`. For accounting purposes, storage space that cross-volume blocks consume is assigned to the FlexVol volume that first wrote to the block.

For information about feature support, see the *Logical Storage Management Guide*.

## Inline data compaction

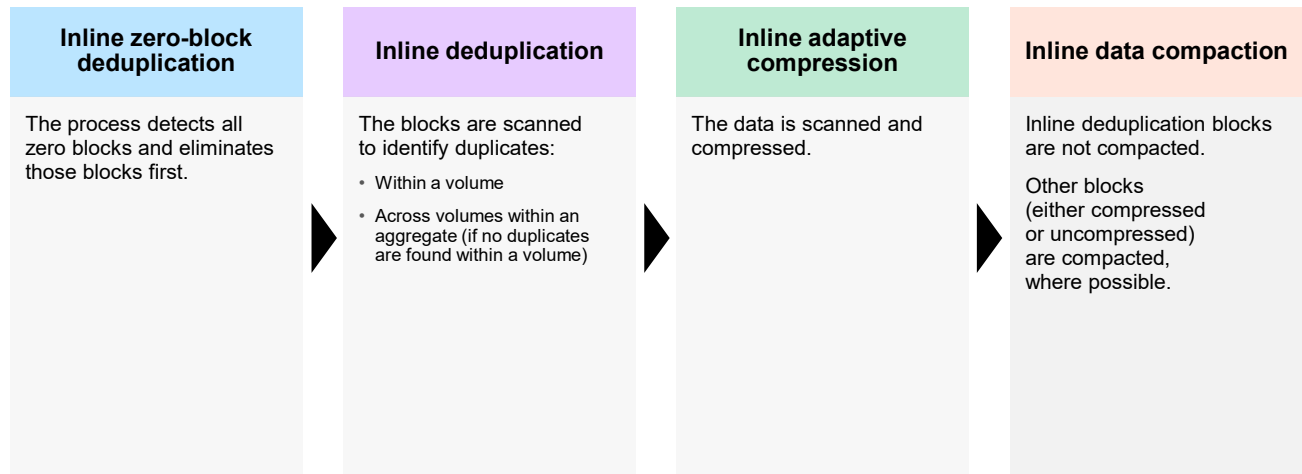
- Writes multiple logical data blocks in the same volume to one 4KB block on storage:
  - Compaction occurs during the consistency point (CP) operation just before the write to media.
  - Compaction occurs after inline adaptive compression and inline deduplication.
- Provides additional space savings with highly compressible data
- Is enabled by default for new AFF systems but is disabled on FAS systems:
  - Optional policy for Flash Pool aggregates
  - Optional policy for hard-disk-only aggregates




Data compaction is disabled by default for FAS systems. To control inline data compaction on FAS systems with Flash Pool (hybrid) aggregates or HDD aggregates at the volume or aggregate level, use the `volume efficiency modify` command. If you enable data compaction at the aggregate level, data compaction is enabled on any new volume that is created with a volume space guarantee of none in the aggregate. Enabling data compaction on a volume on an HDD aggregate uses additional CPU resources.

## AFF inline storage efficiency

ONTAP workflow

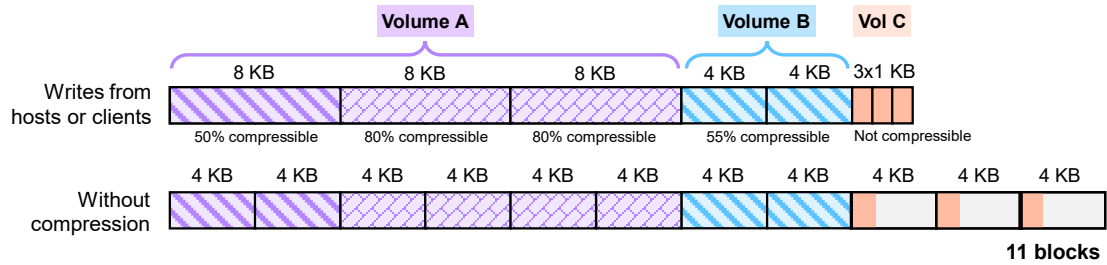


 15 © 2023 NetApp, Inc. All rights reserved.

Aggregate inline deduplication works seamlessly with other efficiency technologies, such as compression and inline zero-block deduplication.

# Storage consumption

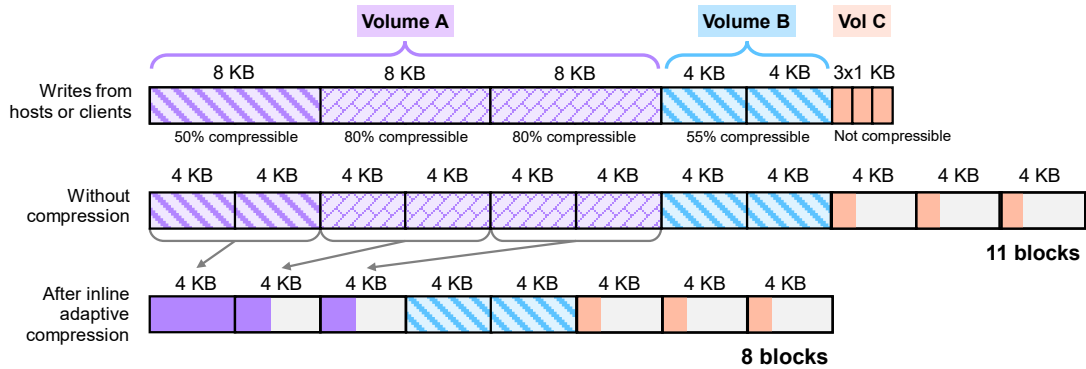
No inline storage efficiency



The figure shows the writes for a host or client and the amount of space on disk when no efficiency features are enabled.

# Storage consumption

## Inline adaptive compression

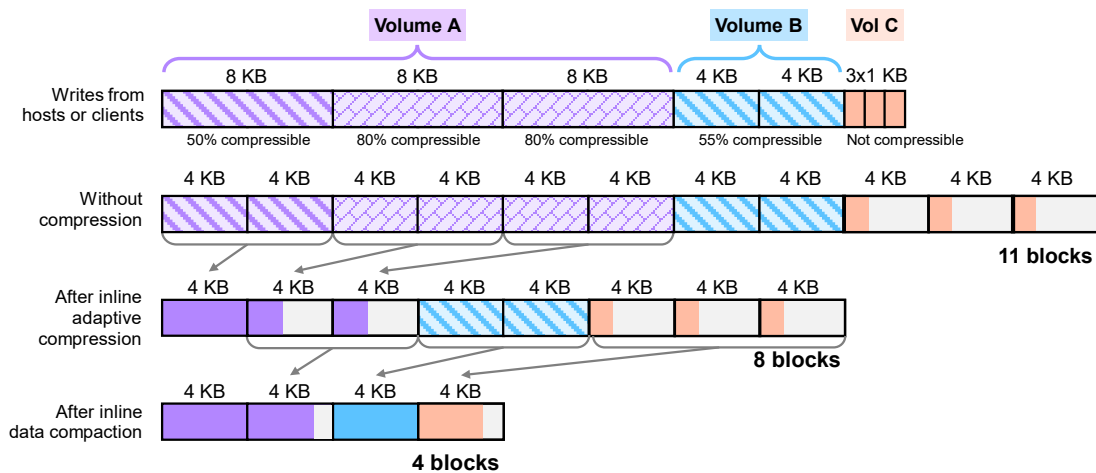


The figure shows the results of inline adaptive compression, which is automatically enabled on AFF systems.

Temperature-sensitive data compression was added to AFF systems in ONTAP 9.8 software. Frequently accessed hot data is compressed inline, using compression groups. Infrequently accessed cold data is then compressed again in the background, using a more aggressive 32KB compression group. These changes mean better performance for hot data and better data reduction ratios for all data.

## Storage consumption

Inline adaptive compression and inline data compaction



NetApp 18 © 2023 NetApp, Inc. All rights reserved.

The figure shows the default policy for AFF systems that run ONTAP 9 software.

Data compaction is an inline operation and occurs after inline compression and inline deduplication. On an AFF system, the order of execution is as follows:

1. Inline zero-block deduplication: All zero blocks are detected. No user data is written to physical storage. Only metadata and reference counts are updated.
2. Inline adaptive compression: 8KB logical blocks are compressed into 4KB physical blocks. Inline adaptive compression efficiently determines the compressibility of the data and does not waste many CPU cycles trying to compress incompressible data.
3. Inline deduplication: Incoming blocks are opportunistically deduplicated to existing blocks on physical storage.
4. Inline adaptive data compaction: Multiple logical blocks of less than 4KB are combined into a single 4KB physical block, which maximizes savings. Also, 4KB logical blocks that inline compression skips are compressed to improve compression savings.

## Default storage efficiency settings

Storage efficiency feature	AFF	FAS
Inline compression	Enabled	Disabled
Background compression	Not supported	Disabled
Inline volume deduplication	Enabled	Flash Pool only
Background volume deduplication	Enabled*	Disabled
Inline aggregate deduplication	Enabled	Not supported
Background aggregate deduplication	Enabled*	Not supported
Inline compaction	Enabled	Disabled

\* Background deduplication operations are disabled when the `-inline-only` storage efficiency policy is applied to a volume.

All inline storage efficiency features are enabled by default on all existing and newly created volumes on all AFF systems. Storage efficiency features include inline deduplication, inline cross-volume deduplication, and inline compression. They are enabled by default on AFF systems as shown in the table.

Aggregate-level deduplication is not supported on FAS systems. However, volume-level storage efficiency features can be enabled as needed.


The `volume efficiency start` command can be used to initiate jobs to compress, deduplicate, and compact data that has already been written to storage.





## Lesson 4


# Logical space reporting

 20 © 2023 NetApp, Inc. All rights reserved.

## The storage service provider conflict

Charge for data stored

- Storage service providers want to charge customers for reserved data space.
- Customers want to pay only for stored data.




Assume that, as a service provider, you provision a 20TB volume to a customer, and the customer stores 10TB of data in the volume. As the provider, you want to charge the customer for reserving 20TB of storage space. As the customer, you want to pay for only storing 10TB of data.

## The storage service provider conflict

Storage efficiencies provide more space

- Storage efficiencies give customers more storage space than they pay for.
- Before ONTAP 9.4 software, ONTAP software provided reporting only for consumed physical storage. Storage service providers could charge only for consumed space, not for stored data.



**NetApp** 22 © 2023 NetApp, Inc. All rights reserved.

With ONTAP storage-efficiency technologies, if 10 TB of data is reduced to 4 TB, the actual used space in the volume is shown as 4 TB.


The customer sees that the available space is 16 TB. This situation does not help you to charge the customer based on the actual amount of data stored, regardless of storage efficiencies.

## Logical space reporting

Volume option

`-is-space-reporting-logical` [true | false]  
shows customers consumed logical space rather than consumed physical space.

```
volume modify -vserver SVM-name -volume volume-name  
-is-space-reporting-logical true
```

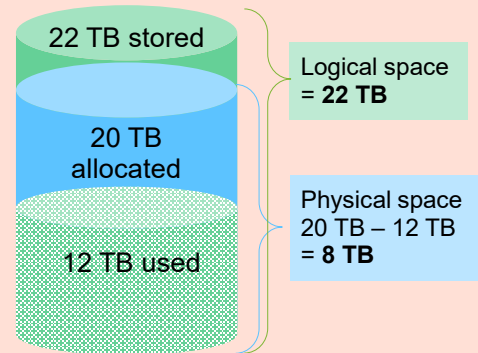


Enable logical space reporting to show the amount of consumed logical storage space rather than the used physical storage space after storage-efficiency features are run.

## Logical space enforcement

- Storage efficiencies enable customers to store more data than space was allocated for.
- Volume option  
`-is-space-enforcement-logical` [true | false] ensures that customers cannot store more than the allocated logical space, regardless of the consumed physical space.
- Error messages are generated when stored data reaches 95%, 98%, and 100% of logical space limits.

```
volume modify -vserver SVM-name -volume volume-name  
-is-space-enforcement-logical true
```



With ONTAP storage efficiencies, the customer can store more data than the logical available space. The logical space reporting shows that the amount of stored customer data is larger than the size of the provisioned volume.

To overcome this issue, ONTAP introduced the enforcement of logical space. With the logical space enforcement feature, customers cannot store data in a volume if the logical space limit is reached. Thus, the customer cannot store more than 20 TB of data even though physical space is available. ONTAP systems trigger error messages as the customer reaches the logical space limit at 95%, 98%, and 100%. These space limits are predefined and nonconfigurable. Use an external monitoring application to set alerts for custom space limits.

Any new writes to the volume when the logical space used is 100% return an ENOSPC (out of space) error message.




## Complete an exercise

Module 9  
Storage efficiencies

### Managing Storage Efficiency

- Access your lab equipment.
- Open your Exercise Guide to Module 9.
- Complete Exercise 1.
- Share your results.

This exercise requires approximately  
**15 minutes.**


 25 © 2023 NetApp, Inc. All rights reserved.

See the instructions in your Exercise Guide.




## Lesson 5

### Volume and file clones

 26 © 2023 NetApp, Inc. All rights reserved.

## FlexClone technology

- FlexClone software uses block pointers to enable you to create multiple, instant dataset clones (files, LUNs, or entire volumes) with no storage overhead.
- FlexClone technology provides dramatic improvement for application testing and development environments:
  - Create an instantaneous replica of a file or LUN (such as an entire database).
  - Provision thousands of virtual machines (VMs) in seconds by cloning *golden images*.
- You can split clones from the source, but then make copies of all shared blocks.



**NetApp** 27 © 2023 NetApp, Inc. All rights reserved.

FlexClone volumes are often referred to as *writable Snapshot copies*. By using block pointers, you can create multiple, instant dataset clones—files, LUNs, or entire volumes—with no initial storage overhead. Only when data is added or changed in a clone is storage space consumed.


Clones are especially useful in test and development environments. Data can be replicated numerous times within seconds and used just like the source data, without concerns of damaging or destroying the source data. FlexClone software is also useful in virtual environments, where golden images of VMs can be cloned thousands of times.

Clones can be split from the source, but then make copies of all shared blocks. This behavior is useful for upgrading or patching an application in a clone and then rolling it out by splitting off the clone and promoting it to production. Rollbacks can be as simple as promoting the source back into production.



## How cloning works

- First, a Snapshot copy of the volume is made.
- A second Snapshot copy of the volume is made for the clone.
  - Initially, the parent and clone snapshot copies share the same data blocks (blocks A–F).
  - Modifications to the original volume (blocks G and H) are separate from modifications to the cloned volume.
  - Modifications to the clone (blocks I–K) are separate from the original volume. These blocks are the only drive space that the clone consumes.



NetApp 28 © 2023 NetApp, Inc. All rights reserved.

FlexClone volumes are managed like regular FlexVol volumes, with a few key differences. FlexClone volumes have the following features:

- FlexClone volumes are point-in-time, writable copies of parent volumes. The FlexClone volume does not reflect changes that are made to the parent volume after the FlexClone volume is created.
- FlexClone volumes are fully functional volumes that are managed, as with the parent volume, by using the `vol` command. As with parent volumes, FlexClone volumes can be cloned.
- FlexClone volumes are always in the same aggregate as parent volumes.
- FlexClone volumes and parent volumes share the drive space for shared data. The process of creating a FlexClone volume is instantaneous and requires no additional drive space (until changes are made to the clone or parent).
- A FlexClone volume is created with the same space guarantee as the parent.
- You can sever the connection between the parent and the clone. This action is called *splitting* the FlexClone volume. Splitting removes all restrictions on the parent volume and causes the FlexClone volume to use its own storage.

**Note:** When you split a FlexClone volume from its parent volume, the following occurs:

- All existing Snapshot copies of the FlexClone volume are deleted.
- Creation of Snapshot copies is disabled while the splitting operation is in progress.
- Quotas that are applied to a parent volume are not automatically applied to the clone.
- When a FlexClone volume is created, existing LUNs in the parent volume are also present in the FlexClone volume, but these LUNs are unmapped and offline.

**Note:** ON AFF systems, when a FlexClone volume is split from the parent volume, the data blocks are not duplicated. Therefore, the clone splitting operation consumes no additional space and is completed quickly. Copying the data blocks is skipped because aggregate (cross-volume) deduplication eliminates the copied data blocks.

## Clone a volume

The screenshot shows the ONTAP System Manager interface. The 'Volumes' page is active, displaying a table of volumes. The 'c1\_svm1\_vol1' volume is selected, and its context menu is open, with the 'Clone' option highlighted. A 'Clone Volume' dialog box is open, showing the following details:

- NAME: c1\_svm1\_vol1\_clone
- Enable thin provisioning
- CLONE PARENT SNAPSHOT COPY:
  - Add a Snapshot copy
  - Use an existing Snapshot copy

Buttons for 'Cancel' and 'Clone' are visible at the bottom of the dialog.

```
::> volume clone create -vserver svm3 -parent-volume exp_svm3_NFS_volume  
-flexclone exp_svm3_NFS_volume_clone
```

NetApp 29 © 2023 NetApp, Inc. All rights reserved.

To create a FlexClone volume, use the `volume clone create` command. If you do not include the `-parent-snapshot` option, a new Snapshot copy of the parent volume is created and used as the parent of the FlexClone volume. This Snapshot copy cannot be deleted while the FlexClone volume exists.

Beginning with ONTAP 9.10.1 you can clone a FlexVol volume which is a constituent of a FlexGroup volume.

## Split a cloned volume

The screenshot shows the ONTAP System Manager interface. The 'Volumes' page is active, displaying a table of volumes. A context menu is open over the 'c1\_svm1\_vol1\_clone' volume, with 'Split Clone' highlighted. A 'Split Clone' dialog box is displayed, showing the selected clone volume 'c1\_svm1\_vol1\_clone' and the parent volume 'c1\_svm1\_vol1'. Below the dialog, a terminal window shows the command: `::> volume clone split start -vserver svm3 -flexclone svm3_vol_002_clone`

Name	Storage VM	Status	Capacity
c1_svm1_root	c1_svm1	Online	1.4 MIB used
c1_svm1_vol1	c1_svm1	Online	86.9 MIB used
c1_svm1_vol1_clone	c1_svm1	Online	97.1 MIB used
	c1_svm2	Online	1.21 MIB used
	c1_svm2	Online	1.04 MIB used
	c1_svm3	Online	1.39 MIB used
	c1_svm3	Online	1.04 MIB used

Terminal output: `::> volume clone split start -vserver svm3 -flexclone svm3_vol_002_clone`

To split a FlexClone volume from its parent volume, use the `volume clone split start` command. ONTAP software will replicate all the data blocks that are shared by the parent volume Snapshot copy and the FlexClone.

## Clone a file or LUN

The screenshot shows the ONTAP System Manager interface. On the left is a navigation sidebar with categories: DASHBOARD, INSIGHTS, and STORAGE. The STORAGE section is expanded, showing options like Overview, Volumes, LUNs (selected), Consistency Groups, Shares, Buckets, Qtrees, Quotas, Storage VMs, and Tiers. The main area displays a table of LUNs with columns for Name, Con..., Storage VM, and Volume. A context menu is open over the first LUN, with the 'Clone' option highlighted. To the right, a 'Clone LUN' dialog box is open, showing the name 'SQL\_prodB\_Titan\_PartsDB\_1\_clone' and host information including 'Windows' for the operating system and 'Existing initiator group' for host mapping. Below the dialog, a terminal window shows the command: 

```
::> volume file clone create -server svm5 -volume svm5_vol_002 -source-path file1 -destination-path file1_clone
```


NetApp 31 © 2023 NetApp, Inc. All rights reserved.

You can also use FlexClone technology to clone individual files or LUNs. This ability is useful in application testing. Unlike FlexClone volumes, cloning of files and LUNs does not require a backup Snapshot copy.

Use the `volume file clone` command to clone individual files.

# Knowledge check

Module 9: Storage efficiency

 32 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

# Which types of data compression are available in NetApp ONTAP software?

- a. inline and external
- b. inline and preprocess
- c. inline and postprocess
- d. inline and reclaimable

## Knowledge check

**True or false: Data can be written to a FlexClone volume.**

- a. true
- b. false

## Knowledge check

**True or false: A FlexClone volume, by definition, shares no data blocks with the parent volume.**

- a. true
- b. false



## References

- NetApp Hardware Universe:  
<http://hwu.netapp.com>
- ONTAP 9 Documentation Center:  
<http://docs.netapp.com/ontap-9/index.jsp>
- *[ONTAP Storage Efficiency  
Technical Presentation](#)*

## Module summary

This module focused on enabling you to do the following:

- Implement storage-efficiency features
- Use FlexClone software volumes



## Complete exercises

Module 9  
Storage efficiencies

### Managing FlexClone volumes

- Access your lab equipment.
- Open your Exercise Guide to Module 9.
- Complete Exercise 2.
- Share your results.

This exercise requires approximately  
**10 minutes.**

See the instructions in your Exercise Guide.



## Share your experiences


Roundtable discussion

- Were you able to observe storage-efficiency benefits in your exercise environment?
- What are some popular uses for FlexClone volumes?



# Addendum

## Inline deduplication status

 40 © 2023 NetApp, Inc. All rights reserved.

## Aggregate inline deduplication


### Status

#### Volume status

```
::> volume efficiency show -vserver svm4 -volume svm4_vol1003
      -fields cross-volume-inline-dedupe
vserver  volume          cross-volume-inline-dedupe
-----
svm4     svm4_vol1003         true
```

#### Aggregate status

```
::> aggregate efficiency show -aggregate n1_data_003
                                     Aggregate: n1_data_003
                                     Node: cluster1-01
Cross Volume Background Deduplication: false
Cross Volume Inline Deduplication: true
Has Cross Volume Deduplication Savings: true
Has Auto Adaptive Compression Savings: false
```

 41 © 2023 NetApp, Inc. All rights reserved.

To display the aggregate inline deduplication status for a volume, use the `volume efficiency show` command. To display the status for an aggregate, use the `aggregate efficiency show` command.

To enable or disable aggregate inline deduplication for a volume, use the `volume efficiency modify -cross-volume-inline-dedupe {true|false}` command.

**Note:** If you try to enable aggregate inline deduplication on a node that is not an AFF node, the following error message appears:

```
::> volume efficiency modify -vserver svm3 -volume smb1_share -cross-volume-inline-dedupe
true
```

```
Error: command failed: Failed to modify efficiency configuration for volume "smb1_share"
of Vserver "svm3":
```

Cross volume deduplication is supported only on volumes that are owned by nodes that are All-Flash optimized personality enabled.

# Aggregate inline deduplication

## Savings

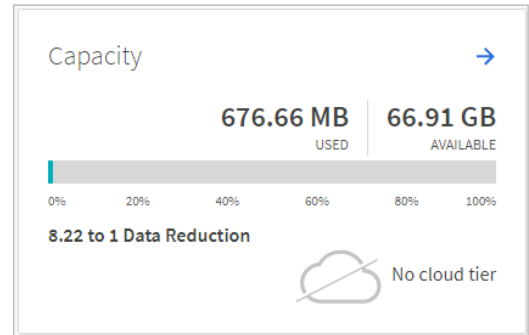
### Aggregate Savings

```
::> aggr show-efficiency -details
Aggregate: cluster1_ssd_001
Node: cluster1-01

Total Storage Efficiency Ratio: 25.22:1
Total Data Reduction Ratio: 2.57:1

Aggregate level Storage Efficiency
(Aggr Dedupe and Data Compaction): 1.33:1
Volume Dedupe Efficiency: 1.40:1
Compression Efficiency: 1.29:1

Snapshot Volume Storage Efficiency: 27.14:1
FlexClone Volume Storage Efficiency: -
```




The data-reduction ratio includes aggregate inline deduplication savings.

Aggregate inline deduplication savings and data compaction savings are combined and reported as one ratio percentage in the ONTAP System Manager dashboard. Use the `aggr show-efficiency -details` command to view the space savings for each storage efficiency feature.

# Module 10

## Cluster maintenance


 1 © 2023 NetApp, Inc. All rights reserved.



## About this module


This module focuses on enabling you to do the following:

- Use the NetApp Active IQ customer dashboard
- Plan NetApp ONTAP software upgrades
- Follow recommended practices for peak performance
- Configure event notifications and alerts
- Prepare to engage NetApp technical support
- Perform cluster maintenance



# Lesson 1

## Data collection, monitoring, and automation tools

 3 © 2023 NetApp, Inc. All rights reserved.

# Alerts

Tools for monitoring the system:

- NetApp ONTAP System Manager
- Event management system
- AutoSupport
- NetApp Active IQ Unified Manager

The screenshot displays the NetApp ONTAP System Manager interface. The left sidebar contains a navigation menu with categories: DASHBOARD, INSIGHTS, STORAGE, NETWORK, EVENTS & JOBS (expanded), and CLUSTER. Under 'EVENTS & JOBS', 'System Alerts' is selected. The main content area is titled 'System Alerts' and shows a table with columns: Alert Name, Time, Node, Monitor, and Resource. A single alert is listed: 'DisabledInuseSASPort\_Alert 20210306\_193452' on 'node3' for 'SAS-connect'. Below the table, a detailed view of the alert is shown, including the Node (cluster1-01), Monitor (node-connect), Subsystem (SAS-connect), and Alert ID (DisabledInuseSASPort\_Alert). The alert severity is Major, with a probable cause of 'Cable\_tamper' and a description: 'SAS node3:1b port is disabled. This might occur if the port has been administratively disabled or the attached cable is faulty.' The possible effect is that the controller might lose a path to storage devices. Corrective actions include verifying the cable connection and replacing the cable if necessary.

Monitoring your system regularly is a best practice.

In the example, a notification from NetApp ONTAP System Manager (formerly OnCommand System Manager) needs to be diagnosed. When there is an alert or event, first try the solution that the monitoring software suggests.

## Event management system

- The event management system (EMS) does the following:
  - Writes events to the event log
  - Sends and routes notifications of events
  - Collects events throughout the cluster
  - Can view events of all nodes from any node
    - `::> event log show`
- Each event contains the following:
  - Message name
  - Severity level
  - Description
  - Corrective action, if applicable

Time	Node	Severity	S...	Event
Thursday, Feb 1...	cluster2-02	error	w...	waf1.cp.toolong: Aggregate aggr0_n2 ex...
Thursday, Feb 1...	cluster2-02	error	w...	waf1.cp.toolong: Aggregate aggr0_n2 ex...
Thursday, Feb 1...	cluster2-02	error	w...	waf1.cp.toolong: Aggregate aggr0_n2 ex...
Thursday, Feb 1...	cluster2-02	error	w...	waf1.cp.toolong: Aggregate aggr0_n2 ex...

NetApp 5 © 2023 NetApp, Inc. All rights reserved.

The event management system collects and displays information about events that occur in a cluster.

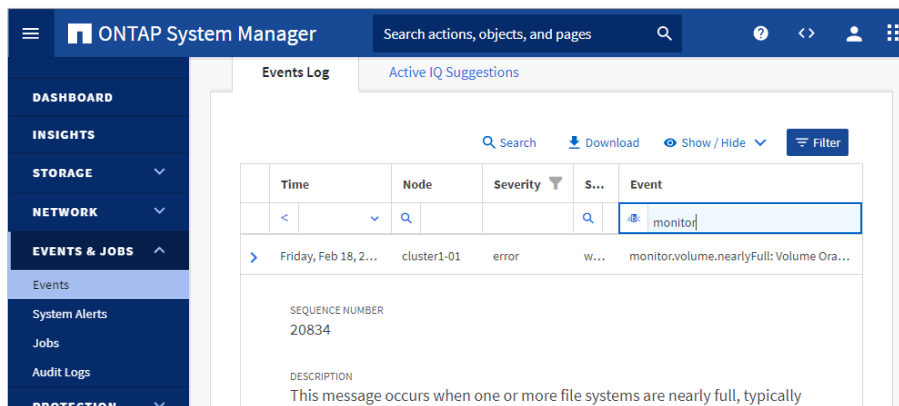
Beginning with ONTAP 9.10.1 software, you can now use ONTAP System Manager to configure email addresses, syslog servers, REST API clients (WebHooks), and SNMP traphosts as event notification destinations.

You can also configure event notification and logging using the CLI. The `event notification create` command creates a new notification of a set of events that is defined by an event filter to one or more notification destinations. Use the `event filter create` command to select the events of interest. An event filter is made up of one or more rules, each of which contains the event (message) name, the event severity, and the event SNMP trap type. The `event notification destination create` command creates a new email or syslog type event notification destination. Configure the SNMP notification destination by using the `system snmp traphost` command.

## Event log filtering

Filter event management system log messages by severity, time, message name, and other criteria.

```
::> event log show -severity {EMERGENCY|ALERT|ERROR|NOTICE|INFORMATIONAL|DEBUG}
::> event log show -time "08/30/2022 10:00:00".."08/30/2022 11:30:00"
::> event log show -severity informational -message-name kern.uptime.filer
```



NetApp 6 © 2023 NetApp, Inc. All rights reserved.

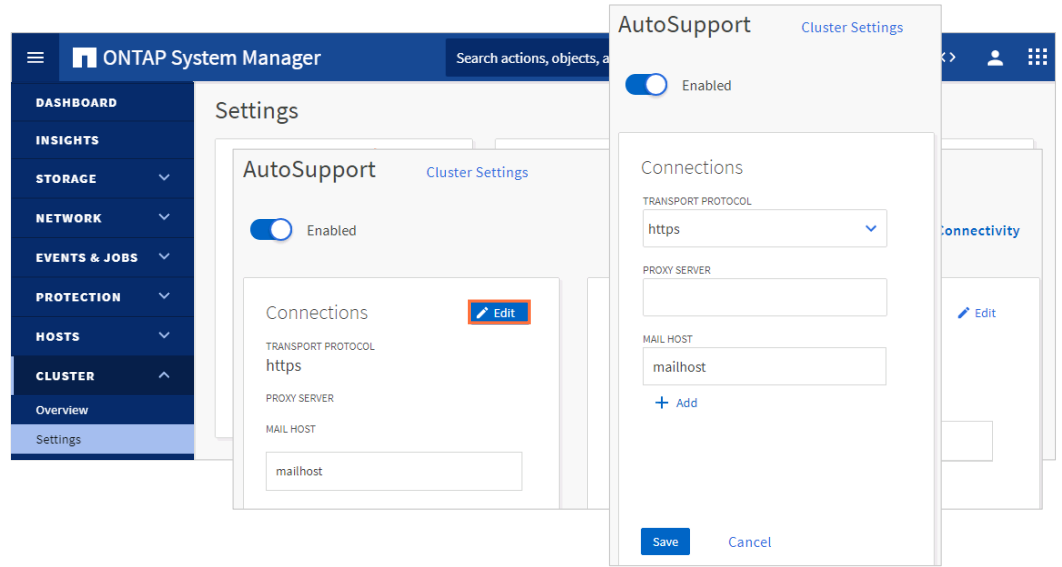
Use event log filtering to easily find relevant error messages that are buried in thousands of normal messages.

```
::> event log show ?
```

```
[ -detail | -detailtime | -instance | -fields <fieldname>, ... ]
[ [-node] <nodename> ] Node name
[ [-seqnum] <Sequence Number> ]. Log Sequence #
[ -time <"MM/DD/YYYY HH:MM:SS"> ] Time of event
[ -severity <level> ] Event Severity
[ -source <text> ]. Source
[ -message-name <Message Name> ] Message Name
[ -event <text> ] Event
[ -action <text> ] Corrective Action
[ -description <text> ] Description
[ -filter-name <text> ] Filter Name
```

# AutoSupport

- Is an integrated monitoring and reporting technology
- Checks the health of NetApp systems
- Should be enabled on all ONTAP clusters



NetApp 7 © 2023 NetApp, Inc. All rights reserved.

AutoSupport is an integrated and efficient monitoring and reporting technology. When enabled on a NetApp system, AutoSupport checks system health on a continual basis. AutoSupport should be enabled on all NetApp ONTAP clusters.

AutoSupport can be enabled or disabled. To configure AutoSupport, select **More options** from the AutoSupport action menu on the ONTAP System Manager Cluster > Settings page. Then, click **Edit** and enter your configuration information.

If you plan any changes to your controllers, you should manually trigger an AutoSupport message before you make the changes. The message provides a “before” Snapshot copy for comparison in case a problem arises later.

## Active IQ

The evolution of AutoSupport

- Actionable intelligence
- Predictive, self-healing care
- Global analytics

### Active IQ: AI-Powered Digital Advisor

NetApp® Active IQ® uses AI/ops to simplify the proactive care and optimization of your NetApp environment, leading to reduced risks and higher availability.

As business continuity plans are rolled out, many organizations are seeing their production systems utilized at a scale or growth trajectory that's beyond normal expectations. Active IQ has deployed new risk signatures to help NetApp customers stay ahead of potential performance and capacity issues. You can learn more about these new risks in this [customer support bulletin](#). Then [login into Active IQ](#) to check system health.

[Login to Active IQ to check system health](#)

#### What can you do with Active IQ?

Identify and remediate system health risks that can cause downtime

Uncover systems reaching performance or capacity limits

Identify and remediate security risks

Plan system software upgrades

Confirm AutoSupport adoption

#### Learn More

[Active IQ product page](#)

[Documentation Resources](#)

[Online Support Page](#)

[Turn on AutoSupport in ONTAP](#)

[Learn about the new Digital Advisor](#)

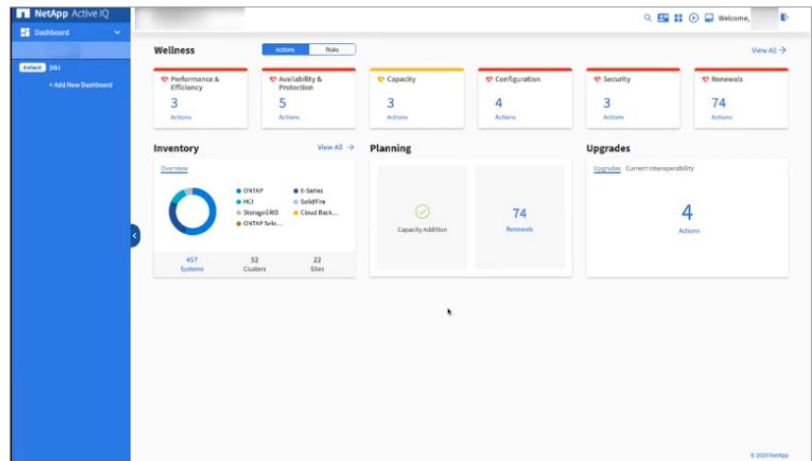

[Active IQ and AutoSupport users community](#)

NetApp Active IQ is a suite of web-based applications that are hosted on the NetApp Support Site and are accessible through your web browser. Active IQ uses data from the AutoSupport tool. Active IQ proactively identifies storage infrastructure issues through a continuous health-check feature. Active IQ also automatically provides guidance about remedial actions that help to increase uptime and avoid disruptions to your business.

For example, Active IQ might find a configuration issue, a bad disk, or version incompatibility on your system. Active IQ might also notify you about end-of-life (EOL) issues or an upcoming support contract expiration date.

# Active IQ Digital Advisor

- Dashboard
  - Wellness
  - Inventory
  - Planning
  - Upgrades
- Mobile app for iOS and Android



**NetApp** 9 © 2023 NetApp, Inc. All rights reserved.

Active IQ provides predictive analytics and support for the hybrid cloud. Along with an inventory of NetApp systems, Active IQ provides a predictive health summary and trends. You also get improved storage efficiency information and a system-risk profile.

Access Active IQ either from the NetApp Support Site or from the Active IQ mobile app.

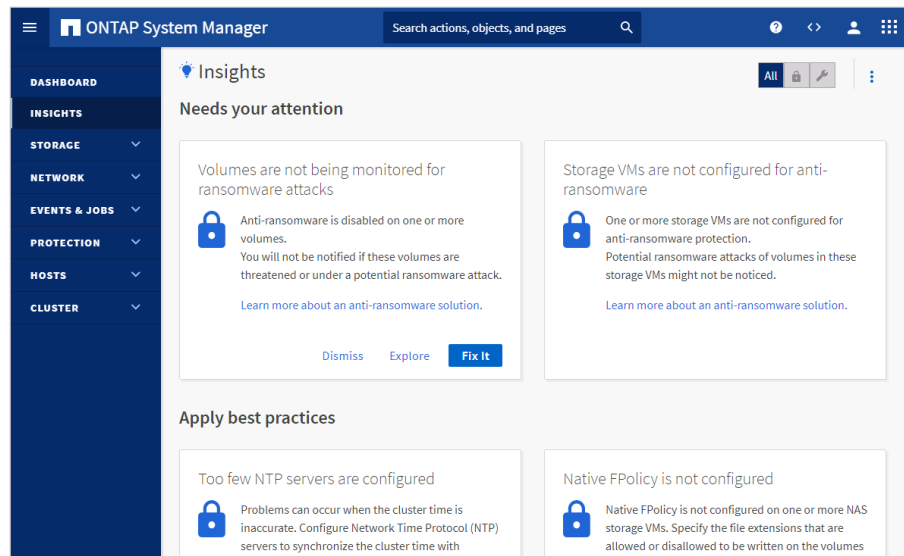
The dashboard Wellness widget displays information about the performance, efficiency, capacity, configuration settings, security vulnerabilities, and others. It proactively determines the systems that have either exceeded the capacity or are near exceeding 90% capacity. In addition, it provides information about software and hardware that have either expired or set to expire in the next 6 months.



# Active IQ in NetApp ONTAP System Manager

## Insights to optimization

- View insights that help you optimize the capacity, security, compliance, and configuration of your system.
- Respond to insights by:
  - Dismissing them
  - Exploring different ways to remediate the problems
  - Initiating the process to fix the problems
- Choose which insights to be informed of.



NetApp 10 © 2023 NetApp, Inc. All rights reserved.

With System Manager, you can view insights that help you optimize the capacity, security, compliance, and configuration of your system.

Based on best practices, these insights are displayed on one page from which you can initiate immediate actions to optimize your system.

In System Manager, you can respond to insights by either dismissing them, exploring different ways to remediate the problems, or initiating the process to fix the problems. Hover over an insight to reveal the buttons to perform the following actions:

- **Dismiss:** Remove the insight from the view.
- **Explore:** Find out various ways to remediate the problem mentioned in the insight. This button appears only if there is more than one method of remediation.
- **Fix:** Initiate the process of remediating the problem mentioned in the insight. You will be asked to confirm whether you want to take the action needed to apply the fix.

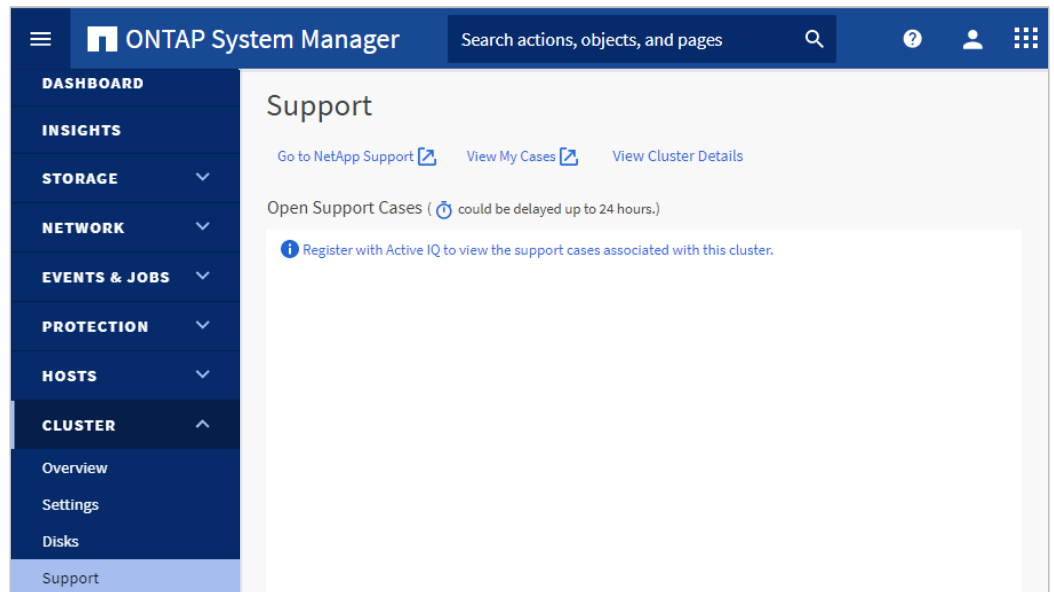
Some of these actions can be initiated from other pages in System Manager, but the Insights page helps you streamline your day-to-day tasks by allowing you to initiate these actions from this one page.

You can customize which insights you will be notified about in System Manager. On the Settings page, ensure there is a check in the check boxes next to the insights you want to be notified about. If you previously dismissed an insight, you can “undismiss” it by selecting its check box.

## Active IQ in NetApp ONTAP System Manager

Manage cases

- View open Active IQ support cases for this cluster
- Generate information to create a new support case
- Quickly access the NetApp My AutoSupport and My Cases pages



NetApp 11 © 2023 NetApp, Inc. All rights reserved.

Starting with ONTAP 9.9.1 software, you can use Active IQ to view support cases that are associated with the cluster. You can also copy necessary cluster details to submit a new support case on the NetApp Support Site.

With ONTAP 9.10.1 software, you can now use System Manager to review risk mitigation recommendations and acknowledge the risks reported by Active IQ.

To view open support cases for the cluster, you must first register the cluster with Active IQ:

1. Click **Go to NetApp Support Site** to navigate to the My AutoSupport page on the NetApp Support Site, where you can view knowledge base articles or submit a new support case.
2. Click **View My Cases** to navigate to the My Cases page on the NetApp Support Site.
3. Click **View Cluster Details** to view and copy information you need when you submit a new support case.

## BlueXP control plane


Deploy NetApp data management and data services in minutes

Manage your data replication from the data center to the cloud

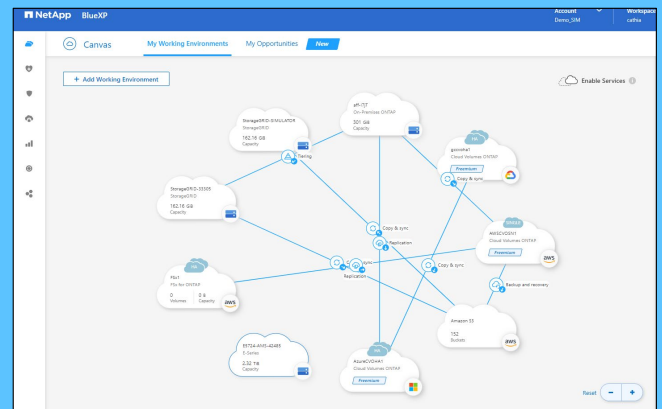
Achieve data security through NetApp managed encryption

Build hybrid environments quickly

Do it all with ease by using the BlueXP intuitive user interface

 12 © 2023 NetApp, Inc. All rights reserved.

**NetApp BlueXP** gives you a single control plane for **NetApp data management and data services.**



Automate data management and application deployment from the edge to the core to cloud.

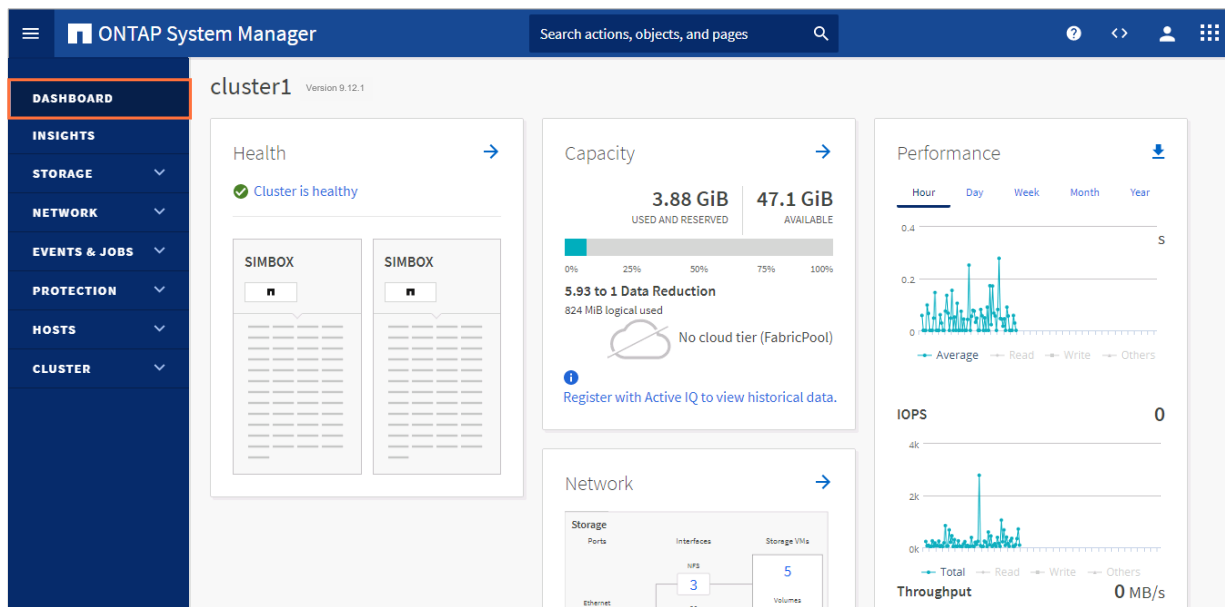
NetApp BlueXP enables you to build, protect, and govern your hybrid multicloud data from a single control plane. BlueXP enables you to manage all your storage and data assets from a single interface. You can use BlueXP to create and administer cloud storage (for example, NetApp Cloud Volumes ONTAP and Azure NetApp Files). You can use data services (for example, NetApp Cloud Backup and Cloud Data Sense) and control many on-premises and edge storage devices. You can use BlueXP to manage NetApp on-premises ONTAP software, E-Series systems, and StorageGRID systems.

Beginning with ONTAP 9.12.1 software, System Manager is fully integrated with BlueXP. From BlueXP, you have access to the System Manager interface that you know and access to BlueXP functionality.

When you connect to the cluster management network interface from your web browser, you are prompted to manage the cluster with System Manager in BlueXP or to use System Manager directly.

If you choose BlueXP, you are prompted to log in by using your NetApp support credentials. If you do not already have a NetApp account, you can create one. Then enter the cluster administrator credentials for the ONTAP cluster. You can now manage the ONTAP cluster from BlueXP.

## ONTAP System Manager dashboard



NetApp 13 © 2023 NetApp, Inc. All rights reserved.

ONTAP System Manager provides a simple graphical interface for the management of ONTAP systems. System Manager is built into ONTAP and available with every deployment.

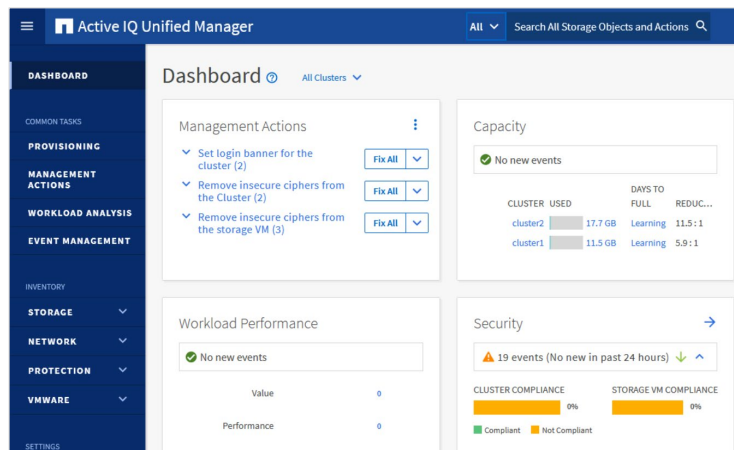
The System Manager dashboard shows at-a-glance system status for a storage system. The dashboard displays vital storage information, including efficiency and capacity use for various storage objects, such as aggregates and volumes.

NetApp and Ansible have partnered to develop a collection of automation and storage configuration management modules to easily provision, deploy, and manage NetApp storage systems. NetApp playbooks combine Ansible modules to deliver full-stack presentations of storage to the host. The recommendation is to use Ansible modules for NetApp, and to use sample playbooks and collections in conjunction with AWX or Ansible Tower for new automation projects. For more information and documentation about NetApp modules for Ansible, go to thePub, and attend the Automate Storage Administration Using ONTAP REST API and Ansible course.

# Active IQ Unified Manager

Application for monitoring multiple NetApp storage systems

- Works with ONTAP System Manager on each storage system
- Supports plug-in modules to extend functionality



To learn more about Unified Manager and how it integrates with ONTAP, enroll in the instructor-led course *Administration of Active IQ Unified Manager*.

NetApp 14 © 2023 NetApp, Inc. All rights reserved.

Think of Unified Manager as the big brother of ONTAP System Manager. Unified Manager can manage multiple clusters and opens System Manager when you navigate to a specific node.

The Unified Manager web UI enables a storage administrator, cluster administrator, or storage VM (storage virtual machine, also known as SVM) administrator to monitor and troubleshoot cluster or storage VM issues that relate to data-storage capacity, availability, performance, and protection.

For the user guide, see [https://library.netapp.com/ecm/ecm\\_download\\_file/ECMP1653271](https://library.netapp.com/ecm/ecm_download_file/ECMP1653271). NetApp Learning Services offers courses that focus on the configuration and use of these tools.

## Cloud Insights

- From public cloud to private data center, you can see your entire application infrastructure stack in one place.
- With NetApp Cloud Insights, you can monitor, troubleshoot, protect, and maximize the efficiencies of all your resources.
- You can reclaim unused, abandoned, or overprovisioned resources.
- Predictive analytics help you to proactively fix issues before they arise.
- You can protect your data with early detection and actionable intelligence on ransomware and other malware incursions.



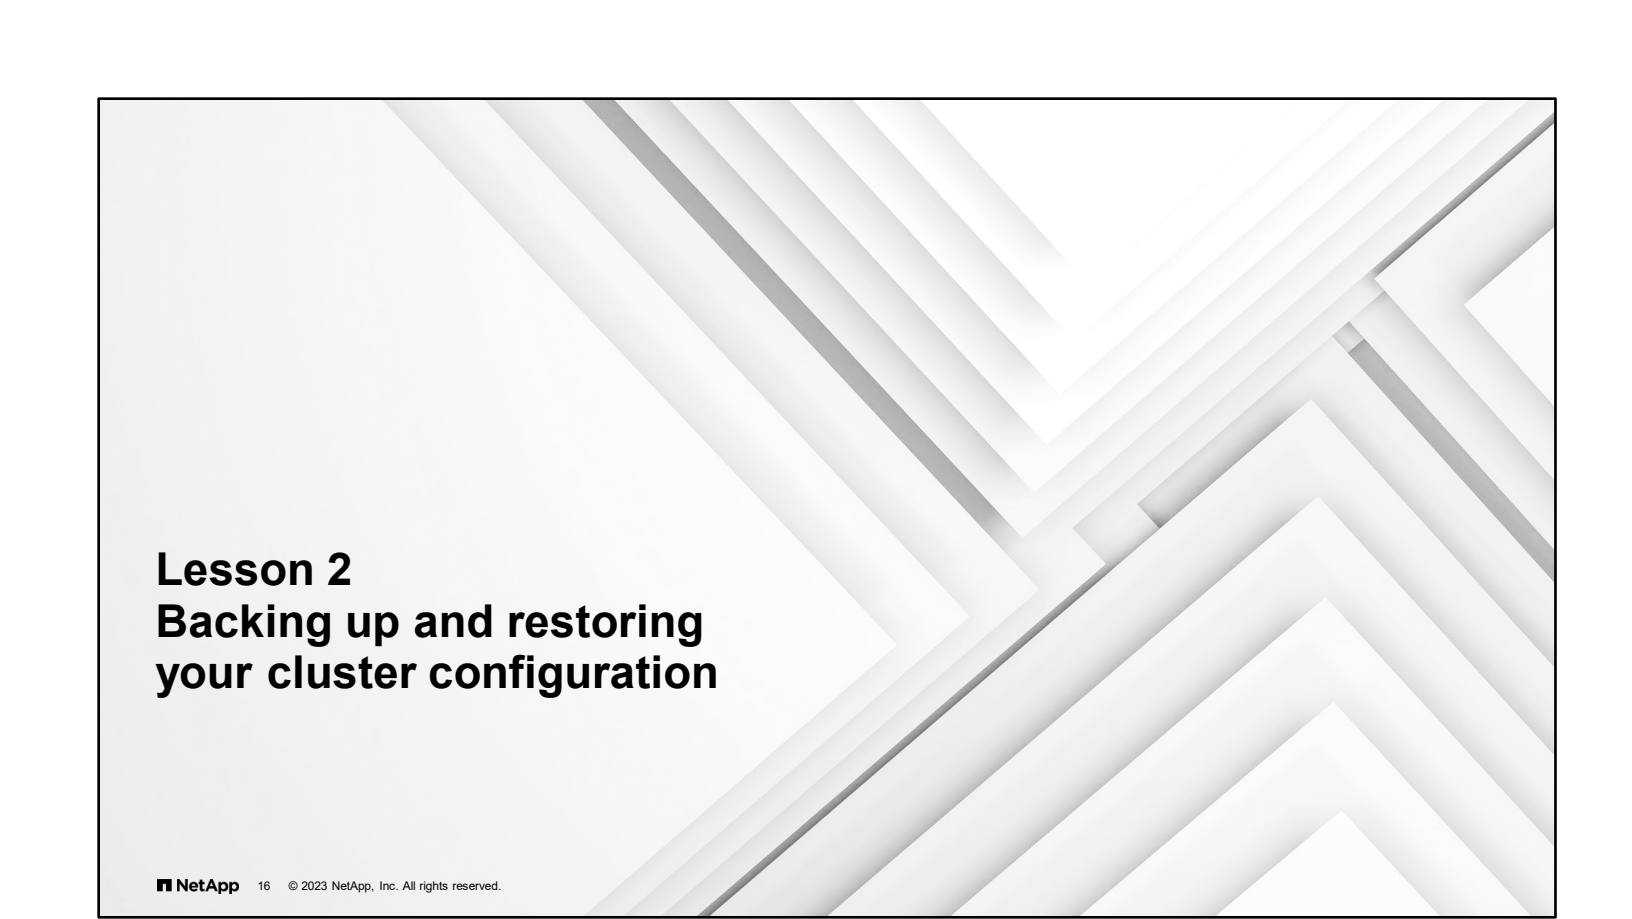
NetApp 15 © 2023 NetApp, Inc. All rights reserved.

NetApp Cloud Insights is a monitoring and reporting tool that you can use for your entire hybrid cloud infrastructure, from application to archive.

Cloud Insights is available through a subscription. You can use Cloud Insights to monitor your entire IT infrastructure, whether physical or virtual, from NetApp or hundreds of other vendors. Cloud Insights includes a gallery of predefined dashboards that enable you to benefit from the value of Cloud Insights immediately. You can customize the dashboards and tailor them to your environment to view only the information that is important to you.


Cloud Insights helps you to identify and fix problems by using end-to-end analytics. This capability can dramatically reduce the resolution time and the effect on users. Cloud Insights can also reduce your exposure to malicious use by protecting your data with actionable intelligence. Cloud Insights provides cost control and resource optimization across your hybrid environment.

NetApp Learning Services has multiple courses that discuss all the features and functionality of Cloud Insights. You should start with the fundamentals course.



## **Lesson 2**

# **Backing up and restoring your cluster configuration**

 16 © 2023 NetApp, Inc. All rights reserved.

## Cluster configuration backup files

Backing up the cluster configuration enables you to restore the configuration of any node or the entire cluster in a disaster or emergency.

- Configuration backup files are archive files (.7z) that contain information for all configurable options that are necessary for the cluster and cluster nodes to operate properly.
- Two types of configuration backup files exist:
  - Node
  - Cluster
- Configuration backup files do *not* include any user data.

Configuration backup files are archive files (.7z). These files contain information for all the configurable options that are necessary for a cluster and cluster nodes to operate properly. These files store the local configuration of each node and the cluster-wide replicated configuration. Use configuration backup files to restore the configuration of your cluster.

There are two types of configuration backup files:

### Node configuration backup file

Each healthy node in the cluster includes a node configuration backup file. This file contains all the configuration information and metadata that are necessary for the node to operate as healthy in the cluster.

### Cluster configuration backup file

These files include an archive of all the node configuration backup files in the cluster and the replicated cluster configuration information (the replicated database, or RDB files). Cluster configuration backup files enable you to restore the configuration of the entire cluster or of any node in the cluster. The cluster configuration backup schedules create these files automatically and store them on several nodes in the cluster.


**NOTE:** Configuration backup files contain only configuration information. They do not include any user data. For information about restoring user data, see the *Data Protection Power Guide*.



## Cluster backup scheduling

- ONTAP software automatically creates the configuration backup files every 8 hours, daily, and weekly.
- Use the `system configuration backup` command to manage cluster and node configuration backup files and backup schedules and to perform a configuration restore.
- Before you restore a node or cluster configuration, *always* see the *ONTAP 9 System Administration Reference* and contact NetApp technical support.

There might be discrepancies between the configuration backup file and the configuration that is present in the cluster.



Three separate schedules automatically create cluster and node configuration backup files and replicate them among the nodes in the cluster. The configuration backup files are automatically created according to the following schedules:

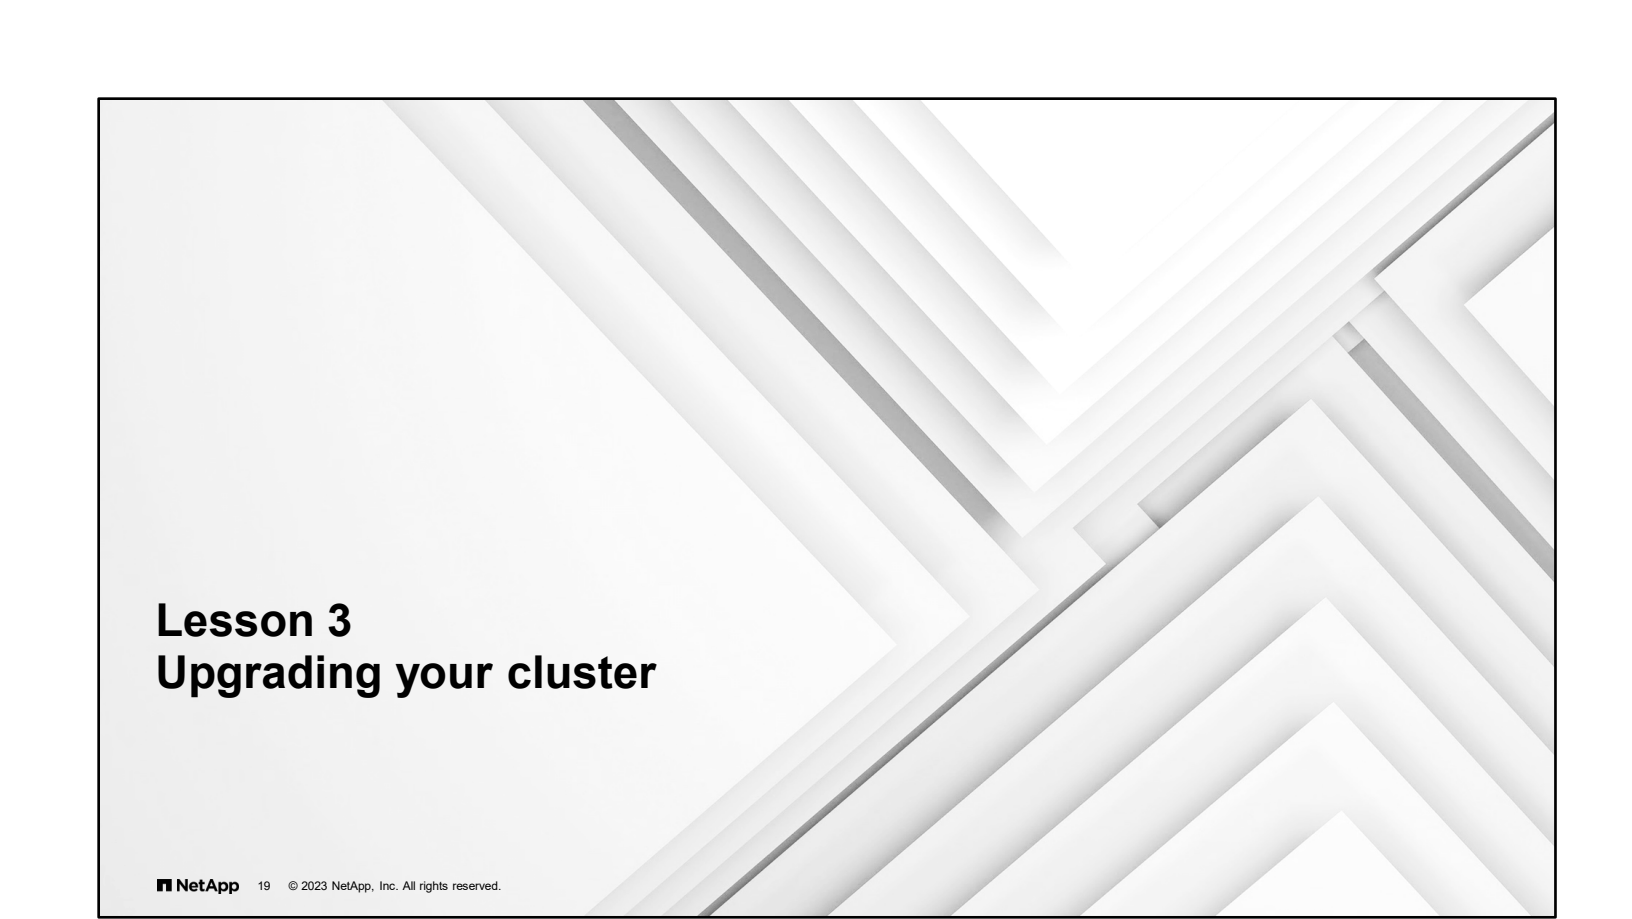
- Every 8 hours
- Daily
- Weekly

At each of these times, a node configuration backup file is created on each healthy node in the cluster. All these backup files are then collected into one cluster configuration backup file with the replicated cluster configuration, which is saved on one or more nodes in the cluster.

You can specify the configuration backup destination during software setup. After setup, you can use ONTAP commands to modify those settings.


Use the `system configuration backup` commands to manage cluster and node configuration backup files and backup schedules and to perform a configuration restore. You should only perform a restore to recover from a disaster that resulted in the loss of the entire cluster configuration.

**Attention:** If you are re-creating the cluster from a configuration backup file, you must contact technical support to resolve any discrepancies between the configuration backup file and the configuration in the cluster.

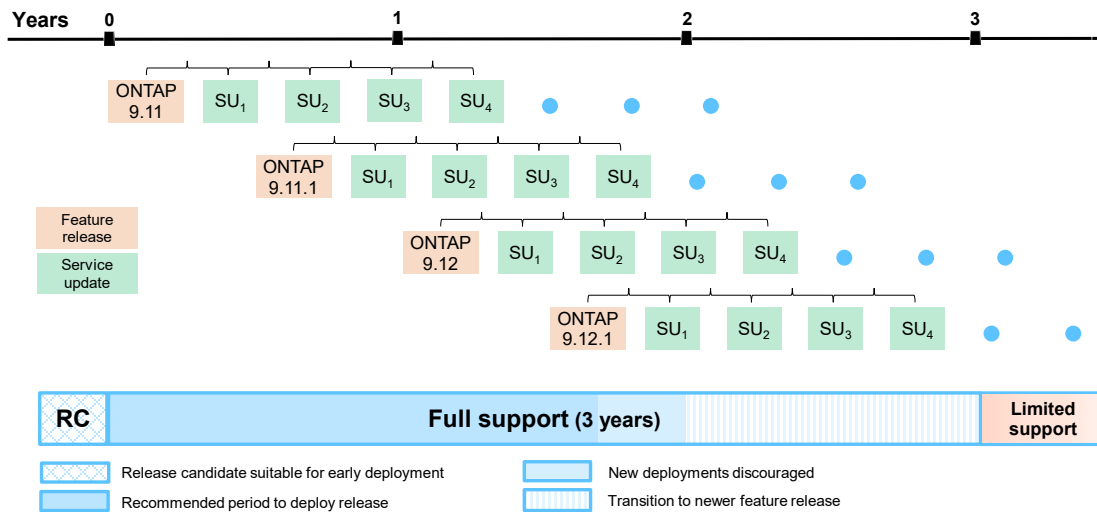


## Lesson 3

# Upgrading your cluster

 19 © 2023 NetApp, Inc. All rights reserved.

## Release support



NetApp 20 © 2023 NetApp, Inc. All rights reserved.

ONTAP software is keeping the twice-yearly release cadence by providing spring and fall releases to keep up with the need for new and rapidly developing features. In addition, the cloud versions of ONTAP software use different release numbers. ONTAP 9.12 software signifies the cloud release of ONTAP software and ONTAP 9.12.1 software is the on-premises release. With this consistent release schedule, you receive new versions of ONTAP software that contain bug fixes and new features. These new features help you to realize the benefits of a data fabric enabled by NetApp.

### Feature Release

Each feature release includes a set of new market-driven features and fixes for bugs that customers have encountered in earlier releases.

A feature release typically includes the following:

- New ONTAP feature content
- Key infrastructure changes, such as fundamental changes to NetApp WAFL operation or RAID operation
- Support of new NetApp engineered systems and replacement hardware components

Each feature release is numbered ONTAP x.y; for example, ONTAP 9.11 or ONTAP 9.12.

### Service Update

A service update release delivers timely fixes for critical field-encountered bugs for customers who cannot wait for the next feature release. These service updates are issued after a feature release is designated as general availability (GA). NetApp recommends that customers always adopt the latest service updates into their environment.

Each service update is numbered ONTAP x.yPz; for example, ONTAP 9.12P1 or ONTAP 9.12P2.

Bug fixes are cumulative. Every fix that is contained in ONTAP 9.12P1 is also contained in ONTAP 9.12P2.

As a best practice, you should run the recommended release up to 2 years after GA. During that period, NetApp strongly advises you to apply service updates every 6 months (or more often). By applying service updates, you incorporate timely fixes for critical bugs.

For more information about the software version cadence for ONTAP 9 software, see

<https://blog.netapp.com/ontap-upgrades?linkId=100000013429710>.

## Learning about new ONTAP features

- NetApp recommends upgrading your systems to the latest ONTAP general availability (GA) release.
- To simplify upgrades, ONTAP now allows you to skip two major revisions.
- How can you learn what changed since your current running version?
  - The release notes in the ONTAP documentation
  - The *What Is New in ONTAP <version #>* online courses
  - The CLI Comparison Tool: <https://mysupport.netapp.com/site/info/cli-comparison>

To benefit from the latest features and fixes, NetApp recommends that you upgrade your systems to the latest GA feature release of ONTAP software. Starting with ONTAP 9.6 software, every feature release is supported as a Long-Term Service release.

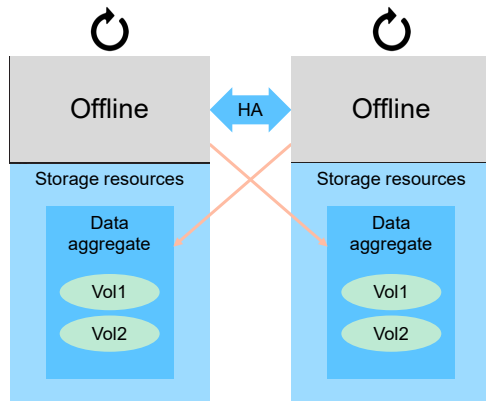
Before you decide to upgrade, do your due diligence and learn whether any changes will affect your environment, either positively or negatively. Each major version of ONTAP software maintains a set of release notes, which is expanded with each minor release. For a short overview of the key new features and changes in a release, take the associated *What's New* online course, available through NetApp Learning Services. If you use the command line extensively, the CLI Comparison Tool is a great resource for comparing changes that are made to commands between releases.

# Upgrade Advisor

Upgrade Advisor, which is part of Active IQ, simplifies the process of planning ONTAP upgrades. NetApp strongly recommends that you generate an upgrade plan from Upgrade Advisor before you upgrade your cluster.

You submit your system identification and target release to Upgrade Advisor. The tool compares AutoSupport data about your cluster to known requirements and limitations of the target release. Upgrade Advisor then generates an upgrade plan (and optionally a backout plan) with recommended preparation and execution procedures.

## Rolling upgrade



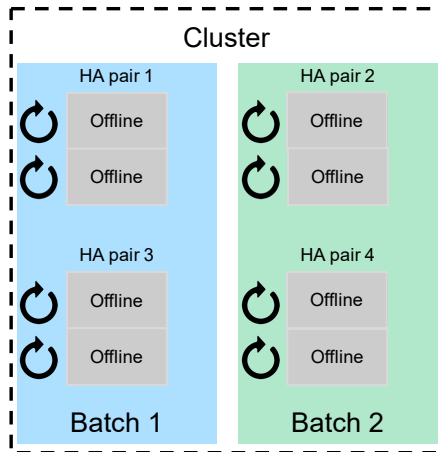
To upgrade software in a cluster of two or more nodes, complete the following steps:

1. Have the high-availability (HA) partner take control of the storage resources.
2. Take the node that is being upgraded offline.
3. Wait as the node reboots and is upgraded.
4. After the upgrade is complete, verify that the failed-over resources are returned home.
5. Repeat the process on the other node of the HA pair.
6. Repeat the process on other HA pairs.

You can perform rolling upgrades on clusters of two or more nodes. The upgrade runs on one node of a high-availability (HA) pair at a time. This approach makes it easier to roll back an upgrade in the unlikely event that an issue occurs.

The cluster does not switch over to the new version of ONTAP software until all nodes have installed the new version.

## Batch upgrade



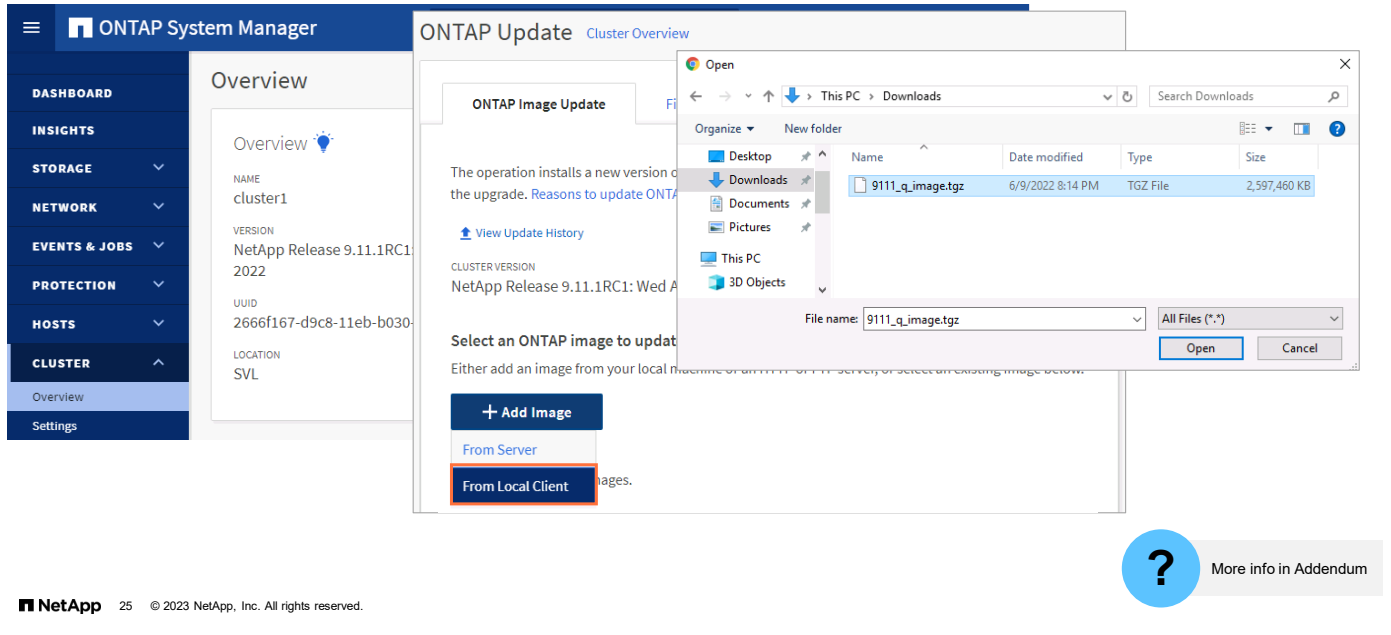
To upgrade software in a cluster of eight or more nodes, complete the following steps:

1. Separate the cluster into two batches, each of which contains multiple HA pairs.
2. In the first batch, take one node in each HA pair offline and upgrade the nodes while the partner nodes take over the storage.
3. After upgrades are completed on the first nodes, upgrade the other nodes of the HA pairs.
4. Repeat the process on the second batch.

You can perform batch upgrades on clusters of eight or more nodes. Unlike rolling upgrades, batch upgrades can run on more than one HA pair at a time.

As in rolling upgrades, in a batch upgrade, the cluster does not switch over to the new version of ONTAP software until all nodes have installed the new version.

# Automated nondisruptive upgrade




Use CLI commands to perform rolling upgrades and batch upgrades. If your cluster meets all the conditions, you can use ONTAP System Manager to perform an automated nondisruptive upgrade (NDU) instead of using the CLI. Read the ONTAP 9 Upgrade documentation (<https://docs.netapp.com/us-en/ontap/upgrade/index.html>) to prepare your cluster, and then follow the simple wizard to get the package, validate, and start the upgrade process.

Beginning with ONTAP 9.10.1 software, if you have a cluster with eight or more nodes, you can select to have them updated one HA pair at a time in a rolling upgrade fashion. This enables you, if needed, to correct upgrade issues on the first HA pair before moving to subsequent pairs.



## Install and upgrade from a USB drive

- Many FAS and AFF systems support the installation of ONTAP software and firmware from a FAT32 formatted USB device to do the following:
  - Perform boot device recovery from the LOADER prompt
  - Copy ONTAP software for installation
  - Copy service images for a firmware update
- Use the `system node image command`.



More info in Addendum

You can also install ONTAP software and firmware from an external USB device on most FAS and AFF systems that shipped since late 2016.



## Try this task

From the clustershell on cluster1, type:

```
system node image show -instance
```

1. How many nodes are in your cluster?
2. Which version of ONTAP software is current on each node?
3. Can you tell which image is booted?

1. The cluster contains two nodes.
2. Some revision of ONTAP 9 software should be installed, but the revision varies.
3. Verify the image in the Image Name field.

## Cluster expansion

Complete the following steps in the CLI to add nodes to a multinode switched cluster:

1. Verify that the nodes are configured as HA pairs and are connected to the cluster interconnect.
2. Power on both nodes of the HA pair.
3. Start the Cluster Setup wizard on one of the nodes.
4. Use the `join` command and follow the wizard.
5. Repeat Steps 3 and 4 on the partner node.

```
::> cluster setup
```

```
Welcome to the cluster setup wizard.
```

```
You can enter the following commands at any time:  
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the cluster setup wizard.  
Any changes you made before quitting will be saved.
```

```
You can return to cluster setup at any time by typing "cluster  
setup".  
To accept a default or omit a question, do not enter a value.
```

```
Do you want to create a new cluster or join an existing cluster?  
{create, join}: join
```


You can expand an existing cluster by using the CLI to nondisruptively add nodes to the cluster.

You must add nodes from HA pairs that are connected to the cluster interconnect. Nodes are joined to the cluster one at a time.



## **Lesson 4**

# **Recommended practices for performance**

 29 © 2023 NetApp, Inc. All rights reserved.

## Performance considerations



Workloads



I/O operation types:

- Random
- Sequential



Quality of service (QoS)

Storage system performance calculations vary widely based on the kinds of operations, or workloads, that are being managed.

The storage system sends and receives information in the form of I/O operations. I/O operations can be categorized as either random or sequential. Random operations, such as database operations, are usually small, lack any pattern, and happen quickly. In contrast, sequential operations, such as video files, are large and have multiple parts that must be accessed in a particular order.

Some applications have more than one dataset. For example, a database application data files and log files might have different requirements. Data requirements might also change over time. For example, data might start with specific requirements that change as the data ages.

If more than one application shares the storage resources, each workload might need to have quality of service (QoS) restrictions imposed. QoS restrictions prevent applications or tenants from being either bullies or victims.

## Analyzing I/O

### IOPS

- I/O is measured in IOPS.
- IOPS measures *how many* requests are managed in 1 second.
- IOPS data is most useful if I/O has any of the following features:
  - I/O request patterns are random.
  - I/O requests are small.
  - Multiple I/O sources must be managed.

 31 © 2023 NetApp, Inc. All rights reserved.



Factors that affect IOPS include the balance of read and write operations in the system and whether traffic is sequential, random, or mixed. Other factors that affect IOPS include the application type, operating system, background operations, and I/O size.

Applications with a random I/O profile, such as databases and email servers, usually have requirements that are based on an IOPS value.

## Analyzing I/O

Latency and response time



Latency is measured in microseconds and milliseconds.




Latency is a measurement of how long data processing takes.



Response time is the elapsed time between an inquiry and the response to that inquiry.

Response time is a sum of all latency that is encountered between the inquiry and receipt of a response.

 32 © 2023 NetApp, Inc. All rights reserved.

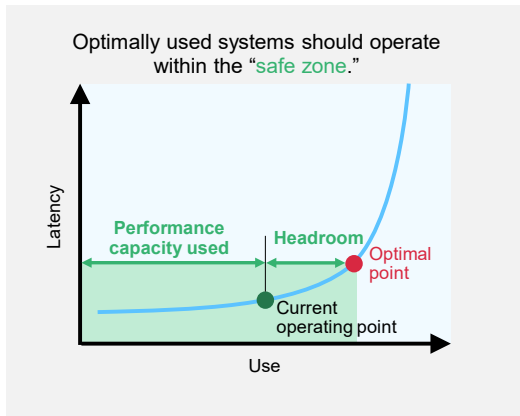
Latency is the measurement of how long a storage system takes to process an I/O task. Smaller latency values are better.

Latency for hard drives is typically measured in milliseconds. Because solid-state media is much faster than hard drives, the latency of the media is measured in submilliseconds or microseconds.

Response time is the elapsed time between an inquiry on a system and the response to that inquiry. Every mechanical and digital component along the way introduces some latency. All the latencies are added together to constitute the response time.

## Headroom and performance capacity used

Key for optimal use of a system



- **Optimal point:**
  - The maximum optimal operating point for a system
  - The point beyond which a small increase results in a bigger increase in latency
- **Headroom:**
  - A metric that is used in ONTAP 9 software
  - The remaining useful capacity of a resource, when measured from the optimal point
- **Performance capacity used:**
  - A metric that is used in Unified Manager
  - Equal to the optimal point minus headroom
  - Performance metric for node and aggregate

NetApp 33 © 2023 NetApp, Inc. All rights reserved.


QoS is effective in optimally used systems.

If you know the available performance capacity in the cluster, you can better provision to balance workflows. Performance capacity is how much work you can place on a node or an aggregate before latency affects the performance of all workloads. You can use Active IQ Unified Manager to identify available performance capacity.



## Maintain optimal operating point

Adding and relocating resources



Relocating resources nondisruptively:

- Moving volumes and LUNs
- Moving an aggregate between the nodes of an HA pair

NetApp 34 © 2023 NetApp, Inc. All rights reserved.


As well as discussing performance at the node level, discussing performance at the cluster level is important.

In the example, an administrator creates volumes on a 2-node cluster that is used for file services. The system is configured with NL\_SAS disks to meet the workload requirements.

After some time, the administrator needs to add a volume for a database application. The NL\_SAS disks do not meet the requirements for the new workload. The administrator decides, for future growth, to nondisruptively add another HA pair with SSDs. With new nodes with SSDs active in the cluster, the administrator can nondisruptively move the volume to the faster drives.

## Maintain optimal operating point

Quality of service



SVM = storage virtual machine or storage VM

- Key capability to *manage and control* performance
- Effective in *optimally* used systems
- Increasingly sought by both enterprise and service provider market segments

Use cases:

- Contain “runaway” workloads (QoS Max)
- Experience dedicated workload performance (QoS Min)
- Enable performance services classes

**NetApp** 35 © 2023 NetApp, Inc. All rights reserved.

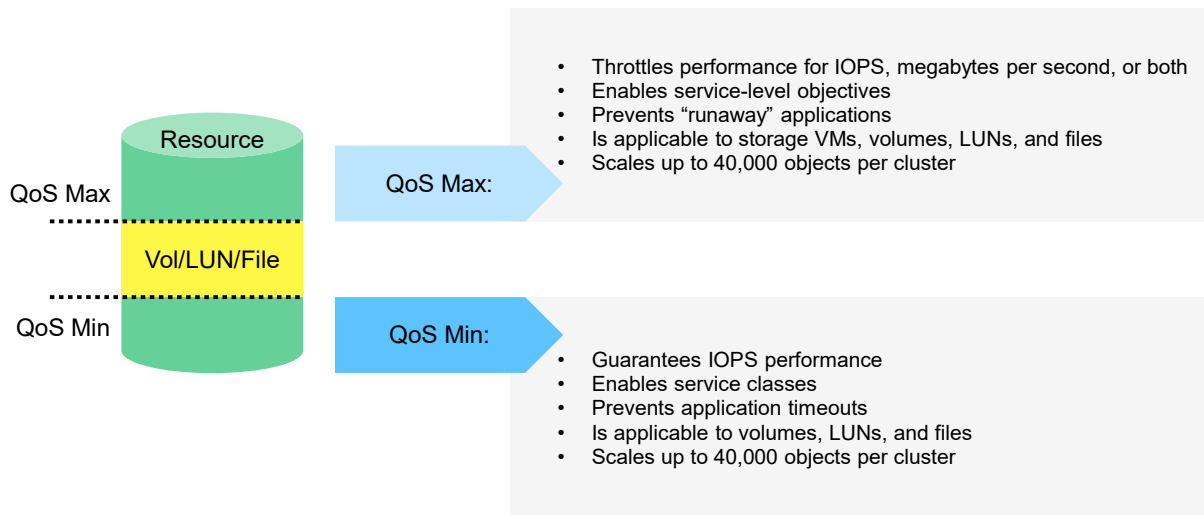
You can use storage QoS to deliver consistent performance by monitoring and managing application workloads.

You can configure the storage QoS feature to prevent user workloads or tenants from affecting one another. You can configure the feature to isolate and throttle resource-intensive workloads. The feature can also enable critical applications to achieve consistent performance expectations.

Essentially, QoS is about managing and controlling performance in heavily used systems. Both enterprise and service-provider market segments increasingly seek QoS.

## Managing workloads

Guaranteeing performance



NetApp 36 © 2023 NetApp, Inc. All rights reserved.

The goal of controlling performance in a shared storage environment is to provide dedicated performance for business-critical workloads against all other workloads. To guarantee performance, you must apply QoS policies on these resources.

QoS Max, which is used to contain runaway workloads, was introduced in Data ONTAP 8.2 software and has been continually enhanced. QoS Min, which provides a throughput floor, was introduced with ONTAP 9.2 software.

QoS Min (sometimes called a throughput floor or TP Floor) has the same scalability as QoS Max, with up to 12,000 policy groups managing up to 40,000 workloads per cluster. The major difference is that QoS Max can limit IOPS, megabytes per second, or both. QoS Min only guarantees IOPS performance. Also, QoS Min is applicable to volumes, LUNs, and files in a cluster. Storage VMs are not supported.

## Maximizing performance



Ways to minimize performance issues:

- Correctly size and follow recommended practices for the specific workload.
- Verify the supported minimums and maximums.
- Adhere to the ONTAP storage system mixing rules (Hardware Universe).
- Verify the compatibility of components, host operating system, applications, and ONTAP software using the (NetApp Interoperability Matrix Tool [IMT]).



Potential performance issues:

- **Controller:** Resource overuse, ONTAP version, offline, or rebooting
- **Storage:** Drive types, aggregate configuration, volume movement, or free space
- **Networking:** Configuration, LIF location, port saturation, port speeds, or indirect access
- **Host or clients:** Application, drivers, network adapter, or user knowledge

## Create free space in an aggregate

### Simple steps

A full aggregate affects performance and might lead to an inability to write new data. Use these no-risk measures to free space:

- Add drives to the aggregate.
- Move some volumes to another aggregate with available space.
- Enable space-saving features, such as deduplication or compression.



Even the most diligently watched storage systems can occasionally have an aggregate that fills up. The situation results in performance degradation and can eventually result in failed writes. You can take three simple steps to free space in a full aggregate:

- The easiest step is to grow the aggregate by adding drives. Be sure to leave adequate spare disks in the spare pool.
- Moving volumes to a less full aggregate takes some time but safely frees up space.
- If you have not enabled deduplication or compression, these efficiency features can make more space but require some time to run.

## Create free space in an aggregate

Complex steps

Use these measures with caution:

- Shrink the size of volume-guaranteed volumes in the aggregate.

You can do so manually, or you can use the `grow_shrink` option of the automatic resize capability.

- Change volume guarantee type to `none` on volumes that use large amounts of space so that the volumes take up less space in the aggregate.
- Delete unneeded volume Snapshot copies if the volume has a guarantee type of `none`.

**Note:** Blocks are returned to free space only when there are no pointers to the block. You might need to delete multiple Snapshot copies before you gain any space.

- Delete unneeded volumes.

The volume recovery queue holds a deleted volume for 12 hours. Contact NetApp technical support if you need to purge the queue sooner.



The following steps can involve some risk of future issues or potential unrecoverable loss of data:


- If space-guaranteed volumes with significant unused space exist, you can resize them to return some of the unused space. The potential risk is that the volume might run out of space and cause failed writes.
- Changing the volume guarantees to `none` removes space reservations in the aggregate for those volumes.
- Deleting old or unneeded Snapshot copies might free space. Only blocks that no longer have any pointers to them are returned to the free space. If multiple Snapshot copies reference a block, the block is not released until all the Snapshot copies are deleted. After a Snapshot copy is deleted, it can no longer be used to recover data.
- Deleting unneeded volumes carries the biggest risk. If you later discover that the volume is needed, you cannot recover the volume. One exception, which can also cause confusion, is that deleted volumes are held in a recovery queue for 12 hours. The recovery queue provides you with time to realize that a volume was deleted by mistake and recover it. If you and your users are certain that the volume is no longer needed and do not want to wait 12 hours, contact NetApp technical support for the procedure to purge the queue.

When freeing up space in an aggregate, follow the maxim to “measure twice and cut once” to avoid making the situation worse by deleting useful data.



# Lesson 5

## Technical support

 40 © 2023 NetApp, Inc. All rights reserved.

## System logs

- Log messages can be sent to the following:
  - Console
  - Message log
- You can access the message log by using the following:
  - `event log` command
  - System Manager
  - Active IQ OneCollect
  - Web browser:  
`https://<cluster-mgmt-ip>/spi/<nodename>/etc/log/`



Use the `event log show` command to browse the `messages.log` file.

The system log contains information and error messages that the storage system displays on the console and logs in message files.

On active systems the event log can become very large. Use the filtering capability of System Manager and the `event log show` command to locate log messages by time, severity level, cluster node, type, or other criteria.



## Manage log files

- Consider setting up log forwarding of systems and servers to a central system log (syslog) server.

Use the `cluster log-forwarding` command to set up forwarding on ONTAP clusters.

- Establish a schedule to roll forward the log files of applications, like PuTTY, that do not have the capability built in.
- Add the creation of dedicated log files to your maintenance checklists.



In addition to monitoring and reporting applications that watch over your systems, consider creating a dedicated system log (or syslog) server. A syslog server acts as a repository for log files and as a failsafe to your primary monitoring tools. You can configure your ONTAP clusters to forward their logs to up to 10 destinations by using the `cluster log-forwarding` command.

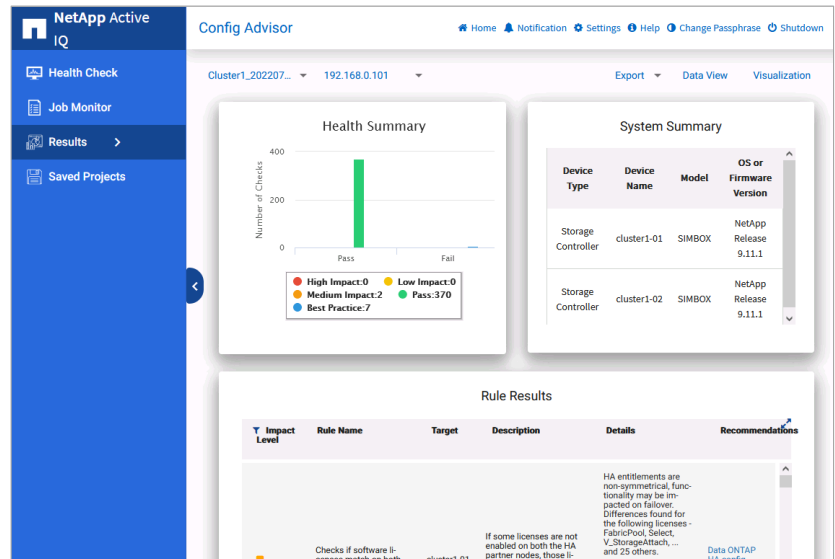
Log rolling is the process of closing out a log file before it becomes too large and cumbersome and opening a new log file. Many commercial applications that generate large or numerous log files do this automatically and retain three or more old log files. Applications like PuTTY, which are used intermittently, do not have this capability. To keep the log files from becoming unwieldy, create a schedule to manually roll the logs forward every month, quarter, or at least once each year. Archive the old log files so that you maintain a history. This information can be vital in tracking down a long-term issue that might have gone unnoticed.

Every time that you perform maintenance, include a copy of the log files with your records. Doing so is easier if you make the creation of a new log file at the start and end of a maintenance cycle part of your process. If you need to send log files to technical support, a dedicated log file has less noise for the technical support team to read through.

# Use Config Advisor before and after maintenance

## What is Config Advisor?

- Use to verify or troubleshoot cabling and configuration of cluster and switches.
- If desired, configure to run on a schedule.
- Download from the Support site and run from PC that is connected to the serial port or over the network.



NetApp 43 © 2023 NetApp, Inc. All rights reserved.

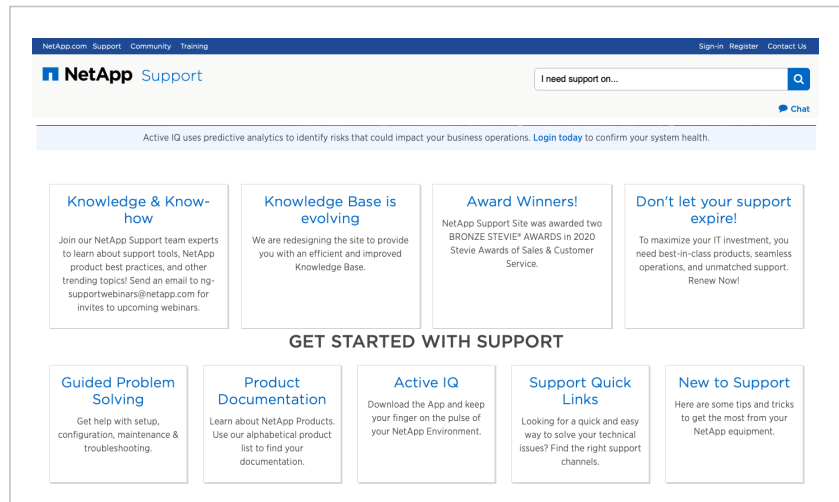
NetApp Active IQ Config Advisor is a configuration validation and health check tool for NetApp systems. You can deploy this tool at secure sites and nonsecure sites for data collection and analysis. You can use Config Advisor to check a NetApp system or FlexPod solution for the correctness of hardware installation and conformance to NetApp recommended settings. Config Advisor collects data and runs a series of commands on the hardware. The tool then checks for cabling, configuration, availability, and best practice issues.

The time that Config Advisor spends collecting data depends on how large the cluster is but is usually just minutes.

The View and Analyze tab shows the results of the data collection. The first panel enables you to look into error messages based on severity. The next panel shows an inventory of all queried devices. The Visualization panel is a visual depiction of how the systems are cabled. The last panel displays total storage available and how it is used.

Download Config Advisor from ToolChest on the NetApp Support Site.

# NetApp Support



- **NetApp Support Site:** [mysupport.netapp.com](https://mysupport.netapp.com)
- **Hardware Universe:** [hwu.netapp.com](https://hwu.netapp.com)
- **IMT:** [mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix)

For support information, documentation, software downloads, and access to Active IQ, see NetApp Support at [mysupport.netapp.com](https://mysupport.netapp.com).

For system configuration information, see the NetApp Hardware Universe at [hwu.netapp.com](https://hwu.netapp.com).

To determine the compatibility between various NetApp and third-party products that are officially supported, see the NetApp Interoperability Matrix Tool (IMT) at [mysupport.netapp.com/matrix](https://mysupport.netapp.com/matrix).

## Bug tools and reports

Stay up to date about bugs and bug fixes with the tools on the NetApp Support Site.

<https://mysupport.netapp.com/site/bugs-online/product>

- Bug Search
- Release Bug Comparison Tool
- Release Bug Advisor
- Bug Watcher Summary
- New Bug Alerts Profiler


The screenshot shows the NetApp Support Site interface for searching bugs. The page title is "Selected product: ONTAP". There is a search bar with "Search by Bug ID" and a dropdown menu set to "9.8". A "NEW SEARCH" button is visible. Below the search bar, there are options to "Export to CSV" and "Export to Excel". A table titled "My favored ONTAP Bugs" is shown, but it contains no data. A dropdown menu is open over the table, showing options: "All", "New", "In Progress", "Rejected", "Fixed", and "Duplicate".

Like all software applications, ONTAP software occasionally comes with “undocumented features” that might negatively affect your cluster. The NetApp Support Site contains many ways to research and stay current on bugs and find known solutions or workarounds.

See <https://mysupport.netapp.com/site/bugs-online/product>.

## Recommended preventive maintenance checklist


- ✓ Replace failed components as soon as possible.
- ✓ View weekly AutoSupport and health checks in Active IQ on the NetApp Support Site.
- ✓ Run Config Advisor once each month to detect cabling issues.
- ✓ Read the release notes for new versions of ONTAP software to determine whether you can benefit from new features or bug fixes.
- ✓ Twice each year, verify Return Material Authorization (RMA) contact information and the expiration date of the support contracts.
- ✓ Change the Cluster and SVM Admin passwords at least twice each year.

 46 © 2023 NetApp, Inc. All rights reserved.

A properly configured NetApp storage system can be run with a set-it-and-forget-it mentality. But just like an automobile, the system runs better and more reliably with regular maintenance.

# Knowledge check

Module 10: Cluster maintenance

 47 © 2023 NetApp, Inc. All rights reserved.

## Knowledge check

**Where do you find the Upgrade Advisor tool, which you can use to plan a NetApp ONTAP upgrade?**

- a. in NetApp ONTAP System Manager
- b. on the Tools page on the NetApp Support Site
- c. on the Downloads page at [upgradeontap.netapp.com](https://upgradeontap.netapp.com)
- d. on the Active IQ page on the NetApp Support Site

## Knowledge check

**Which three intervals does NetApp ONTAP software follow when creating cluster configuration backup files? (Choose three.)**

- a. every hour
- b. every 8 hours
- c. every 12 hours
- d. daily
- e. weekly
- f. monthly



## References

- ONTAP 9 Documentation Center  
<http://docs.netapp.com/ontap-9/index.jsp>

## Module summary

This module focused on enabling you to do the following:

- Use the Active IQ customer dashboard
- Plan ONTAP software upgrades
- Follow recommended practices for peak performance
- Configure event notifications and alerts
- Prepare to engage NetApp technical support
- Perform cluster maintenance



## Complete an exercise

Module 10  
Cluster maintenance

### Installing and Configuring Active IQ Config Advisor


- Access your lab equipment.
- Open your Exercise Guide to Module 10.
- Complete Exercise 1.
- Share your results.

This exercise requires approximately  
**20 minutes.**

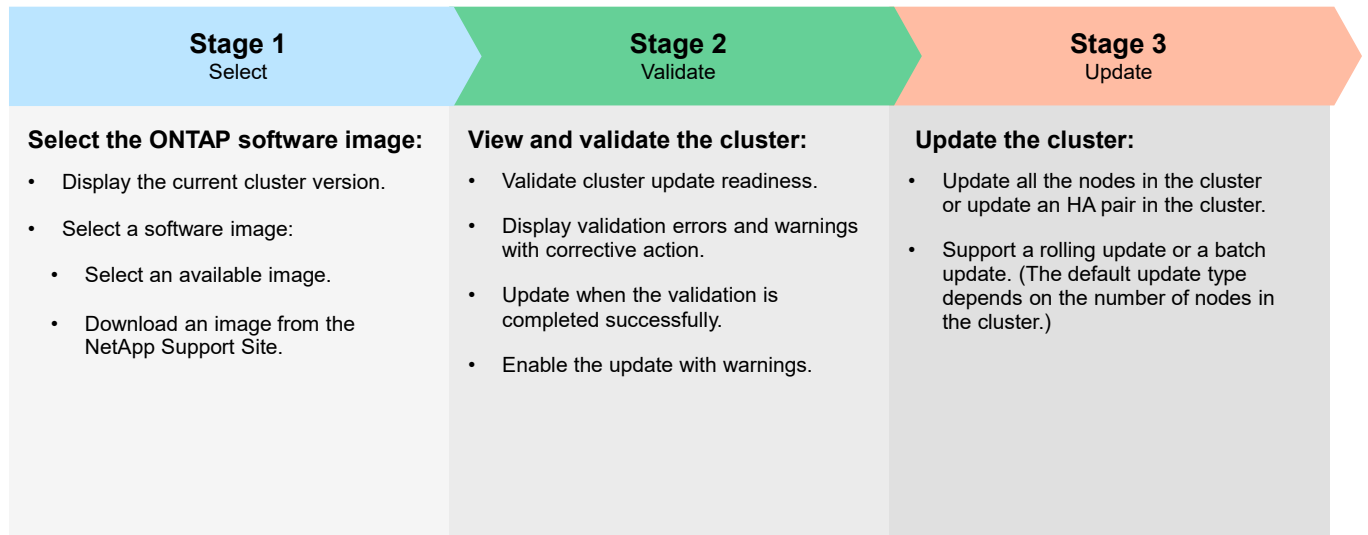
See the instructions in your Exercise Guide.



# Addendum ONTAP software upgrades

 53 © 2023 NetApp, Inc. All rights reserved.

## Stages of an automated upgrade



54 © 2023 NetApp, Inc. All rights reserved.

The automated upgrades that you can perform by using System Manager consist of three stages: select, validate, and update.

In the select stage, you select the ONTAP software image. The current version details are displayed for each node or HA pair.

In the validate stage, you view and validate the cluster against the software image version for the update. A pre-update validation helps you to determine whether the cluster is ready for an update. If the validation is completed with errors, a table displays the status of the various components and the required corrective actions. You can perform the update only when the validation is completed successfully.

In the final update stage, you update either all the nodes in the cluster or update an HA pair in the cluster to the selected version of the software image. While the update is in progress, you can pause and then either cancel or resume the update. If an error occurs, the update is paused, and an error message is displayed with the remedial steps. You can resume the update after performing the remedial steps or cancel the update. You can view the table with the node name, uptime, state, and ONTAP software version when the update is successfully completed.

## USB port use cases

Scenario	Prerequisites	Command
Perform boot device recovery from the LOADER prompt.	<ul style="list-style-type: none"> <li>The USB 2.0 device is formatted to FAT32 with the correct ONTAP image.tgz file.</li> <li>The device is not hot-pluggable. After you insert the USB device, you must boot to the LOADER prompt.</li> </ul>	<ul style="list-style-type: none"> <li>At the LOADER prompt, use <code>boot_recovery</code> by using the netboot image.</li> <li>At the boot menu, select the appropriate ONTAP image.</li> </ul>
Copy ONTAP software for installation.	The USB 2.0 device is formatted to FAT32 with the correct ONTAP image.	<ul style="list-style-type: none"> <li>Use the <code>system node image update/get</code> command.</li> <li>From the additional options for the command, copy ONTAP software from the USB device.</li> </ul>
Copy service images for firmware update.	The USB 2.0 device is formatted to FAT32 with the correct service image.	<ul style="list-style-type: none"> <li>Use the <code>system node firmware download</code> command.</li> <li>From the additional options for the command, copy ONTAP software from the USB device.</li> </ul>

The chart shows scenarios in which you can use the USB port. Each scenario has prerequisite considerations. The Command column shows you the commands to use in each scenario.