



Concepts

Cloud Volumes ONTAP

NetApp

February 13, 2026

This PDF was generated from <https://docs.netapp.com/us-en/storage-management-cloud-volumes-ontap/concept-licensing.html> on February 13, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Concepts	1
Licensing	1
Licensing for Cloud Volumes ONTAP	1
Learn more about capacity-based licenses for Cloud Volumes ONTAP	5
Storage	9
Supported client protocols for Cloud Volumes ONTAP	9
Disks and aggregates used for Cloud Volumes ONTAP clusters	10
Learn about support for AWS Elastic Volumes with Cloud Volumes ONTAP	13
Learn about data tiering with Cloud Volumes ONTAP in AWS, Azure, or Google Cloud	19
Cloud Volumes ONTAP storage management	24
Write speed	26
Flash Cache	29
Learn about WORM storage on Cloud Volumes ONTAP	29
High-availability pairs	31
Learn about Cloud Volumes ONTAP HA pairs in AWS	31
Learn about Cloud Volumes ONTAP HA pairs in Azure	38
Learn about Cloud Volumes ONTAP HA pairs in Google Cloud	44
Operations unavailable when a node in Cloud Volumes ONTAP HA pair is offline	48
Learn about Cloud Volumes ONTAP data encryption and ransomware protection	49
Encryption of data at rest	49
ONTAP virus scanning	50
Ransomware protection	50
Learn about performance monitoring for Cloud Volumes ONTAP workloads	51
Performance technical reports	51
CPU performance	51
License management for node-based BYOL	52
BYOL system licenses	52
License management for a new system	52
License expiration	52
License renewal	53
License transfer to a new system	53
Learn how AutoSupport and Digital Advisor are used for Cloud Volumes ONTAP	53
Supported default configurations for Cloud Volumes ONTAP	54
Default setup	54
Internal disks for system data	56

Concepts

Licensing

Licensing for Cloud Volumes ONTAP

Several licensing options are available for Cloud Volumes ONTAP. Each option enables you to choose a consumption model that meets your needs.

Licensing overview

The following licensing options are available for new customers.

Capacity-based licensing

Pay for multiple Cloud Volumes ONTAP systems in your NetApp account by provisioned capacity. Includes the ability to purchase add-on cloud data services. For more information about consumption models or purchase options in capacity-based licenses, refer to [Learn more about capacity-based licenses](#).

Keystone Subscription

A pay-as-you-grow subscription-based service that delivers a seamless hybrid cloud experience for High Availability (HA) pairs.

The following sections provide more details about each of these options.



Support is not available for the use of licensed features without a license.

Capacity-based licensing

Capacity-based licensing packages enable you to pay for Cloud Volumes ONTAP per TiB of capacity. The license is associated with your NetApp account and enables you to charge multiple systems against the license, as long as enough capacity is available through the license.

For example, you could purchase a single 20 TiB license, deploy four Cloud Volumes ONTAP systems, and then allocate a 5 TiB volume to each system, for a total of 20 TiB. The capacity is available to the volumes on each Cloud Volumes ONTAP system deployed in that account.

Capacity-based licensing is available in the form of a *package*. When you deploy a Cloud Volumes ONTAP system, you can choose from several licensing packages based on your business needs.



While the actual usage and metering for the products and services managed in the NetApp Console are always calculated in GiB and TiB, the terms GB/GiB and TB/TiB are used interchangeably. This is reflected in the Cloud marketplace listings, price quotes, listing descriptions, and in other supporting documentation.

Packages

The following capacity-based packages are available for Cloud Volumes ONTAP. For more information about capacity-based license packages, refer to [Learn more about capacity-based licenses](#).

For a list of supported VM types with the following capacity-based packages, refer to:

- [Supported configurations in Azure](#)
- [Supported configurations in Google Cloud](#)

Freemium

Provides all Cloud Volumes ONTAP features free of charge from NetApp (cloud provider charges still apply). A Freemium package has these characteristics:

- No license or contract is needed.
- Support from NetApp is not included.
- You're limited to 500 GiB of provisioned capacity per Cloud Volumes ONTAP system.
- You can use up to 10 Cloud Volumes ONTAP systems with the Freemium offering per NetApp account, for any cloud provider.
- If the provisioned capacity for a Cloud Volumes ONTAP system exceeds 500 GiB, the Console converts the system to an Essentials package.

As soon as a system is converted to the Essentials package, [minimum charging](#) applies to it.

A Cloud Volumes ONTAP system that has been converted into an Essentials package cannot be switched back to Freemium even if the provisioned capacity is reduced to less than 500 GiB. Other systems with less than 500 GiB of provisioned capacity stay on Freemium (as long as they were deployed using the Freemium offering).

Essentials

You can pay by capacity in a number of different configurations:

- Choose your Cloud Volumes ONTAP configuration:
 - A single node or HA system
 - File and block storage or secondary data for disaster recovery (DR)
- Add on any of NetApp's cloud data services at extra cost

Professional

Pay by capacity for any type of Cloud Volumes ONTAP configuration with unlimited backups.

- Provides licensing for any Cloud Volumes ONTAP configuration

Single node or HA with capacity charging for primary and secondary volumes at the same rate

- Includes unlimited volume backups using NetApp Backup and Recovery, but only for Cloud Volumes ONTAP systems that use the Professional package.



A pay-as-you-go (PAYGO) subscription is required for Backup and Recovery, however no charges will be incurred for using this service. For more information on setting up licensing for Backup and Recovery, refer to [Set up licensing for Backup and Recovery](#).

- Add on any of NetApp's cloud data services at extra cost

Availability of capacity-based licenses

The availability of the PAYGO and BYOL licenses for Cloud Volumes ONTAP systems requires the Console agent to be up and running.

[Learn about Console agents.](#)



NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAPP](#).

How to get started

Learn how to get started with capacity-based licensing:

- [Set up licensing for Cloud Volumes ONTAP in AWS](#)
- [Set up licensing for Cloud Volumes ONTAP in Azure](#)
- [Set up licensing for Cloud Volumes ONTAP in Google Cloud](#)

Keystone Subscription

A pay-as-you-grow subscription-based service that delivers a seamless hybrid cloud experience for those preferring OpEx consumption models to upfront CapEx or leasing.

Charging is based on the size of your committed capacity for one or more Cloud Volumes ONTAP HA pairs in your Keystone Subscription.

The provisioned capacity for each volume is aggregated and compared to the committed capacity on your Keystone Subscription periodically, and any overages are charged as burst on your Keystone Subscription.

[Learn more about NetApp Keystone.](#)

Supported configurations

Keystone Subscriptions are supported with HA pairs. This licensing option isn't supported with single-node systems at this time.

Capacity limit

In the capacity-based licensing model, each Cloud Volumes ONTAP system supports tiering to object storage, and the total tiered capacity can scale up to the cloud provider's bucket limit. Although the license does not impose capacity restrictions, follow the [FabricPool Best Practices](#) to ensure optimal performance, reliability, and cost efficiency when configuring and managing tiering.

For information about the capacity limits of each cloud provider, refer to their documentation:

- [AWS documentation](#)
- [Azure documentation for managed disks](#) and [Azure documentation for blob storage](#)
- [Google Cloud documentation](#)

How to get started

Learn how to get started with a Keystone Subscription:

- Set up licensing for Cloud Volumes ONTAP in AWS
- Set up licensing for Cloud Volumes ONTAP in Azure
- Set up licensing for Cloud Volumes ONTAP in Google Cloud

Node-based licensing

Node-based licensing is the previous generation licensing model that enabled you to license Cloud Volumes ONTAP by node. This licensing model is not available for new customers. By-node charging has been replaced with the by-capacity charging methods described above.

NetApp has planned the end of availability (EOA) and support (EOS) of node-based licensing. After the EOA and EOS, node-based licenses will need to be converted to capacity-based licenses.

For information, refer to [Customer communique: CPC-00589](#).

End of availability of node-based licenses

Beginning on 11 November, 2024, the limited availability of node-based licenses has been terminated. The support for node-based licensing ends on 31 December, 2024.

If you have a valid node-based contract that extends beyond the EOA date, you can continue to use the license until the contract expires. Once the contract expires, it will be necessary to transition to the capacity-based licensing model. If you don't have a long-term contract for a Cloud Volumes ONTAP node, it is important to plan your conversion before the EOS date.

Learn more about each license type and the impact of EOA on it from this table:

License type	Impact after EOA
Valid node-based license purchased through bring your own license (BYOL)	License remains valid till expiration. Existing unused node-based licenses can be used for deploying new Cloud Volumes ONTAP systems.
Expired node-based license purchased through BYOL	You won't be entitled to deploy new Cloud Volumes ONTAP systems using this license. The existing systems might continue to work, but you won't receive any support or updates for your systems post the EOS date.
Valid node-based license with PAYGO subscription	Will cease to receive NetApp support post the EOS date, until you transition to a capacity-based license.

Exclusions

NetApp recognizes that certain situations require special consideration, and EOA and EOS of node-based licensing will not apply to the following cases:

- U.S. Public Sector customers
- Deployments in private mode
- China region deployments of Cloud Volumes ONTAP in AWS

For these particular scenarios, NetApp will offer support to address the unique licensing requirements in

compliance with contractual obligations and operational needs.



Even in these scenarios, new node-based licenses and license renewals are valid for a maximum of one year from the date of approval.

License conversion

The Console enables a seamless conversion of node-based licenses to capacity based through the license conversion tool. For information about EOA of node-based licensing, refer to [End of availability of node-based licenses](#).

Before transitioning, it is good to familiarize yourself with the difference between the two licensing models. Node-based licensing includes fixed capacity for each ONTAP instance, which can restrict flexibility. Capacity-based licensing, on the other hand, allows for a shared pool of storage across multiple instances, offering enhanced flexibility, optimizing resource utilization, and reducing the potential for financial penalties when redistributing workloads. Capacity-based charging seamlessly adjusts to changing storage requirements.

To know how you can perform this conversion, refer to [Convert a Cloud Volumes ONTAP node-based license to capacity-based license](#).



Conversion of a system from capacity-based to node-based licensing is not supported.

Learn more about capacity-based licenses for Cloud Volumes ONTAP

You should be familiar with the charging and capacity usage for capacity-based licenses.

Consumption models or license purchase options

Capacity-based licensing packages are available with the following consumption models or purchase options:

- **BYOL:** Bring your own license (BYOL). A license purchased from NetApp that can be used to deploy Cloud Volumes ONTAP in any cloud provider.



NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAP](#).

- **PAYGO:** A pay-as-you-go (PAYGO) subscription is an hourly subscription from your cloud provider's marketplace.
- **Annual:** An annual contract from your cloud provider's marketplace.

Note the following:

- If you purchase a license from NetApp (BYOL), you also need to subscribe to the PAYGO offering from your cloud provider's marketplace. NetApp has restricted BYOL licensing. When your BYOL licenses expire, you are required to replace them with cloud marketplace subscriptions.

Your license is always charged first, but you'll be charged from the hourly rate in the marketplace in these cases:

- If you exceed your licensed capacity
- If the term of your license expires

- If you have an annual contract from a marketplace, *all* Cloud Volumes ONTAP systems that you deploy are charged against that contract. You can't mix and match an annual marketplace contract with BYOL.
- Only single-node systems with BYOL are supported in China regions. China region deployments are exempt from BYOL licensing restrictions.

Changing of license packages

After deployment, you can change the package for a Cloud Volumes ONTAP system that uses capacity-based licensing. For example, if you deployed a Cloud Volumes ONTAP system with the Essentials package, you can change it to the Professional package if your business needs changed.

[Learn how to change charging methods.](#)

For information about converting node-based licenses to capacity-based, see

How you are charged for supported storage types and packages

Charging in Cloud Volumes ONTAP is based on a number of factors, such as packages and volume types. Capacity-based licensing packages are available with Cloud Volumes ONTAP 9.7 and later.

For details about pricing, go to the [NetApp Console website](#).

Storage VMs

- There are no extra licensing costs for additional data-serving storage VMs (SVMs), but there is a 4 TiB minimum capacity charge per data-serving SVM.
- Disaster recovery SVMs are charged according to the provisioned capacity.

HA pairs

For HA pairs, you're only charged for the provisioned capacity on a node. You aren't charged for data that is synchronously mirrored to the partner node.

FlexClone and FlexCache volumes

- You won't be charged for the capacity used by FlexClone volumes.
- Source and destination FlexCache volumes are considered primary data and charged according to the provisioned space.

Read/write volumes

If you create or use a writable (read/write) volume, it is considered a primary volume and is charged for the provisioned capacity based on the minimum charge per storage VM (SVM). Examples include FlexVol read/write volumes, SnapLock audit volumes, and CIFS/NFS audit volumes. All user-created data volumes are charged per your subscription and package type. ONTAP internal volumes that are automatically created and cannot store data, such as SVM root volumes, are not charged.

Essentials packages

With the Essentials package, you're billed by the deployment type (HA or single node) and the volume type (primary or secondary). Pricing from high to low is in the following order: *Essentials Primary HA*, *Essentials Primary Single Node*, *Essentials Secondary HA*, and *Essentials Secondary Single Node*. Alternately, when you purchase a marketplace contract or accept a private offer, capacity charges are the same for any deployment or volume type.

Licensing is based entirely on the volume type created within Cloud Volumes ONTAP systems:

- Essentials Single Node: Read/write volumes created on a Cloud Volumes ONTAP system using one ONTAP node only.
- Essentials HA: Read/write volumes using two ONTAP nodes that can fail over to each other for non-disruptive data access.
- Essentials Secondary Single Node: Data Protection (DP) type volumes (typically SnapMirror or SnapVault destination volumes that are read-only) created on a Cloud Volumes ONTAP system using one ONTAP node only.



If a read-only/DP volume becomes a primary volume, the Console considers it as primary data and the charging costs are calculated based on the time the volume was in read/write mode. When the volume is again made read-only/DP, it considers the volume as secondary data again and charges accordingly using the best matching license in the Console.

- Essentials Secondary HA: Data Protection (DP) type volumes (typically SnapMirror or SnapVault destination volumes that are read-only) created on a Cloud Volumes ONTAP system using two ONTAP nodes that can fail over to each other for non-disruptive data access.

Capacity limit

In the capacity-based licensing model, each Cloud Volumes ONTAP system supports tiering to object storage, and the total tiered capacity can scale up to the cloud provider's bucket limit. Although the license does not impose capacity restrictions, follow the [FabricPool Best Practices](#) to ensure optimal performance, reliability, and cost efficiency when configuring and managing tiering.

For information about the capacity limits of each cloud provider, refer to their documentation:

- [AWS documentation](#)
- [Azure documentation for managed disks](#) and [Azure documentation for blob storage](#)
- [Google Cloud documentation](#)

Max number of systems

With capacity-based licensing, the maximum number of Cloud Volumes ONTAP systems is limited to 24 per NetApp Console organization. A *system* is a Cloud Volumes ONTAP HA pair, a Cloud Volumes ONTAP single-node system, or any additional storage VMs that you create. The default storage VM does not count against the limit. This limit applies to all licensing models.

For example, let's say you have three systems:

- A single node Cloud Volumes ONTAP system with one storage VM (this is the default storage VM that's created when you deploy Cloud Volumes ONTAP)

This system counts as one system.

- A single node Cloud Volumes ONTAP system with two storage VMs (the default storage VM, plus one additional storage VM that you created)

This system counts as two systems: one for the single-node system and one for the additional storage VM.

- A Cloud Volumes ONTAP HA pair with three storage VMs (the default storage VM, plus two additional storage VMs that you created)

This system counts as three systems: one for the HA pair and two for the additional storage VMs.

That's six systems in total. You would then have room for an additional 14 systems in your organization.

If you have a large deployment that requires more than 24 systems, contact your accounts representative or sales team.

[Learn about storage limits for AWS, Azure, and Google Cloud.](#)

Minimum charge

There is a 4 TiB minimum charge for each data-serving storage VM that has at least one primary (read-write) volume. If the sum of the primary volumes is less than 4 TiB, then the Console applies the 4 TiB minimum charge to that storage VM.

If you haven't provisioned any volumes yet, then the minimum charge doesn't apply.

For the Essentials package, the 4 TiB minimum capacity charge doesn't apply to storage VMs that contain secondary (data protection) volumes only. For example, if you have a storage VM with 1 TiB of secondary data, then you're charged just for that 1 TiB of data. With the Professional package type, the minimum capacity charging of 4 TiB applies regardless of the volume type.

Billing preferences and overages

You can choose how you want to be charged in the **Licenses and subscriptions** section in the Console. Overages occur when your usage exceeds the capacity specified in your license package or annual subscription.

- **NetApp licenses first:** In this model, your usage is first charged against the capacity of your license package (BYOL). If you exceed your license capacity, overages are charged based on your annual marketplace subscription or marketplace on-demand hourly rates (PAYGO). If your BYOL license expires, you must transition to a capacity-based licensing model through the cloud marketplaces. For information, refer to [Convert a Cloud Volumes ONTAP node-based license to a capacity-based license](#).
- **Marketplace subscriptions only:** In this model, your usage is first charged against your annual marketplace subscription. Any additional usage is charged at marketplace on-demand hourly rates (PAYGO). Any unused license capacity is disregarded for billing.

For more information about billing preferences, refer to [Learn about billing preferences for licenses and subscriptions](#).

How overages are charged for Essentials licenses

If you purchase an Essentials license from NetApp (BYOL) and you exceed the licensed capacity for a specific Essentials package, the Console charges overages against a higher-priced Essentials license (if you have one with available capacity). The Console first uses the available capacity that you've paid for before charging against the marketplace. If there is no available capacity with your BYOL license, the exceeded capacity is charged at marketplace on-demand hourly rates (PAYGO) and added to your monthly bill.

Similarly, if you have an annual marketplace contract or a private offer with multiple Essentials packages, and your usage exceeds the committed capacity for a deployment and volume type of a specific package, the Console charges overages against a higher-priced Essentials package based on its available capacity. After that capacity is exhausted, the remaining overage is billed at marketplace on-demand (PAYGO) hourly rates and added to your monthly bill.

For information about Essentials licenses charging, refer to [Essentials packages](#).

Here's an example. Let's say you have the following licenses for the Essentials package:

- A 500 TiB *Essentials Secondary HA* license that has 500 TiB of committed capacity
- A 500 TiB *Essentials Single Node* license that only has 100 TiB of committed capacity

Another 50 TiB is provisioned on an HA pair with secondary volumes. Instead of charging that 50 TiB to PAYGO, the Console charges the 50 TiB overage against the *Essentials Single Node* license. That license is priced higher than *Essentials Secondary HA*, but it's making use of a license you have already purchased, and it will not add costs to your monthly bill.

In **Administration > Licenses and subscriptions**, you can see 50 TiB charged against the *Essentials Single Node* license.

Here's another example. Let's say you have the following licenses for the Essentials package:

- A 500 TiB *Essentials Secondary HA* license that has 500 TiB of committed capacity
- A 500 TiB *Essentials Single Node* license that only has 100 TiB of committed capacity

Another 100 TiB is provisioned on an HA pair with primary volumes. The license you purchased doesn't have *Essentials Primary HA* committed capacity. The *Essentials Primary HA* license is priced higher than both the *Essentials Primary Single Node* and *Essentials Secondary HA* licenses.

In this example, the Console charges overages at the marketplace rate for the additional 100 TiB. The overage charges will appear on your monthly bill.

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.



The screenshot shows the 'Select Disk Type' step of a four-step wizard. The steps are numbered 1 through 4: 1 Disk Type (highlighted in blue), 2 Aggregate details, 3 Tiering Data, and 4 Review. The 'Disk Type' dropdown is set to 'GP3 - General Purpose SSD Dynamic Performance'. Below it, a detailed view of the 'General Purpose SSD (gp3) Disk Properties' is shown, including a description: 'General purpose SSD volume that balances price and performance (performance level is independent of storage capacity)', an IOPS Value of 12000, and a Throughput MB/s of 250.

[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.

INFO		CAPACITY	
Disk Type	GP3 3000 IOPS	Provisioned size	907.12 GiB
Disks	4	EBS Used	1.13 GiB
Volumes	2	S3 Used	0 GiB
Elastic Volumes	Enabled		
S3 Tiering	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.

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](#)

High-availability pairs

Learn about Cloud Volumes ONTAP HA pairs in AWS

A Cloud Volumes ONTAP high-availability (HA) configuration provides nondisruptive operations and fault tolerance. In AWS, data is synchronously mirrored between the two nodes.

HA components

In AWS, Cloud Volumes ONTAP HA configurations include the following components:

- Two Cloud Volumes ONTAP nodes whose data is synchronously mirrored between each other.
- A mediator instance that provides a communication channel between the nodes to assist in storage takeover and giveback processes.

Mediator

Here are some key details about the mediator instance in AWS:

Instance type

t3-micro

Disks

Two st1 disks of 8 GiB and 4 GiB

Operating system

Debian 11



For Cloud Volumes ONTAP 9.10.0 and earlier, Debian 10 was installed on the mediator.

Upgrades

When you upgrade Cloud Volumes ONTAP, the NetApp Console also updates the mediator instance as needed.

Access to the instance

When you create a Cloud Volumes ONTAP HA pair from the Console, you're prompted to provide a key pair for the mediator instance. You can use that key pair for SSH access using the `admin` user.

Third-party agents

Third-party agents or VM extensions are not supported on the mediator instance.

Storage takeover and giveback

If a node goes down, the other node can serve data for its partner to provide continued data service. Clients can access the same data from the partner node because the data was synchronously mirrored to the partner.

After the node reboots, the partner must resync data before it can return the storage. The time that it takes to resync data depends on how much data was changed while the node was down.

Storage takeover, resync, and giveback are all automatic by default. No user action is required.

RPO and RTO

An HA configuration maintains high-availability of your data as follows:

- The recovery point objective (RPO) is 0 seconds.
Your data is transactionally consistent with no data loss.
- The recovery time objective (RTO) is 120 seconds.
In the event of an outage, data should be available in 120 seconds or less.

HA deployment models

You can ensure the high availability of your data by deploying an HA configuration across multiple availability zones (AZs) or in a single availability zone (AZ). You should review more details about each configuration to choose which best fits your needs.

Multiple availability zones

Deploying an HA configuration in multiple availability zones (AZs) ensures high availability of your data if a failure occurs with an AZ or an instance that runs a Cloud Volumes ONTAP node. You should understand how NAS IP addresses impact data access and storage failover.

NFS and CIFS data access

When an HA configuration is spread across multiple Availability Zones, *floating IP addresses* enable NAS client access. The floating IP addresses, which must be outside of the CIDR blocks for all VPCs in the region, can migrate between nodes when failures occur. They aren't natively accessible to clients that are outside of the VPC, unless you [set up an AWS transit gateway](#).

If you can't set up a transit gateway, private IP addresses are available for NAS clients that are outside the VPC. However, these IP addresses are static—they can't failover between nodes.

You should review requirements for floating IP addresses and route tables before you deploy an HA configuration across multiple availability zones. You must specify the floating IP addresses when you deploy the configuration. The private IP addresses are automatically created.

For more information, refer to [AWS networking requirements for Cloud Volumes ONTAP HA in multiple AZs](#).

iSCSI data access

Cross-VPC data communication is not an issue since iSCSI does not use floating IP addresses.

Takeover and giveback for iSCSI


For iSCSI, Cloud Volumes ONTAP uses multipath I/O (MPIO) and Asymmetric Logical Unit Access (ALUA) to manage path failover between the active-optimized and non-optimized paths.



For information about which specific host configurations support ALUA, refer to the [NetApp Interoperability Matrix Tool](#) and the [SAN hosts and cloud clients guide](#) for your host operating system.

Takeover and giveback for NAS

When takeover occurs in a NAS configuration using floating IPs, the node's floating IP address that clients use to access data moves to the other node. The following image depicts storage takeover in a NAS configuration using floating IPs. If node 2 goes down, the floating IP address for node 2 moves to node 1.



NAS data IPs used for external VPC access cannot migrate between nodes if failures occur. If a node goes offline, you must manually remount volumes to clients outside the VPC by using the IP address on the other node.

After the failed node comes back online, remount clients to volumes using the original IP address. This step is needed to avoid transferring unnecessary data between two HA nodes, which can cause significant performance and stability impact.

You can locate the correct IP address from the Console by selecting the volume and clicking **Mount Command**.

Single availability zone

Deploying an HA configuration in a single availability zone (AZ) can ensure high availability of your data if an instance that runs a Cloud Volumes ONTAP node fails. All data is natively accessible from outside of the VPC.




The Console creates an [AWS Documentation: AWS spread placement group](#) and launches the two HA nodes in that placement group. The placement group reduces the risk of simultaneous failures by spreading the instances across distinct underlying hardware. This feature improves redundancy from a compute perspective and not from disk failure perspective.

Data access

Because this configuration is in a single AZ, it does not require floating IP addresses. You can use the same IP address for data access from within the VPC and from outside the VPC.

The following image shows an HA configuration in a single AZ. Data is accessible from within the VPC and from outside the VPC.



Takeover and giveback

For iSCSI, Cloud Volumes ONTAP uses multipath I/O (MPIO) and Asymmetric Logical Unit Access (ALUA) to manage path failover between the active-optimized and non-optimized paths.



For information about which specific host configurations support ALUA, refer to the [NetApp Interoperability Matrix Tool](#) and the [SAN hosts and cloud clients guide](#) for your host operating system.

For NAS configurations, the data IP addresses can migrate between HA nodes if failures occur. This ensures client access to storage.

AWS Local Zones

AWS Local Zones are an infrastructure deployment where storage, compute, database, and other select AWS services are located close to large cities and industry areas. With AWS Local Zones, you can bring AWS services closer to you which improves latency for your workloads and maintain databases locally. On Cloud Volumes ONTAP,

You can deploy a single AZ or multiple AZ configuration in AWS Local Zones.



AWS Local Zones are supported when using the Console in standard and private modes. At this time, AWS Local Zones are not supported in restricted mode.

Example AWS Local Zone configurations

Cloud Volumes ONTAP in AWS supports only high availability (HA) mode in a single availability zone. Single node deployments are not supported.

Cloud Volumes ONTAP does not support data tiering, cloud tiering, and unqualified instances in AWS Local Zones.

The following are example configurations:

- Single availability zone: Both cluster nodes and the mediator are in the same Local Zone.
- Multiple availability zones
In multiple availability zone configurations, there are three instances, two nodes and one mediator. One instance out of the three instances must be in a separate zone. You can choose how you set this up.

Here are three example configurations:

- Each cluster node is in a different Local Zone and the mediator in a public availability zone.
- One cluster node in a Local Zone, the mediator in a Local Zone, and the second cluster node is in an availability zone.
- Each cluster node and the mediator are in separate Local Zones.

Supported disk and instance types

The only supported disk type is GP2. The following EC2 instance type families with sizes xlarge to 4xlarge are currently supported:

- M5
- C5
- C5d
- R5
- R5d



Cloud Volumes ONTAP supports only these configurations. Selecting unsupported disk types or unqualified instances in AWS Local Zone configuration might result in deployment failure. Data tiering to Amazon Simple Storage Service (Amazon S3) is not supported if your Cloud Volumes ONTAP system is in an AWS Local Zone, because accessing the Amazon S3 buckets outside of the Local Zone involves higher latency and impacts Cloud Volumes ONTAP activities.

[AWS Documentation: EC2 instance types in Local Zones.](#)

How storage works in an HA pair

Unlike an ONTAP cluster, storage in a Cloud Volumes ONTAP HA pair is not shared between nodes. Instead, data is synchronously mirrored between the nodes so that the data is available in the event of failure.

Storage allocation

When you create a new volume and additional disks are required, the Console allocates the same number of disks to both nodes, creates a mirrored aggregate, and then creates the new volume. For example, if two disks are required for the volume, the Console allocates two disks per node for a total of four disks.

Storage configurations

You can use an HA pair as an active-active configuration, in which both nodes serve data to clients, or as an active-passive configuration, in which the passive node responds to data requests only if it has taken over storage for the active node.



You can set up an active-active configuration only when using the Console in the Storage System View.

Performance expectations

A Cloud Volumes ONTAP HA configuration synchronously replicates data between nodes, which consumes network bandwidth. As a result, you can expect the following performance in comparison to a single-node Cloud Volumes ONTAP configuration:

- For HA configurations that serve data from only one node, read performance is comparable to the read performance of a single-node configuration, whereas write performance is lower.
- For HA configurations that serve data from both nodes, read performance is higher than the read performance of a single-node configuration, and write performance is the same or higher.

For more details about Cloud Volumes ONTAP performance, refer to [Performance](#).


Client access to storage

Clients should access NFS and CIFS volumes by using the data IP address of the node on which the volume resides. If NAS clients access a volume by using the IP address of the partner node, traffic goes between both nodes, which reduces performance.



If you move a volume between nodes in an HA pair, you should remount the volume by using the IP address of the other node. Otherwise, you can experience reduced performance. If clients support NFSv4 referrals or folder redirection for CIFS, you can enable those features on the Cloud Volumes ONTAP systems to avoid remounting the volume. For details, refer to the ONTAP documentation.

You can easily identify the correct IP address through the *Mount Command* option under the manage volumes panel in.




Learn about Cloud Volumes ONTAP HA pairs in Azure

A Cloud Volumes ONTAP high-availability (HA) pair provides enterprise reliability and continuous operations in case of failures in your cloud environment. In Azure, storage is shared between the two nodes.

HA components

HA single availability zone configuration with page blobs

A Cloud Volumes ONTAP HA page blob configuration in Azure includes the following components:



Resource group

Note the following about the Azure components that the NetApp Console deploys for you:

Azure Standard Load Balancer

The load balancer manages incoming traffic to the Cloud Volumes ONTAP HA pair.

VMs in single availability zones

Beginning with Cloud Volumes ONTAP 9.15.1, you can create and manage heterogeneous virtual machines (VMs) in a single availability zone (AZ). You can deploy high-availability (HA) nodes in separate fault domains within the same AZ, guaranteeing optimal availability. To learn more about the flexible orchestration mode that enables this capability, refer to the [Microsoft Azure documentation: Virtual Machine Scale Sets](#).

Disks

Customer data resides on Premium Storage page blobs. Each node has access to the other node's storage. Additional storage is also required for [boot, root, and core data](#).

Storage accounts

- One storage account is required for managed disks.
- One or more storage accounts are required for the Premium Storage page blobs, as the disk capacity limit per storage account is reached.

[Microsoft Azure documentation: Azure Storage scalability and performance targets for storage accounts.](#)


- One storage account is required for data tiering to Azure Blob storage.
- Starting with Cloud Volumes ONTAP 9.7, the storage accounts that the Console creates for HA pairs are general-purpose v2 storage accounts.
- You can enable an HTTPS connection from a Cloud Volumes ONTAP 9.7 HA pair to Azure storage accounts when adding a Cloud Volumes ONTAP system. Note that enabling this option can impact write performance. You can't change the setting after you create the system.

 Starting with Cloud Volumes ONTAP 9.15.0P1, Azure page blobs are no longer supported for new high-availability pair deployments. If you currently use Azure page blobs in existing high-availability pair deployments, you can migrate to newer VM instance types in the Edsv4-series VMs and Edsv5-series VMs.

[Learn more about supported configurations in Azure.](#)

HA single availability zone configuration with shared managed disks

A Cloud Volumes ONTAP HA single availability zone configuration running on top of shared managed disk includes the following components:



Resource group

Note the following about the Azure components that the Console deploys for you:

Azure Standard Load Balancer

The load balancer manages incoming traffic to the Cloud Volumes ONTAP HA pair.

VMs in single availability zones

Beginning with Cloud Volumes ONTAP 9.15.1, you can create and manage heterogeneous virtual machines (VMs) in a single availability zone (AZ). You can deploy high-availability (HA) nodes in separate fault domains within the same AZ, guaranteeing optimal availability. To learn more about the flexible orchestration mode that enables this capability, refer to the [Microsoft Azure documentation: Virtual Machine](#)

Scale Sets.

The zonal deployment uses Premium SSD v2 Managed Disks when the following conditions are fulfilled:

- The version of Cloud Volumes ONTAP is 9.15.1 or later.
- The selected region and zone support Premium SSD v2 Managed Disks. For information about the supported regions, refer to [Microsoft Azure website: Products available by region](#).
- The subscription is registered for the Microsoft [Microsoft.Compute/VMOrchestratorZonalMultiFD feature](#).



If you choose Premium SSD Managed Disks for an environment that fulfills the above criteria, the Console automatically deploys Premium SSD v2 Managed Disks. You cannot switch to Premium SSD v1 Managed Disks.

Disks


Customer data resides on locally-redundant storage (LRS) managed disks. Each node has access to the other node's storage. Additional storage is also required for [boot](#), [root](#), [partner root](#), [core](#), and [NVRAM](#) data.

Storage accounts

Storage accounts are used for managed disk based deployments to handle diagnostic logs and tiering to blob storage.

HA multiple availability zone configuration

A Cloud Volumes ONTAP HA multiple availability zone configuration in Azure includes the following components:



Resource group

Note the following about the Azure components that the Console deploys for you:

Azure Standard Load Balancer

The load balancer manages incoming traffic to the Cloud Volumes ONTAP HA pair.

Availability Zones

HA multiple availability zone configuration utilizes a deployment model where two Cloud Volumes ONTAP nodes are deployed into different availability zones, ensuring that the nodes are in different fault domains to provide redundancy and availability. To learn how Virtual Machine Scale Sets in Flexible orchestration mode can use availability zones in Azure, refer to the [Microsoft Azure documentation: Create a Virtual Machine Scale Set that uses Availability Zones](#).

Disks

Customer data resides on zone-redundant storage (ZRS) managed disks. Each node has access to the other node's storage. Additional storage is also required for [boot](#), [root](#), [partner root](#), and [core data](#).

Storage accounts

Storage accounts are used for managed disk based deployments to handle diagnostic logs and tiering to blob storage.

RPO and RTO

An HA configuration maintains high availability of your data as follows:

- The recovery point objective (RPO) is 0 seconds.
Your data is transactionally consistent with no data loss.
- The recovery time objective (RTO) is 120 seconds.
In the event of an outage, data should be available in 120 seconds or less.

Storage takeover and giveback

Similar to a physical ONTAP cluster, storage in an Azure HA pair is shared between nodes. Connections to the partner's storage allows each node to access the other's storage in the event of a *takeover*. Network path failover mechanisms ensure that clients and hosts continue to communicate with the surviving node. The partner *gives back* storage when the node is brought back on line.

For NAS configurations, data IP addresses automatically migrate between HA nodes if failures occur.

For iSCSI, Cloud Volumes ONTAP uses multipath I/O (MPIO) and Asymmetric Logical Unit Access (ALUA) to manage path failover between the active-optimized and non-optimized paths.



For information about which specific host configurations support ALUA, refer to the [NetApp Interoperability Matrix Tool](#) and the [SAN hosts and cloud clients guide](#) for your host operating system.

Storage takeover, resync, and giveback are all automatic by default. No user action is required.

Storage configurations

You can use an HA pair as an active-active configuration, in which both nodes serve data to clients, or as an active-passive configuration, in which the passive node responds to data requests only if it has taken over the storage for the active node.

Learn about Cloud Volumes ONTAP HA pairs in Google Cloud

A Cloud Volumes ONTAP high-availability (HA) configuration provides nondisruptive operations and fault tolerance. In Google Cloud, data is synchronously mirrored between the two nodes.

HA components

Cloud Volumes ONTAP HA configurations in Google Cloud include the following components:

- Two Cloud Volumes ONTAP nodes whose data is synchronously mirrored between each other.

- A mediator instance that provides a communication channel between the nodes to assist in storage takeover and giveback processes.
- One zone or three zones (recommended).

If you choose three zones, the two nodes and mediator are in separate Google Cloud zones.


- Four Virtual Private Clouds (VPCs).

The configuration uses four VPCs because GCP requires that each network interface resides in a separate VPC network.

- Four Google Cloud internal load balancers (TCP/UDP) that manage incoming traffic to the Cloud Volumes ONTAP HA pair.

[Learn about networking requirements](#), including more details about load balancers, VPCs, internal IP addresses, subnets, and more.

The following conceptual image shows a Cloud Volumes ONTAP HA pair and its components:



Mediator

Here are some key details about the mediator instance in Google Cloud:

Instance type

e2-micro (an f1-micro instance was previously used)

Disk

Two standard persistent disks that are 10 GiB each

Operating system

Debian 11



For Cloud Volumes ONTAP 9.10.0 and earlier, Debian 10 was installed on the mediator.

Upgrades

When you upgrade Cloud Volumes ONTAP, the NetApp Console also updates the mediator instance as needed.

Access to the instance

For Debian, the default cloud user is `admin`. Google Cloud creates and adds a certificates for the `admin` user when SSH access is requested through the Google Cloud Console or `gcloud` command line. You can specify `sudo` to gain root privileges.

Third-party agents

Third-party agents or VM extensions are not supported on the mediator instance.

Storage takeover and giveback

If a node goes down, the other node can serve data for its partner to provide continued data service. Clients can access the same data from the partner node because the data was synchronously mirrored to the partner.

After the node reboots, the partner must resync data before it can return the storage. The time that it takes to resync data depends on how much data was changed while the node was down.

Storage takeover, resync, and giveback are all automatic by default. No user action is required.

RPO and RTO

An HA configuration maintains high availability of your data as follows:

- The recovery point objective (RPO) is 0 seconds.

Your data is transactionally consistent with no data loss.

- The recovery time objective (RTO) is 120 seconds.

In the event of an outage, data should be available in 120 seconds or less.

HA deployment models

You can ensure the high availability of your data by deploying an HA configuration in multiple zones or in a single zone.

Multiple zones (recommended)

Deploying an HA configuration across three zones ensures continuous data availability if a failure occurs within a zone. Note that write performance is slightly lower compared to using a single zone, but it's minimal.

Single zone

When deployed in a single zone, a Cloud Volumes ONTAP HA configuration uses a spread placement policy. This policy ensures that an HA configuration is protected from a single point of failure within the zone, without having to use separate zones to achieve fault isolation.

This deployment model does lower your costs because there are no data egress charges between zones.

How storage works in an HA pair

Unlike an ONTAP cluster, storage in a Cloud Volumes ONTAP HA pair in GCP is not shared between nodes. Instead, data is synchronously mirrored between the nodes so that the data is available in the event of failure.

Storage allocation

When you create a new volume and additional disks are required, the Console allocates the same number of disks to both nodes, creates a mirrored aggregate, and then creates the new volume. For example, if two disks are required for the volume, the Console allocates two disks per node for a total of four disks.

Storage configurations

You can use an HA pair as an active-active configuration, in which both nodes serve data to clients, or as an active-passive configuration, in which the passive node responds to data requests only if it has taken over storage for the active node.

Performance expectations for an HA configuration

A Cloud Volumes ONTAP HA configuration synchronously replicates data between nodes, which consumes network bandwidth. As a result, you can expect the following performance in comparison to a single-node Cloud Volumes ONTAP configuration:

- For HA configurations that serve data from only one node, read performance is comparable to the read performance of a single-node configuration, whereas write performance is lower.
- For HA configurations that serve data from both nodes, read performance is higher than the read performance of a single-node configuration, and write performance is the same or higher.

For more details about Cloud Volumes ONTAP performance, refer to [Performance](#).


Client access to storage

Clients should access NFS and CIFS volumes by using the data IP address of the node on which the volume resides. If NAS clients access a volume by using the IP address of the partner node, traffic goes between both nodes, which reduces performance.

If you move a volume between nodes in an HA pair, you should remount the volume by using the IP address of the other node. Otherwise, you can experience reduced performance. If clients support NFSv4 referrals or folder redirection for CIFS, you can enable those features on the Cloud Volumes ONTAP systems to avoid remounting the volume. For details, refer to ONTAP documentation.



You can locate the correct IP address from the Console by selecting the volume and clicking **Mount Command**.



Related links

- [Learn about networking requirements](#)
- [Learn how to get started in GCP](#)

Operations unavailable when a node in Cloud Volumes ONTAP HA pair is offline

When a node in an HA pair isn't available, the other node serves data for its partner to provide continued data service. This is called *storage takeover*. Several actions are unavailable until storage giveback is complete.



When a node in an HA pair is unavailable, the state of the system in the NetApp Console is *Degraded*.

The following actions are unavailable from storage takeover:

- Support registration
- License changes
- Instance or VM type changes
- Write speed changes
- CIFS setup
- Changing the location of configuration backups
- Setting the cluster password
- Managing disks and aggregates (advanced allocation)

These actions are available again after storage giveback completes and the state of the system changes back to normal.

Learn about Cloud Volumes ONTAP data encryption and ransomware protection

Cloud Volumes ONTAP supports data encryption and provides protection against viruses and ransomware.

Encryption of data at rest

Cloud Volumes ONTAP supports the following encryption technologies:

- NetApp encryption solutions (NVE and NAE)
- AWS Key Management Service
- Azure Storage Service Encryption
- Google Cloud Platform default encryption

You can use NetApp encryption solutions with native encryption from your cloud provider, which encrypts data at the hypervisor level. Doing so would provide double encryption, which might be desired for very sensitive data. When the encrypted data is accessed, it's unencrypted twice—once at the hypervisor-level (using keys from the cloud provider) and then again using NetApp encryption solutions (using keys from an external key manager).

NetApp encryption solutions (NVE and NAE)

Cloud Volumes ONTAP supports [NetApp Volume Encryption \(NVE\)](#) and [NetApp Aggregate Encryption \(NAE\)](#). NVE and NAE are software-based solutions that enable (FIPS) 140-2-compliant data-at-rest encryption of volumes. Both NVE and NAE use AES 256-bit encryption.

- NVE encrypts data at rest one volume at a time. Each data volume has its own unique encryption key.
- NAE is an extension of NVE—it encrypts data for each volume, and the volumes share a key across the aggregate. NAE also allows common blocks across all volumes in the aggregate to be deduplicated.

Cloud Volumes ONTAP supports both NVE and NAE with external key management services (EKMs) provided by AWS, Azure, and Google Cloud, including third-party solutions, such as Fortanix. Unlike ONTAP, for Cloud Volumes ONTAP, encryption keys are generated at the cloud provider's side, not in ONTAP. Cloud Volumes ONTAP doesn't support [Onboard Key Manager](#).

Cloud Volumes ONTAP uses the standard Key Management Interoperability Protocol (KMIP) services that ONTAP uses. For more information about the supported services, refer to the [Interoperability Matrix Tool](#).

If you use NVE, you have the option to use your cloud provider's key vault to protect ONTAP encryption keys:

- AWS Key Management Service (KMS)
- Azure Key Vault (AKV)
- Google Cloud Key Management Service

New aggregates have NetApp Aggregate Encryption (NAE) enabled by default after you set up an external key manager. New volumes that aren't part of an NAE aggregate have NVE enabled by default (for example, if you have existing aggregates that were created before setting up an external key manager).

Setting up a supported key manager is the only required step. For set up instructions, refer to [Encrypt volumes with NetApp encryption solutions](#).

AWS Key Management Service

When you launch a Cloud Volumes ONTAP system in AWS, you can enable data encryption using the [AWS Key Management Service \(KMS\)](#). The NetApp Console requests data keys using a customer master key (CMK).



You can't change the AWS data encryption method after you create a Cloud Volumes ONTAP system.

If you want to use this encryption option, then you must ensure that the AWS KMS is set up appropriately. For information, refer to [Setting up the AWS KMS](#).

Azure Storage Service Encryption

Data is automatically encrypted on Cloud Volumes ONTAP in Azure using [Azure Storage Service Encryption](#) with a Microsoft-managed key.

You can use your own encryption keys if you prefer. [Learn how to set up Cloud Volumes ONTAP to use a customer-managed key in Azure](#).

Google Cloud Platform default encryption

[Google Cloud Platform data-at-rest encryption](#) is enabled by default for Cloud Volumes ONTAP. No setup is required.

While Google Cloud Storage always encrypts your data before it's written to disk, you can use the Console APIs to create a Cloud Volumes ONTAP system that uses *customer-managed encryption keys*. These are keys that you generate and manage in GCP using the Cloud Key Management Service. [Learn more](#).

ONTAP virus scanning

You can use integrated antivirus functionality on ONTAP systems to protect data from being compromised by viruses or other malicious code.

ONTAP virus scanning, called *Vscan*, combines best-in-class third-party antivirus software with ONTAP features that give you the flexibility you need to control which files get scanned and when.

For information about the vendors, software, and versions supported by Vscan, refer to the [NetApp Interoperability Matrix](#).

For information about how to configure and manage the antivirus functionality on ONTAP systems, refer to the [ONTAP 9 Antivirus Configuration Guide](#).

Ransomware protection

Ransomware attacks can cost a business time, resources, and reputation. The Console enables you to implement the NetApp solution for ransomware, which provides effective tools for visibility, detection, and remediation.

- The Console identifies volumes that are not protected by a Snapshot policy and enables you to activate the default Snapshot policy on those volumes.

Snapshot copies are read-only, which prevents ransomware corruption. They can also provide the granularity to create images of a single file copy or a complete disaster recovery solution.

- The Console also enables you to block common ransomware file extensions by enabling ONTAP's FPolicy solution.

Ransomware Protection

Ransomware attacks can cost a business time, resources, and reputation. The NetApp solution for ransomware provides effective tools for visibility, detection, and remediation. [Learn More](#)

1 Enable Snapshot Copy Protection 



1 Volumes without a Snapshot Policy

To protect your data, activate the default Snapshot policy for these volumes 

[Activate Snapshot Policy](#)

2 Block Ransomware File Extensions 



ONTAP's native FPolicy configuration monitors and blocks file operations based on a file's extension.

[View Denied File Names](#) 

[Activate FPolicy](#)

[Learn how to implement the NetApp solution for ransomware.](#)

Learn about performance monitoring for Cloud Volumes ONTAP workloads

You can review performance results to help you decide which workloads are appropriate for Cloud Volumes ONTAP.

Performance technical reports

- Cloud Volumes ONTAP for AWS

[NetApp Technical Report 4383: Performance Characterization of Cloud Volumes ONTAP in Amazon Web Services with Application Workloads](#)

- Cloud Volumes ONTAP for Microsoft Azure

[NetApp Technical Report 4671: Performance Characterization of Cloud Volumes ONTAP in Azure with Application Workloads](#)

- Cloud Volumes ONTAP for Google Cloud

[NetApp Technical Report 4816: Performance Characterization of Cloud Volumes ONTAP for Google Cloud](#)

CPU performance

Cloud Volumes ONTAP nodes show as highly utilized (over 90%) from your cloud provider's monitoring tools. This is because ONTAP reserves all vCPUs presented to the virtual machine so that they are available when needed.

For information, refer to the [NetApp knowledgebase article about how to monitor ONTAP CPU utilization using the CLI](#)

License management for node-based BYOL

Each Cloud Volumes ONTAP system that has a node-based bring your own license (BYOL) must have a system license installed with an active subscription. The NetApp Console simplifies the process by managing licenses for you and by displaying a warning before they expire.



A node-based license is the previous generation license for Cloud Volumes ONTAP. A node-based license could be procured from NetApp (BYOL) and is available for license renewals, only in specific cases.

[Learn more about Cloud Volumes ONTAP licensing options.](#)

[Learn more about how to manage node-based licenses.](#)

BYOL system licenses

Node-based licenses could be procured from NetApp. The number of licenses that you can purchase for a single-node system or HA pair is unlimited.



NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAPP](#).

A node-based license provides up to 368 TiB of capacity for a single node or HA pair. You might have purchased multiple licenses for a Cloud Volumes ONTAP BYOL system to allocate more than 368 TiB of capacity. For example, you might have two licenses to allocate up to 736 TiB of capacity to Cloud Volumes ONTAP. Or you could have four licenses to get up to 1.4 PiB.

Be aware that disk limits can prevent you from reaching the capacity limit by using disks alone. You can go beyond the disk limit by [tiering inactive data to object storage](#). For information about disk limits, refer to [storage limits in the Cloud Volumes ONTAP Release Notes](#).

License management for a new system

When you create a node-based BYOL system, the Console prompts you for the serial number of your license and your NetApp Support Site account. The Console uses the account to download the license file from NetApp and to install it on the Cloud Volumes ONTAP system.

[Learn how to add NetApp Support Site accounts to the Console.](#)

If the Console can't access the license file over the secure internet connection, you can [obtain the file yourself and then manually upload the file to the Console](#).

License expiration

The Console displays a warning 30 days before a node-based license is due to expire and again when the license expires. The following image shows a 30-day expiration warning that appears in the user interface:



You can select the system to review the message.

The Console includes a license expiration warning in the Cloud Volumes ONTAP report that's emailed to you, if you are an organization or account admin and you enabled the option. The emailed report includes the license expiration warning every 2 weeks.

If you don't renew the license in time, the Cloud Volumes ONTAP system shuts itself down. If you restart it, it shuts itself down again.

License renewal

If you renew a node-based BYOL subscription by contacting a NetApp representative, the Console automatically obtains the new license from NetApp and installs it on the Cloud Volumes ONTAP system.

If the Console can't access the license file over the secure internet connection, you can [obtain the file yourself and then manually upload the file to the Console](#).

License transfer to a new system

A node-based BYOL license is transferable between Cloud Volumes ONTAP systems when you delete an existing system and then create a new one using the same license.

For example, you might want to delete an existing licensed system and then use the license with a new BYOL system in a different VPC/VNet or cloud provider. Note that only *cloud-agnostic* serial numbers work in any cloud provider. Cloud-agnostic serial numbers start with the 908xxxx prefix.

It's important to note that your BYOL license is tied to your company and a specific set of NetApp Support Site credentials.

Learn how AutoSupport and Digital Advisor are used for Cloud Volumes ONTAP

The AutoSupport component of ONTAP collects telemetry and sends it for analysis. Active IQ Digital Advisor (also known as Digital Advisor) analyzes the data from AutoSupport and provides proactive care and optimization. Using artificial intelligence, Digital Advisor can identify potential problems and help you resolve them before they impact your business.

Digital Advisor enables you to optimize your data infrastructure across your global hybrid cloud by delivering actionable predictive analytics and proactive support through a cloud-based portal and mobile app. Data-driven insights and recommendations from Digital Advisor are available to all NetApp customers with an active SupportEdge contract (features vary by product and support tier).

Here are some things you can do with Digital Advisor:

- Plan upgrades.

Digital Advisor identifies issues in your environment that can be resolved by upgrading to a newer version of ONTAP and the Upgrade Advisor component helps you plan for a successful upgrade.

- View system wellness.

Your Digital Advisor dashboard reports any issues with wellness and helps you correct those issues. Monitor system capacity to make sure you never run out of storage space. View support cases for your system.

- Manage performance.

Digital Advisor shows system performance over a longer period than you can see in ONTAP System Manager. Identify configuration and system issues that are impacting your performance.

Maximize efficiency. View storage efficiency metrics and identify ways to store more data in less space.

- View inventory and configuration.

Digital Advisor displays complete inventory and software and hardware configuration information. See when service contracts are expiring and renew them to ensure you remain supported.

Related links

- [NetApp Documentation: Digital Advisor](#)
- [Launch Digital Advisor](#)
- [SupportEdge Services](#)

Supported default configurations for Cloud Volumes ONTAP

Understanding how Cloud Volumes ONTAP is configured by default can help you set up and administer your systems, especially if you are familiar with ONTAP because the default setup for Cloud Volumes ONTAP is different than ONTAP.

Default setup

- The NetApp Console creates one data-serving storage VM when it deploys Cloud Volumes ONTAP. Some configurations support additional storage VMs. [Learn more about managing storage VMs](#).

Beginning with the 3.9.5 release, logical space reporting is enabled on the initial storage VM. When space is reported logically, ONTAP reports the volume space such that all the physical space saved by the storage efficiency features are also reported as used. For information about inline storage efficiency features, refer to the knowledge base article [KB: What Inline Storage Efficiency features are supported with CVO?](#)

- The Console automatically installs the following ONTAP feature licenses on Cloud Volumes ONTAP:
 - CIFS
 - FlexCache
 - FlexClone
 - iSCSI

- Multi-tenant Encryption Key Management (MTEKM), starting with Cloud Volumes ONTAP 9.12.1 GA
- NetApp Volume Encryption (only for bring your own license (BYOL) or registered pay-as-you-go (PAYGO) systems)
- NFS
- ONTAP S3

Starting with Cloud Volumes ONTAP 9.11.0 in AWS

Starting with Cloud Volumes ONTAP 9.9.1 in Azure

- SnapMirror
- SnapRestore
- SnapVault

- Several network interfaces are created by default:

- A cluster management LIF
- An intercluster LIF
- An SVM management LIF on HA systems in Azure
- An SVM management LIF on HA systems in Google Cloud
- An SVM management LIF on single-node systems in AWS
- A node management LIF

In Google Cloud, this LIF is combined with the intercluster LIF.

- An iSCSI data LIF
- A CIFS and NFS data LIF



LIF failover is disabled by default for Cloud Volumes ONTAP due to cloud provider requirements. Migrating a LIF to a different port breaks the external mapping between IP addresses and network interfaces on the instance, making the LIF inaccessible.

- Cloud Volumes ONTAP sends configuration backups to the Console agent using HTTP.

The backups are accessible from `http://ipaddress/occm/offboxconfig/` where `ipaddress` is the IP address of the host of the Console agent.

You can use the backups for reconfiguring your Cloud Volumes ONTAP system. For more information about configuration backups, refer to the [ONTAP documentation](#).

- The Console sets a few volume attributes differently than other management tools (ONTAP System Manager or the ONTAP CLI, for example).

The following table lists the volume attributes set differently from the defaults:

Attribute	Value that the Console configures
Autosize mode	grow

Attribute	Value that the Console configures
Maximum autosize	1,000 percent <p> The organization or account admin can modify this value from the Settings page.</p>
Security style	NTFS for CIFS volumes UNIX for NFS volumes
Space guarantee style	none
UNIX permissions (NFS only)	777

For information about these attributes, refer to [ONTAP volume create man page](#).

Internal disks for system data

In addition to the storage for user data, the Console also purchases cloud storage for system data.

AWS

- Three disks per node for boot, root, and core data:
 - 47 GiB io1 disk for boot data
 - 140 GiB gp3 disk for root data
 - 540 GiB gp2 disk for core data
- For HA pairs:
 - Two st1 EBS volumes for the mediator instance, one of approximately 8 GiB as root disk, and one of 4 GiB as data disk
 - One 140 GiB gp3 disk in each node to contain a copy of the root data of the other node



In some zones, the available EBS disk type can only be gp2.

- One EBS snapshot for each boot disk and root disk



Snapshots are created automatically upon reboot.

- When you enable data encryption in AWS using the Key Management Service (KMS), the boot and root disks for Cloud Volumes ONTAP are encrypted, as well. This includes the boot disk for the mediator instance in an HA pair. The disks are encrypted using the CMK that you select when you add a Cloud Volumes ONTAP system.



In AWS, NVRAM is on the boot disk.

Azure (single node)

- Three Premium SSD disks:
 - One 10 GiB disk for boot data
 - One 140 GiB disk for root data
 - One 512 GiB disk for NVRAM

If the virtual machine that you chose for Cloud Volumes ONTAP supports Ultra SSDs, then the system uses a 32 GiB Ultra SSD for NVRAM, rather than a Premium SSD.

- One 1024 GiB Standard HDD disk for saving cores
- One Azure snapshot for each boot disk and root disk
- Every disk by default in Azure is encrypted at rest.

If the virtual machine that you chose for Cloud Volumes ONTAP supports Premium SSD v2 Managed Disk as data disks, the system uses a 32 GiB Premium SSD v2 Managed Disk for NVRAM, and another one as the root disk.

Azure (HA pair)

HA pairs with page blob

- Two 10 GiB Premium SSD disks for the boot volume (one per node)
- Two 140 GiB Premium Storage page blobs for the root volume (one per node)
- Two 1024 GiB Standard HDD disks for saving cores (one per node)
- Two 512 GiB Premium SSD disks for NVRAM (one per node)
- One Azure snapshot for each boot disk and root disk



Snapshots are created automatically upon reboot.

- Every disk by default in Azure is encrypted at rest.

HA pairs with shared managed disks in multiple availability zones

- Two 10 GiB Premium SSD disks for the boot volume (one per node)
- Two 512 GiB Premium SSD disks for the root volume (one per node)
- Two 1024 GiB Standard HDD disks for saving cores (one per node)
- Two 512 GiB Premium SSD disks for NVRAM (one per node)
- One Azure snapshot for each boot disk and root disk



Snapshots are created automatically upon reboot.

- Every disk by default in Azure is encrypted at rest.

HA pairs with shared managed disks in single availability zones

- Two 10 GiB Premium SSD disks for the boot volume (one per node)
- Two 512 GiB Premium SSD Shared Managed disks for the root volume (one per node)
- Two 1024 GiB Standard HDD disks for saving cores (one per node)

- Two 512 GiB Premium SSD Managed disks for NVRAM (one per node)

If your virtual machine supports Premium SSD v2 Managed Disks as data disks, it uses 32 GiB Premium SSD v2 Managed Disks for NVRAM and 512 GiB Premium SSD v2 Shared Managed disks for the root volume.

You can deploy HA pairs in a single single availability zone and use Premium SSD v2 Managed Disks when the following conditions are fulfilled:

- The version of Cloud Volumes ONTAP is 9.15.1 or later.
- The selected region and zone support Premium SSD v2 Managed Disks. For information about the supported regions, refer to [Microsoft Azure website: Products available by region](#).
- The subscription is registered for the Microsoft [Microsoft.Compute/VMOrchestratorZonalMultiFD](#) feature.

Google Cloud (single node)

- One 10 GiB SSD persistent disk for boot data
- One 64 GiB SSD persistent disk for root data
- One 500 GiB SSD persistent disk for NVRAM
- One 315 GiB Standard persistent disk for saving cores
- Snapshots for boot and root data



Snapshots are created automatically upon reboot.

- Boot and root disks are encrypted by default.

Google Cloud (HA pair)

- Two 10 GiB SSD persistent disks for boot data
- Four 64 GiB SSD persistent disk for root data
- Two 500 GiB SSD persistent disk for NVRAM
- Two 315 GiB Standard persistent disk for saving cores
- One 10 GiB Standard persistent disk for mediator data
- One 10 GiB Standard persistent disk for mediator boot data
- Snapshots for boot and root data



Snapshots are created automatically upon reboot.

- Boot and root disks are encrypted by default.

Where the disks reside

Storage layout:

- Boot data resides on a disk attached to the instance or virtual machine.

This disk, which contains the boot image, is not available to Cloud Volumes ONTAP.

- Root data, which contains the system configuration and logs, resides in aggr0.

- The storage virtual machine (SVM) root volume resides in aggr1.
- Data volumes also reside in aggr1.

Copyright information

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

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

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

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

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

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

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

Trademark information

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