



Storage

Cloud Volumes ONTAP

NetApp
February 27, 2026

Table of Contents

- Storage 1
 - Supported client protocols for Cloud Volumes ONTAP 1
 - iSCSI 1
 - NFS 1
 - SMB 1
 - S3 1
 - NVMe-TCP 1
 - Disks and aggregates used for Cloud Volumes ONTAP clusters 1
 - Overview 2
 - AWS storage 2
 - Azure storage 3
 - Google Cloud storage 4
 - RAID type 4
- Learn about support for AWS Elastic Volumes with Cloud Volumes ONTAP 4
 - Benefits 4
 - Supported configurations 5
 - Required AWS permissions 5
 - How support for Elastic Volumes works 5
 - Limitations 8
 - How to work with Elastic Volumes 8
- Learn about data tiering with Cloud Volumes ONTAP in AWS, Azure, or Google Cloud 10
 - Data tiering in AWS 11
 - Data tiering in Azure 12
 - Data tiering in Google Cloud 13
 - Data tiering and capacity limits 13
 - Volume tiering policies 14
 - Setting up data tiering 15
- Cloud Volumes ONTAP storage management 15
 - Storage provisioning 15
 - Capacity management 16
- Write speed 17
 - Normal write speed 17
 - High write speed 17
 - How to select a write speed 18
 - What to expect if data loss occurs 19
 - How to stop data access if data loss occurs 19
- Flash Cache 20
 - What's Flash Cache? 20
 - Supported configurations 20
 - Limitations 20
- Learn about WORM storage on Cloud Volumes ONTAP 20
 - How WORM storage works 21
 - Activating WORM storage 21

Committing files to WORM	21
Enabling WORM on a Cloud Volumes ONTAP system	21
Deleting WORM files	22
WORM and data tiering	22
Limitations	22

Storage

Supported client protocols for Cloud Volumes ONTAP

Cloud Volumes ONTAP supports the iSCSI, NFS, SMB, NVMe-TCP, and S3 client protocols.

iSCSI

iSCSI is a block protocol that can run on standard Ethernet networks. Most client operating systems offer a software initiator that runs over a standard Ethernet port.

NFS

NFS is the traditional file access protocol for UNIX and LINUX systems. Clients can access files in ONTAP volumes using the NFSv3, NFSv4, and NFSv4.1 protocols. You can control file access using UNIX-style permissions, NTFS-style permissions, or a mix of both.

Clients can access the same files using both NFS and SMB protocols.

SMB

SMB is the traditional file access protocol for Windows systems. Clients can access files in ONTAP volumes using the SMB 2.0, SMB 2.1, SMB 3.0, and SMB 3.1.1 protocols. Just like with NFS, a mix of permission styles are supported.

S3

Cloud Volumes ONTAP supports S3 as an option for scale-out storage. S3 protocol support enables you to configure S3 client access to objects contained in a bucket in a storage VM (SVM).

[ONTAP documentation: Learn how S3 multiprotocol works.](#)

[ONTAP documentation: Learn how to configure and manage S3 object storage services in ONTAP.](#)

NVMe-TCP

Beginning with ONTAP version 9.12.1, NVMe-TCP is supported for all cloud providers. Cloud Volumes ONTAP supports NVMe-TCP as a block protocol for storage VMs (SVMs) during deployment, and installs the required NVMe licenses automatically.

NetApp Console does not provide any management capabilities for NVMe-TCP.

For more information on configuring NVMe through ONTAP, refer to the [ONTAP documentation: Configure a storage VM for NVMe.](#)

Disks and aggregates used for Cloud Volumes ONTAP clusters


Understanding how Cloud Volumes ONTAP uses cloud storage can help you understand your storage costs.



You must create and delete all disks and aggregates from the NetApp Console. You should not perform these actions from another management tool. Doing so can impact system stability, hamper the ability to add disks in the future, and potentially generate redundant cloud provider fees.

Overview

Cloud Volumes ONTAP uses cloud provider storage as disks and groups them into one or more aggregates. Aggregates provide storage to one or more volumes.



Several types of cloud disks are supported. You choose the disk type when you create a volume and the default disk size when you deploy Cloud Volumes ONTAP.



The total amount of storage purchased from a cloud provider is the *raw capacity*. The *usable capacity* is less because approximately 12 to 14 percent is overhead that is reserved for Cloud Volumes ONTAP use. For example, if the Console creates a 500 GiB aggregate, the usable capacity is 442.94 GiB.

AWS storage

In AWS, Cloud Volumes ONTAP uses EBS storage for user data and local NVMe storage as Flash Cache on some EC2 instance types.

EBS storage

In AWS, an aggregate can contain up to 6 disks that are all the same size. But if you have a configuration that supports the Amazon EBS Elastic Volumes feature, then an aggregate can contain up to 8 disks. [Learn more about support for Elastic Volumes.](#)

The maximum disk size is 16 TiB.

The underlying EBS disk type can be either General Purpose SSDs (gp3 or gp2), Provisioned IOPS SSD (io1), or Throughput Optimized HDD (st1). You can pair an EBS disk with Amazon Simple Storage Service (Amazon S3) to [low-cost object storage](#).



Tiering data to object storage is not recommended when using Throughput Optimized HDDs (st1).

Local NVMe storage

Some EC2 instance types include local NVMe storage, which Cloud Volumes ONTAP uses as [Flash Cache](#).

Related links

- [AWS documentation: EBS Volume Types](#)
- [Learn how to choose disk types and disk sizes for your systems in AWS](#)
- [Review storage limits for Cloud Volumes ONTAP in AWS](#)
- [Review supported configurations for Cloud Volumes ONTAP in AWS](#)

Azure storage

In Azure, an aggregate can contain up to 12 disks that are all the same size. The disk type and maximum disk size depends on whether you use a single-node system or an HA pair:

Single-node systems

Single-node systems can use these types of Azure Managed Disks:

- *Premium SSD Managed Disks* provide high performance for I/O-intensive workloads at a higher cost.
- *Premium SSD v2 Managed Disks* provide higher performance with lower latency at a lower cost for both single node and HA pairs, compared to Premium SSD Managed Disks.
- *Standard SSD Managed Disks* provide consistent performance for workloads that require low IOPS.
- *Standard HDD Managed Disks* are a good choice if you don't need high IOPS and want to reduce your costs.

Each managed disk type has a maximum disk size of 32 TiB.

You can pair a managed disk with Azure Blob storage to [low-cost object storage](#).

HA pairs

HA pairs use two types of disks which provide high performance for I/O-intensive workloads at a higher cost:

- *Premium page blobs* with a maximum disk size of 8 TiB
- *Managed disks* with a maximum disk size of 32 TiB

Related links

- [Learn how to choose disk types and disk sizes for your systems in Azure](#)
- [Launch a Cloud Volumes ONTAP HA pair in Azure](#)

- [Microsoft Azure documentation: Azure managed disk types](#)
- [Microsoft Azure documentation: Overview of Azure page blobs](#)
- [Review storage limits for Cloud Volumes ONTAP in Azure](#)

Google Cloud storage

In Google Cloud, an aggregate can contain up to 6 disks that are all the same size. The maximum disk size is 64 TiB.

The disk type can be either *Zonal SSD persistent disks*, *Zonal Balanced persistent disks*, or *Zonal standard persistent disks*. You can pair persistent disks with a Google Storage bucket to [low-cost object storage](#).

Related links

- [Google Cloud documentation: Storage Options](#)
- [Review storage limits for Cloud Volumes ONTAP in Google Cloud](#)

RAID type

The RAID type for each Cloud Volumes ONTAP aggregate is RAID0 (striping). Cloud Volumes ONTAP relies on the cloud provider for disk availability and durability. No other RAID types are supported.

Hot spares

RAID0 doesn't support the use of hot spares for redundancy.

Creating unused disks (hot spares) attached to a Cloud Volumes ONTAP instance is an unnecessary expense and may prevent provisioning additional space as needed. Therefore, it's not recommended.

Learn about support for AWS Elastic Volumes with Cloud Volumes ONTAP

Support for the Amazon EBS Elastic Volumes feature with a Cloud Volumes ONTAP aggregate provides better performance and additional capacity, while enabling the NetApp Console to automatically increase the underlying disk capacity as needed.

Benefits

- Dynamic disk growth

The Console can dynamically increase the size of disks while Cloud Volumes ONTAP is running and while disks are still attached.

- Better performance

Aggregates that are enabled with Elastic Volumes can have up to eight disks that are equally utilized across two RAID groups. This configuration provides more throughput and consistent performance.

- Larger aggregates

Support for eight disks provides a maximum aggregate capacity of 128 TiB. These limits are higher than

the six disk limit and 96 TiB limit for aggregates that aren't enabled with the Elastic Volumes feature.

Note that total system capacity limits remain the same.

[AWS Documentation: Learn more about Elastic Volumes from AWS](#)

Supported configurations

The Amazon EBS Elastic Volumes feature is supported with specific Cloud Volumes ONTAP versions and specific EBS disk types.

Cloud Volumes ONTAP version

The Elastic Volumes feature is supported with *new* Cloud Volumes ONTAP systems created from version 9.11.0 or later. The feature is *not* supported with existing Cloud Volumes ONTAP systems that were deployed prior to 9.11.0.

For example, the Elastic Volumes feature is not supported if you created a Cloud Volumes ONTAP 9.9.0 system and then later upgraded that system to version 9.11.0. It must be a new system deployed using version 9.11.0 or later.

EBS disk types

The Elastic Volumes feature is automatically enabled at the aggregate level when using General Purpose SSDs (gp3) or Provisioned IOPS SSDs (io1). The Elastic Volumes feature is not supported with aggregates that use any other disk type.

Required AWS permissions


Starting with the 3.9.19 release, the Console agent requires the following permissions to enable and manage the Elastic Volumes feature on a Cloud Volumes ONTAP aggregate:

- ec2:DescribeVolumesModifications
- ec2:ModifyVolume

These permissions are included in [the policies provided by NetApp](#)

How support for Elastic Volumes works

An aggregate that has the Elastic Volumes feature enabled is comprised of one or two RAID groups. Each RAID group has four identical disks that have the same capacity. Here's an example of a 10 TiB aggregate that has four disks that are 2.5 TiB each:




When the Console creates an aggregate, it starts with one RAID group. If additional capacity is needed, it grows the aggregate by increasing the capacity of all disks in the RAID group by the same amount. The capacity increase is either a minimum of 256 GiB or 10% of the aggregate's size.

For example, if you have a 1 TiB aggregate, each disk is 250 GiB. 10% of the aggregate's capacity is 100 GiB. That's lower than 256 GiB, so the size of the aggregate is increased by the 256 GiB minimum (or 64 GiB for each disk).

The Console increases the size of the disks while the Cloud Volumes ONTAP system is running and while the disks are still attached. The change is non-disruptive.

If an aggregate reaches 64 TiB (or 16 TiB on each disk), the Console creates a second RAID group for additional capacity. This second RAID group works just like the first one: it has four disks that have the exact same capacity and it can grow up to 64 TiB. That means an aggregate can have a maximum capacity of 128 TiB.

Here's an example of an aggregate with two RAID groups. The capacity limit has been reached on the first RAID group, while the disks in the second RAID group have plenty of free space.



What happens when you create a volume

If you create a volume that uses gp3 or io1 disks, the Console creates the volume on an aggregate as follows:

- If there is an existing gp3 or io1 aggregate that has Elastic Volumes enabled, the Console creates the volume on that aggregate.

- If there are multiple gp3 or io1 aggregates that have Elastic Volumes enabled, the Console creates the volume on the aggregate that requires the least amount of resources.
- If the system only has gp3 or io1 aggregates that aren't enabled for Elastic Volumes, then the volume is created on that aggregate.



While this scenario is unlikely, it's possible in two cases:

- You explicitly disabled the Elastic Volumes feature when creating an aggregate from the API.
- You created a new Cloud Volumes ONTAP system from the user interface, in which case the Elastic Volumes feature is disabled on the initial aggregate. Review [Limitations](#) below to learn more.

- If no existing aggregates have enough capacity, the Console creates the aggregate with Elastic Volumes enabled and then creates the volume on that new aggregate.

The size of the aggregate is based on the requested volume size plus an additional 10% capacity.

Capacity Management Mode

The Capacity Management Mode for a Console agent works with Elastic Volumes similar to how it works with other types of aggregates:

- When Automatic mode is enabled (this is the default setting), the Console automatically increases the size of aggregates if additional capacity is needed.
- If you change the capacity management mode to Manual, the Console asks for your approval to purchase additional capacity.

[Learn more about the Capacity Management Mode.](#)

Limitations

Increasing the size of an aggregate can take up to 6 hours. During that time, the Console can't request any additional capacity for that aggregate.

How to work with Elastic Volumes

You can perform these tasks with Elastic Volumes:

- Create a new system that has Elastic Volumes enabled on the initial aggregate when using gp3 or io1 disks

[Learn how to create Cloud Volumes ONTAP system](#)

- Create a new volume on an aggregate that has Elastic Volumes enabled

If you create a volume that uses gp3 or io1 disks, the Console automatically creates the volume on an aggregate that has Elastic Volumes enabled. For more details, refer to [What happens when you create a volume.](#)


[Learn how to create volumes.](#)

- Create a new aggregate that has Elastic Volumes enabled

Elastic Volumes is automatically enabled on new aggregates that use gp3 or io1 disks, as long as the Cloud Volumes ONTAP system was created from version 9.11.0 or later.

When you create the aggregate, the Console prompts you for the aggregate's capacity size. This is different than other configurations where you choose a disk size and number of disks.

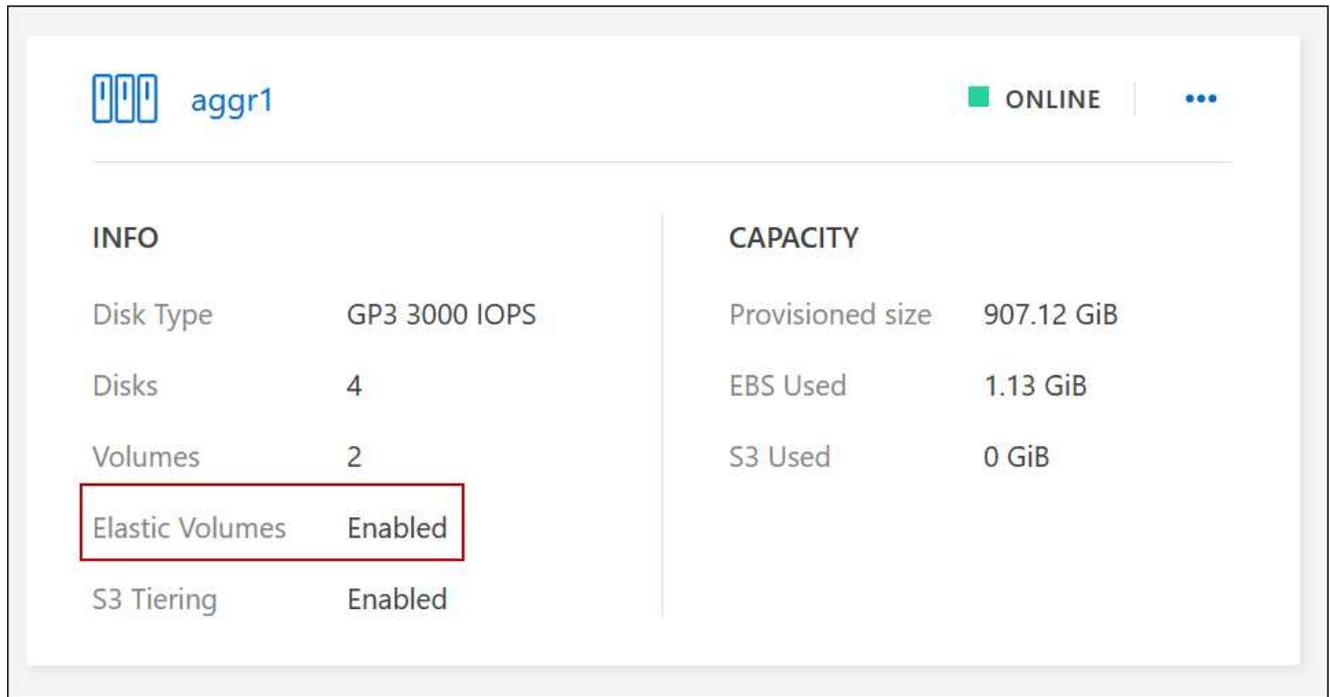
The following screenshot shows an example of a new aggregate comprised of gp3 disks.



[Learn how to create aggregates.](#)

- Identify aggregates that have Elastic Volumes enabled

When you go to the Advanced Allocation page, you can identify whether the Elastic Volumes feature is enabled on an aggregate. In the following example, aggr1 has Elastic Volumes enabled.



- Add capacity to an aggregate

While the Console automatically adds capacity to aggregates as needed, you can manually increase the capacity yourself.

[Learn how to increase aggregate capacity.](#)

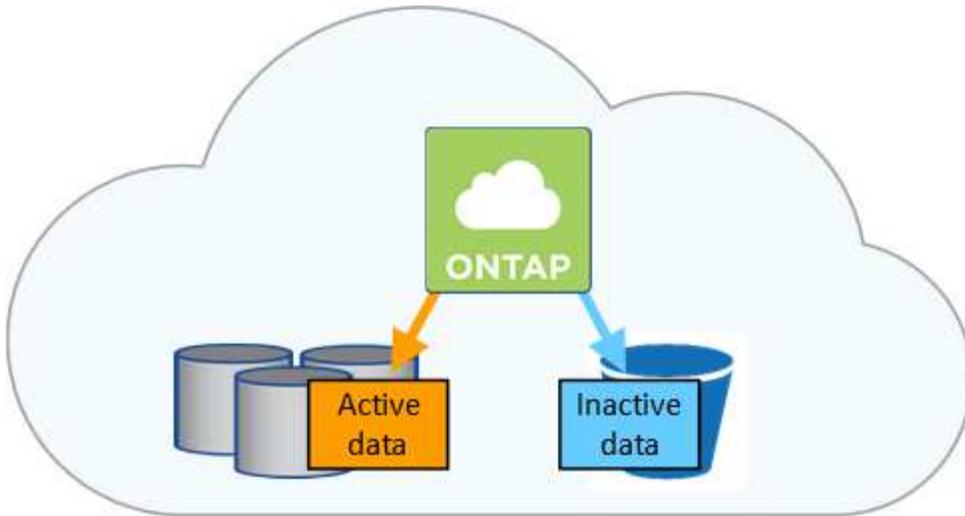
- Replicate data to an aggregate that has Elastic Volumes enabled

If the destination Cloud Volumes ONTAP system supports Elastic Volumes, a destination volume will be placed on an aggregate that has Elastic Volumes enabled (as long as you choose a gp3 or io1 disk).

[Learn how to set up data replication](#)

Learn about data tiering with Cloud Volumes ONTAP in AWS, Azure, or Google Cloud

Reduce your storage costs by enabling automated tiering of inactive data to low-cost object storage. Active data remains in high-performance SSDs or HDDs, while inactive data is tiered to low-cost object storage. This enables you to reclaim space on your primary storage and shrink secondary storage.



Data tiering is powered by FabricPool technology. Cloud Volumes ONTAP provides data tiering for all Cloud Volumes ONTAP clusters without an additional license. When you enable data tiering, data tiered to object storage incurs charges. Refer to your cloud provider’s documentation for details about object storage costs.

Data tiering in AWS

When you enable data tiering in AWS, Cloud Volumes ONTAP uses EBS as a performance tier for hot data and Amazon Simple Storage Service (Amazon S3) as a capacity tier for inactive data.

Performance tier

The performance tier can be General Purpose SSDs (gp3 or gp2) or Provisioned IOPS SSDs (io1).

Tiering data to object storage is not recommended when using Throughput Optimized HDDs (st1).

Capacity tier

A Cloud Volumes ONTAP system tiers inactive data to a single S3 bucket.

The NetApp Console creates a single S3 bucket for each system and names it *fabric-pool-cluster unique identifier*. A different S3 bucket is not created for each volume.

When the Console creates the S3 bucket, it uses the following default settings:

- Storage class: Standard
- Default encryption: Disabled
- Block public access: Block all public access
- Object ownership: ACLs enabled
- Bucket versioning: Disabled
- Object lock: Disabled

Storage classes

The default storage class for tiered data in AWS is *Standard*. Standard is ideal for frequently accessed data stored across multiple Availability Zones.

If you don’t plan to access the inactive data, you can reduce your storage costs by changing the storage class to one of the following: *Intelligent Tiering*, *One-Zone Infrequent Access*, *Standard-Infrequent Access*, or *S3 Glacier Instant Retrieval*. When you change the storage class, inactive data starts in the Standard

storage class and transitions to the storage class that you selected, if the data is not accessed after 30 days.

Access costs are higher if you access the data, so consider this before changing the storage class. [Amazon S3 documentation: Learn more about Amazon S3 storage classes](#).

You can select a storage class when you create the system and you can change it any time afterwards. For instructions on changing the storage class, refer to [Tier inactive data to low-cost object storage](#).

The storage class for data tiering is system wide—it's not per volume.

Data tiering in Azure

When you enable data tiering in Azure, Cloud Volumes ONTAP uses Azure managed disks as a performance tier for hot data and Azure Blob storage as a capacity tier for inactive data.

Performance tier

The performance tier can be either SSDs or HDDs.

Capacity tier

A Cloud Volumes ONTAP system tiers inactive data to a single Blob container.

The Console creates a new storage account with a container for each Cloud Volumes ONTAP system. The name of the storage account is random. A different container is not created for each volume.

The Console creates the storage account with the following settings:

- Access tier: Hot
- Performance: Standard
- Redundancy: Accordingly to Cloud Volume ONTAP Deployment
 - Single availability zone: Locally-redundant storage (LRS)
 - Multiple availability zone: Zone-redundant storage (ZRS)
- Account: StorageV2 (general purpose v2)
- Require secure transfer for REST API operations: Enabled
- Storage account key access: Enabled
- Minimum TLS version: Version 1.2
- Infrastructure encryption: Disabled

Storage access tiers

The default storage access tier for tiered data in Azure is the *hot* tier. The hot tier is ideal for frequently accessed data in the capacity tier.

If you don't plan to access the inactive data in the capacity tier, you can choose the *cool* storage tier, where the inactive data is retained for a minimum of 30 days. You can also opt for the *cold* tier, where the inactive data is stored for a minimum of 90 days. Based on your storage requirements and cost considerations, you can select the tier that best suits your needs. When you change the storage tier to *cool* or *cold*, the inactive capacity tier data moves directly to the cool or cold storage tier. The cool and cold tiers offer lower storage costs compared to the hot tier, but they come with higher access costs, so take that into consideration before you change the storage tier. Refer to [Microsoft Azure documentation: Learn more about Azure Blob storage access tiers](#).

You can select a storage tier when you add a Cloud Volumes ONTAP system and you can change it any time afterwards. For details about changing the storage tier, refer to [Tier inactive data to low-cost object storage](#).

The storage access tier for data tiering is system wide—it's not per volume.

Data tiering in Google Cloud

When you enable data tiering in Google Cloud, Cloud Volumes ONTAP uses persistent disks as a performance tier for hot data and a Google Cloud Storage bucket as a capacity tier for inactive data.

Performance tier

The performance tier can be either SSD persistent disks, balanced persistent disks, or standard persistent disks.

Capacity tier

A Cloud Volumes ONTAP system tiers inactive data to a single Google Cloud Storage bucket.

The Console creates a bucket for each system and names it *fabric-pool-cluster unique identifier*. A different bucket is not created for each volume.

When the Console creates the bucket, it uses the following default settings:

- Location type: Region
- Storage class: Standard
- Public access: Subject to object ACLs
- Access control: Fine-grained
- Protection: None
- Data encryption: Google-managed key

Storage classes

The default storage class for tiered data is the *Standard Storage* class. If the data is infrequently accessed, you can reduce your storage costs by changing to *Nearline Storage* or *Coldline Storage*. When you change the storage class, subsequent inactive data moves directly to the class that you selected.



Any existing inactive data will maintain the default storage class when you change the storage class. To change the storage class for existing inactive data, you must perform the designation manually.

The access costs are higher if you do access the data, so take that into consideration before you change the storage class. To learn more, refer to the [Google Cloud documentation: Storage classes](#).

You can select a storage tier when you create the system and you can change it any time afterwards. For details about changing the storage class, refer to [Tier inactive data to low-cost object storage](#).

The storage class for data tiering is system wide—it's not per volume.

Data tiering and capacity limits

If you enable data tiering, a system's capacity limit stays the same. The limit is spread across the performance tier and the capacity tier.

Volume tiering policies

To enable data tiering, you must select a volume tiering policy when you create, modify, or replicate a volume. You can select a different policy for each volume.

Some tiering policies have an associated minimum cooling period, which sets the time that user data in a volume must remain inactive for the data to be considered "cold" and moved to the capacity tier. The cooling period starts when data is written to the aggregate.



You can change the minimum cooling period and default aggregate threshold of 50% (more on that below). [Learn how to change the cooling period](#) and [learn how to change the threshold](#).

The Console enables you to choose from the following volume tiering policies when you create or modify a volume:

Snapshot Only

After an aggregate has reached 50% capacity, Cloud Volumes ONTAP tiers cold user data of Snapshot copies that are not associated with the active file system to the capacity tier. The cooling period is approximately 2 days.

If read, cold data blocks on the capacity tier become hot and are moved to the performance tier.

All

All data (not including metadata) is immediately marked as cold and tiered to object storage as soon as possible. There is no need to wait 48 hours for new blocks in a volume to become cold. Note that blocks located in the volume prior to the All policy being set require 48 hours to become cold.

If read, cold data blocks on the cloud tier stay cold and are not written back to the performance tier. This policy is available starting with ONTAP 9.6.

Auto

After an aggregate has reached 50% capacity, Cloud Volumes ONTAP tiers cold data blocks in a volume to a capacity tier. The cold data includes not just Snapshot copies but also cold user data from the active file system. The cooling period is approximately 31 days.

This policy is supported starting with Cloud Volumes ONTAP 9.4.

If read by random reads, the cold data blocks in the capacity tier become hot and move to the performance tier. If read by sequential reads, such as those associated with index and antivirus scans, the cold data blocks stay cold and do not move to the performance tier.

None

Keeps data of a volume in the performance tier, preventing it from being moved to the capacity tier.

Replication

When you replicate a volume, you can choose whether to tier the data to object storage. If you do, the Console applies the **Backup** policy to the data protection volume. Starting with Cloud Volumes ONTAP 9.6, the **All** tiering policy replaces the backup policy. When a replication relationship is deleted, the destination volume retains the tiering policy that was in effect during replication.

Turning off Cloud Volumes ONTAP impacts the cooling period

Data blocks are cooled by cooling scans. During this process, blocks that haven't been used have their block

temperature moved (cooled) to the next lower value. The default cooling time depends on the volume tiering policy:

- Auto: 31 days
- Snapshot Only: 2 days

Cloud Volumes ONTAP must be running for the cooling scan to work. If Cloud Volumes ONTAP is turned off, cooling will stop, as well. As a result, you can experience longer cooling times.



When Cloud Volumes ONTAP is turned off, the temperature of each block is preserved until you restart the system. For example, if the temperature of a block is 5 when you turn the system off, the temp is still 5 when you turn the system back on.

Setting up data tiering

For instructions and a list of supported configurations, refer to [Tier inactive data to low-cost object storage](#).

Cloud Volumes ONTAP storage management

The NetApp Console provides simplified and advanced management of Cloud Volumes ONTAP storage.



You must create and delete all disks and aggregates directly from the Console. You should not perform these actions from another management tool. Doing so can impact system stability, hamper the ability to add disks in the future, and potentially generate redundant cloud provider fees.

Storage provisioning

The Console makes storage provisioning for Cloud Volumes ONTAP easy by purchasing disks and managing aggregates for you. You only need to create volumes. You can use an advanced allocation option to provision aggregates yourself, if you want.

Simplified provisioning

Aggregates provide cloud storage to volumes. The Console creates aggregates for you when you launch an instance, and when you provision additional volumes.

When you create a volume, the Console does one of three things:

- It places the volume on an existing aggregate that has sufficient free space.
- It places the volume on an existing aggregate by purchasing more disks for that aggregate.

In the case of an aggregate in AWS that supports Elastic Volumes, it also increases the size of the disks in a RAID group. [Learn more about support for Elastic Volumes](#).

- It purchases disks for a new aggregate and places the volume on that aggregate.

The Console determines where to place a new volume by looking at several factors: an aggregate's maximum size, whether thin provisioning is enabled, and free space thresholds for aggregates.

Disk size selection for aggregates in AWS

When the Console creates new aggregates for Cloud Volumes ONTAP in AWS, it gradually increases disk sizes as aggregate numbers increase to maximize system capacity before reaching AWS data disk limits.

For example, the Console might choose the following disk sizes:

Aggregate number	Disk size	Max aggregate capacity
1	500 GiB	3 TiB
4	1 TiB	6 TiB
6	2 TiB	12 TiB



This behavior does not apply to aggregates that support the Amazon EBS Elastic Volumes feature. Aggregates that have Elastic Volumes enabled are comprised of one or two RAID groups. Each RAID group has four identical disks that have the same capacity. [Learn more about support for Elastic Volumes.](#)

You can choose the disk size yourself by using the advanced allocation option.

Advanced allocation

You can also manage aggregates. [From the Advanced allocation page](#), you can create new aggregates that include a specific number of disks, add disks to an existing aggregate, and create volumes in specific aggregates.

Capacity management

The organization or account admin can configure the Console to notify you of storage capacity decisions or whether to automatically manage capacity requirements for you.

This behavior is determined by the *Capacity Management Mode* on a Console agent. The Capacity Management Mode affects all Cloud Volumes ONTAP systems managed by that Console agent. If you have another Console agent, it can be configured differently.

Automatic capacity management

The Capacity Management Mode is set to automatic by default. In this mode, the Console checks the free space ratio every 15 minutes to determine if the free space ratio falls below the specified threshold. If more capacity is needed, it initiates purchase of new disks, deletes unused collections of disks (aggregates), moves volumes between aggregates as required, and attempts to prevent disk failure.

The following examples illustrate how this mode works:

- If an aggregate reaches the capacity threshold and it has room for more disks, the Console automatically purchases new disks for that aggregate so volumes can continue to grow.

In the case of an aggregate in AWS that supports Elastic Volumes, it also increases the size of the disks in a RAID group. [Learn more about support for Elastic Volumes.](#)

- If an aggregate reaches the capacity threshold and it can't support any additional disks, the Console automatically moves a volume from that aggregate to an aggregate with available capacity or to a new

aggregate.

If the Console creates a new aggregate for the volume, it chooses a disk size that accommodates the size of that volume.

Note that free space is now available on the original aggregate. Existing volumes or new volumes can use that space. The space can't be returned to the cloud provider in this scenario.

- If an aggregate contains no volumes for more than 12 hours, the Console deletes it.

Management of LUNs with automatic capacity management

The Console's automatic capacity management doesn't apply to LUNs. When it creates a LUN, it disables the autogrow feature.

Manual capacity management

If the organization or account admin sets the **Capacity Management Mode** to manual, the Console informs you to take appropriate actions for capacity decisions. The same examples described in the automatic mode apply to the manual mode, but it is up to you to accept the actions.

Learn more

[Learn how to modify the capacity management mode.](#)

Write speed

NetApp Console enables you to choose normal or high write speed for most Cloud Volumes ONTAP configurations. Before you choose a write speed, you should understand the differences between the normal and high settings and risks and recommendations when using high write speed.

Normal write speed

When you choose normal write speed, data is written directly to disk. When data is written directly to disk, reduces the likelihood of data loss in the event of an unplanned system outage, or a cascading failure involving an unplanned system outage (HA pairs only).

Normal write speed is the default option.

High write speed

When you choose high write speed, data is buffered in memory before it is written to disk, which provides faster write performance. Due to this caching, there is the potential for data loss if an unplanned system outage occurs.

The amount of data that can be lost in the event of an unplanned system outage is the span of the last two consistency points. A consistency point is the act of writing buffered data to disk. A consistency point occurs when the write log is full or after 10 seconds (whichever comes first). However, the performance of the storage provided by your cloud provider can affect consistency point processing time.

When to use high write speed

High write speed is a good choice if fast write performance is required for your workload and you can withstand the risk of data loss in the event of an unplanned system outage, or a cascading failure involving an unplanned system outage (HA pairs only).

Recommendations when using high write speed

If you enable high write speed, you should ensure write protection at the application layer, or that the applications can tolerate data loss, if it occurs.

High write speed with an HA pair in AWS

If you plan to enable high write speed on an HA pair in AWS, you should understand the difference in protection levels between a multiple Availability Zone (AZ) deployment and a single AZ deployment. Deploying an HA pair across multiple AZs provides more resiliency and can help to mitigate the chance of data loss.

[Learn more about HA pairs in AWS.](#)

Configurations that support high write speed

Not all Cloud Volumes ONTAP configurations support high write speed. Those configurations use normal write speed by default.

AWS

If you use a single-node system, Cloud Volumes ONTAP supports high write speed with all instance types.

Starting with the 9.8 release, Cloud Volumes ONTAP supports high write speed with HA pairs when using almost all supported EC2 instance types, except for m5.xlarge and r5.xlarge.

[Learn more about the Amazon EC2 instances that Cloud Volumes ONTAP supports.](#)

Azure

If you use a single-node system, Cloud Volumes ONTAP supports high write speed with all VM types.

If you use an HA pair, Cloud Volumes ONTAP supports high write speed with several VM types, starting with the 9.8 release. Go to the [Cloud Volumes ONTAP Release Notes](#) to view the VM types that support high write speed.

Google Cloud

If you use a single-node system, Cloud Volumes ONTAP supports high write speed with all machine types.

If you use an HA pair, Cloud Volumes ONTAP supports high write speed with several VM types, starting with the 9.13.0 release. Go to the [Cloud Volumes ONTAP Release Notes](#) to view the VM types that support high write speed.

[Learn more about the Google Cloud machine types that Cloud Volumes ONTAP supports.](#)

How to select a write speed

You can choose a write speed when you add a new Cloud Volumes ONTAP system and you can [change the write speed for an existing system](#).

What to expect if data loss occurs

If data loss occurs due to high write speed, the Event Management System (EMS) reports the following two events:

- Cloud Volumes ONTAP 9.12.1 or later

```
NOTICE nv.data.loss.possible: An unexpected shutdown occurred while in high write speed mode, which possibly caused a loss of data.
```

- Cloud Volumes ONTAP 9.11.0 to 9.11.1

```
DEBUG nv.check.failed: NVRAM check failed with error "NVRAM disabled due to dirty shutdown with High Write Speed mode"
```

```
ERROR wafl.root.content.changed: Contents of the root volume '' might have changed. Verify that all recent configuration changes are still in effect..
```

- Cloud Volumes ONTAP 9.8 to 9.10.1

```
DEBUG nv.check.failed: NVRAM check failed with error "NVRAM disabled due to dirty shutdown"
```

```
ERROR wafl.root.content.changed: Contents of the root volume '' might have changed. Verify that all recent configuration changes are still in effect.
```

When this happens, Cloud Volumes ONTAP should be able to boot up and continue to serve data without user intervention.

How to stop data access if data loss occurs

If you are concerned about data loss, want the applications to stop running upon data loss, and the data access to be resumed after the data loss issue is properly addressed, you can use the NVFAIL option from the CLI to achieve that goal.

To enable the NVFAIL option

```
vol modify -volume <vol-name> -nvfail on
```

To check NVFAIL settings

```
vol show -volume <vol-name> -fields nvfail
```

To disable the NVFAIL option

```
vol modify -volume <vol-name> -nvfail off
```

When data loss occurs, an NFS or iSCSI volume with NVFAIL enabled should stop serving data (there's no impact to CIFS which is a stateless protocol). For more details, refer to [How NVFAIL impacts access to NFS volumes or LUNs](#).

To check the NVFAIL state

```
vol show -fields in-nvfailed-state
```

After the data loss issue is properly addressed, you can clear the NVFAIL state and the volume will be available for data access.

To clear the NVFAIL state

```
vol modify -volume <vol-name> -in-nvfailed-state false
```

Flash Cache

Some Cloud Volumes ONTAP configurations include local NVMe storage, which Cloud Volumes ONTAP uses as *Flash Cache* for better performance.

What's Flash Cache?

Flash Cache speeds access to data through real-time intelligent caching of recently read user data and NetApp metadata. It's effective for random read-intensive workloads, including databases, email, and file services.

Supported configurations

Flash Cache is supported with specific Cloud Volumes ONTAP configurations. View supported configurations in the [Cloud Volumes ONTAP Release Notes](#)

Limitations

- When configuring Flash Cache for Cloud Volumes ONTAP 9.12.0 or earlier in AWS, compression must be disabled on all volumes to take advantage of the Flash Cache performance improvements. When you deploy or upgrade to Cloud Volumes ONTAP 9.12.1 or later, you don't need to disable compression.

Skip selecting storage efficiency settings when creating a volume from the NetApp Console, or create a volume and then [disable data compression by using the CLI](#).

- Cache rewarming after a reboot is not supported with Cloud Volumes ONTAP.

Related topics

- [Supported configurations for Cloud Volumes ONTAP in AWS](#)
- [Supported configurations for Cloud Volumes ONTAP in Azure](#)
- [Supported configurations for Cloud Volumes ONTAP in Google Cloud](#)

Learn about WORM storage on Cloud Volumes ONTAP

You can activate write once, read many (WORM) storage on a Cloud Volumes ONTAP

system to retain files in unmodified form for a specified retention period. Cloud WORM storage is powered by SnapLock technology, which means WORM files are protected at the file level.

The WORM feature is available for use with bring your own license (BYOL) and marketplace subscriptions for your licenses at no additional cost. Contact your NetApp sales representative to add WORM to your current license.

How WORM storage works

Once a file has been committed to WORM storage, it can't be modified, even after the retention period has expired. A tamper-proof clock determines when the retention period for a WORM file has elapsed.

After the retention period has elapsed, you are responsible for deleting any files that you no longer need.

Activating WORM storage

How you activate WORM storage depends on the Cloud Volumes ONTAP version that you're using.

Version 9.10.1 and later

Beginning with Cloud Volumes ONTAP 9.10.1, you have the option to enable or disable WORM at the volume level.

When you add a Cloud Volumes ONTAP system, you're prompted to enable or disable WORM storage:

- If you enable WORM storage when adding a system, every volume that you create from the NetApp Console has WORM enabled. But you can use ONTAP System Manager or the ONTAP CLI to create volumes that have WORM disabled.
- If you disable WORM storage when adding a system, every volume that you create from the Console, ONTAP System Manager, or the ONTAP CLI has WORM disabled.

Version 9.10.0 and earlier

You can activate WORM storage on a Cloud Volumes ONTAP system when you add a new system. Every volume that you create from the Console has WORM enabled. You can't disable WORM storage on individual volumes.

Committing files to WORM

You can use an application to commit files to WORM over NFS or CIFS, or use the ONTAP CLI to autocommit files to WORM automatically. You can also use a WORM appendable file to retain data that is written incrementally, like log information.

After you activate WORM storage on a Cloud Volumes ONTAP system, you must use the ONTAP CLI for all management of WORM storage. For instructions, refer to the [ONTAP documentation on SnapLock](#).

Enabling WORM on a Cloud Volumes ONTAP system

You can enable WORM storage when creating a Cloud Volumes ONTAP system on the Console. You can also enable WORM on a system if WORM is not enabled on it during creation. After you enable it, you cannot disable WORM.

About this task

- WORM is supported on ONTAP 9.10.1 and later.
- WORM with backup is supported on ONTAP 9.11.1 and later.

Steps

1. On the **Systems** page, double-click the name of the system on which you want to enable WORM.
2. On the Overview tab, click the Features panel and then click the pencil icon next to **WORM**.

If WORM is already enabled on the system, the pencil icon is disabled.

3. On the **WORM** page, set the retention period for the cluster Compliance Clock.

For more information, refer to the [ONTAP documentation: Initialize the Compliance Clock](#).

4. Click **Set**.

After you finish

You can verify the status of **WORM** on the Features panel.

After WORM is enabled, the SnapLock license is automatically installed on the cluster. You can view the SnapLock license on ONTAP System Manager.

Deleting WORM files

You can delete WORM files during the retention period using the privileged delete feature.

For instructions, refer to the [ONTAP documentation](#).

WORM and data tiering

When you create a new Cloud Volumes ONTAP 9.8 system or later, you can enable both data tiering and WORM storage together. Enabling data tiering with WORM storage allows you to tier the data to an object store in the cloud.

You should understand the following about enabling both data tiering and WORM storage:

- Data that is tiered to object storage doesn't include the ONTAP WORM functionality. To ensure end-to-end WORM capability, you'll need to set up the bucket permissions correctly.
- The data that is tiered to object storage doesn't carry the WORM functionality, which means technically anyone with full access to buckets and containers can go and delete the objects tiered by ONTAP.
- Reverting or downgrading to Cloud Volumes ONTAP 9.8 is blocked after enabling WORM and tiering.

Limitations

- WORM storage in Cloud Volumes ONTAP operates under a "trusted storage administrator" model. While WORM files are protected from alteration or modification, volumes can be deleted by a cluster administrator even if those volumes contain unexpired WORM data.
- In addition to the trusted storage administrator model, WORM storage in Cloud Volumes ONTAP also implicitly operates under a "trusted cloud administrator" model. A cloud administrator could delete WORM data before its expiration date by removing or editing cloud storage directly from the cloud provider.

Related link

- Create tamperproof Snapshot copies for WORM storage
- Licensing and charging in Cloud Volumes ONTAP

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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