



Concepts

Cloud Manager 3.6

NetApp
June 10, 2024

Table of Contents

- Concepts 1
 - Cloud Manager and Cloud Volumes ONTAP overview 1
 - NetApp Cloud Central 2
 - Cloud provider accounts and permissions 3
 - Storage 7
 - High-availability pairs 18
 - Evaluating 26
 - Licensing 26
 - Security 27
 - Performance 29

Concepts

Cloud Manager and Cloud Volumes ONTAP overview


OnCommand Cloud Manager enables you to deploy Cloud Volumes ONTAP, which provides enterprise-class features for your cloud storage, and to easily replicate data across hybrid clouds built on NetApp.

Cloud Manager

Cloud Manager was built with simplicity in mind. It guides you through Cloud Volumes ONTAP setup in a few steps, eases data management by offering simplified storage provisioning and automated capacity management, enables drag-and-drop data replication across a hybrid cloud, and more.

Cloud Manager is required to deploy and manage Cloud Volumes ONTAP, but it can also discover and provision storage for on-premises ONTAP clusters. This provides a central point of control for your cloud and on-premises storage infrastructure.

You can run Cloud Manager in the cloud or in your network—it just needs a connection to the networks in which you want to deploy Cloud Volumes ONTAP. The following image shows Cloud Manager running in AWS and managing Cloud Volumes ONTAP systems in AWS and Azure. It also shows data replication across a hybrid cloud.



[Learn more about Cloud Manager](#)

Cloud Volumes ONTAP

Cloud Volumes ONTAP is a software-only storage appliance that runs the ONTAP data management software in the cloud. You can use Cloud Volumes ONTAP for production workloads, disaster recovery, DevOps, file shares, and database management.

Cloud Volumes ONTAP extends enterprise storage to the cloud with the following key features:

- **Storage efficiencies**
Leverage built-in data deduplication, data compression, thin provisioning, and cloning to minimize storage costs.
- **High availability**
Ensure enterprise reliability and continuous operations in case of failures in your cloud environment.
- **Data replication**
Cloud Volumes ONTAP leverages SnapMirror, NetApp's industry-leading replication technology, to replicate on-premises data to the cloud so it's easy to have secondary copies available for multiple use cases.
- **Data tiering**
Switch between high and low-performance storage pools on-demand without taking applications offline.
- **Application consistency**
Ensure consistency of NetApp Snapshot copies using NetApp SnapCenter.



Licenses for ONTAP features are included with Cloud Volumes ONTAP, except for NetApp Volume Encryption.

[View supported Cloud Volumes ONTAP configurations](#)

[Learn more about Cloud Volumes ONTAP](#)

NetApp Cloud Central

[NetApp Cloud Central](#) provides a centralized location to access and manage NetApp cloud data services. These services enable you to run critical applications in the cloud, create automated DR sites, back up your SaaS data, and effectively migrate and control data across multiple clouds.

Cloud Manager's integration with NetApp Cloud Central provides several benefits, including a simplified deployment experience, a single location to view and manage multiple Cloud Manager systems, and centralized user authentication.

With centralized user authentication, you can use the same set of credentials across Cloud Manager systems and between Cloud Manager and other data services, such as Cloud Sync. It's also easy to reset your password if you forgot it.

The following video provides an overview of NetApp Cloud Central:

Hi Kevin Hill, let's get started!



ONTAP Cloud

Loading...

[More Info](#)



Cloud Sync

Loading...

[More Info](#)



Cloud Control

[Go to Cloud Control](#)

[More Info](#)



Azure NFSaaS

[Register for Preview](#)

[More Info](#)



NFS Hybrid For AWS

[Register for Preview](#)

[More Info](#)

API

Cloud provider accounts and permissions

Cloud Manager enables you to choose the *cloud provider account* in which you want to deploy a Cloud Volumes ONTAP system. You should understand the permissions requirements before you add the accounts to Cloud Manager.

AWS accounts and permissions


You can deploy all of your Cloud Volumes ONTAP systems in the initial AWS account, or you can set up additional accounts.

The initial AWS account


When you deploy Cloud Manager from NetApp Cloud Central, you need to use an AWS account that has permissions to launch the Cloud Manager instance. The required permissions are listed in the [NetApp Cloud Central policy for AWS](#).

When Cloud Central launches the Cloud Manager instance in AWS, it creates an IAM role and an instance profile for the instance. It also attaches a policy that provides Cloud Manager with permissions to deploy and manage Cloud Volumes ONTAP in that AWS account. [Review how Cloud Manager uses the permissions](#).

Cloud Central



AWS account



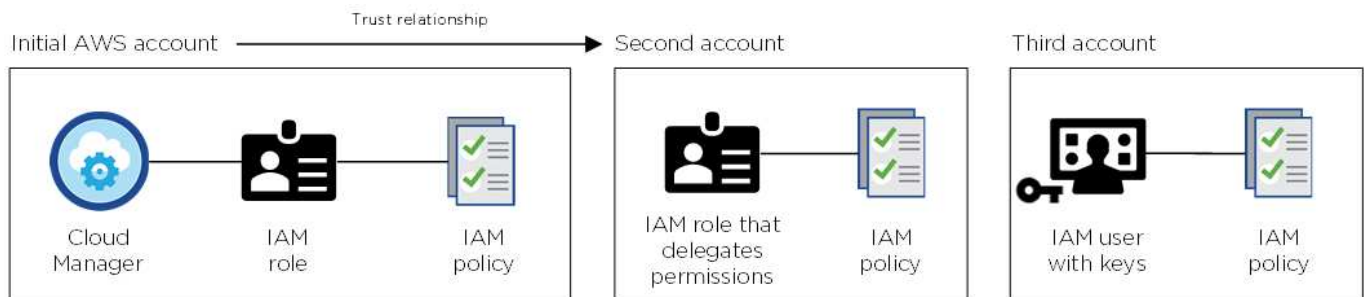
Cloud Manager selects this cloud provider account by default when you create a new working environment:

Details & Credentials

This working environment will be created in Cloud Provider Account: **Instance Profile** | Account ID: XXXXXXXXXX | [Switch Account](#)

Additional AWS accounts

If you want to launch Cloud Volumes ONTAP in different AWS accounts, then you can either [provide AWS keys for an IAM user or the ARN of a role in a trusted account](#). The following image shows two additional accounts, one providing permissions through an IAM role in a trusted account and another through the AWS keys of an IAM user:



You would then [add the cloud provider accounts to Cloud Manager](#) by specifying the Amazon Resource Name (ARN) of the IAM role, or the AWS keys for the IAM user.

After you add another account, you can switch to it when creating a new working environment:

Cloud Provider Profile Name

QA | Account ID: [blurred]

Instance Profile | Account ID: [blurred]

To add a new AWS cloud provider account, go to the [Cloud Provider Account Settings](#).

Apply

Cancel

Azure accounts and permissions

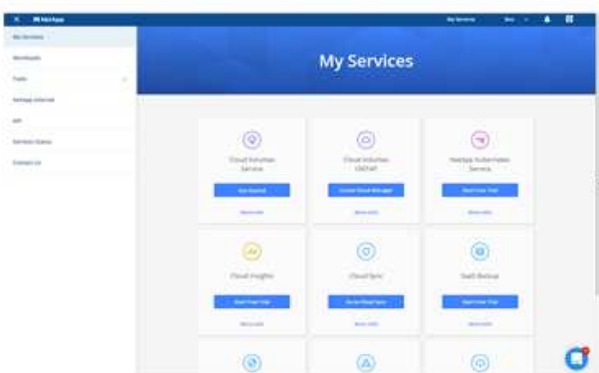
You can deploy all of your Cloud Volumes ONTAP systems in the initial Azure account, or you can set up additional accounts.

The initial Azure account

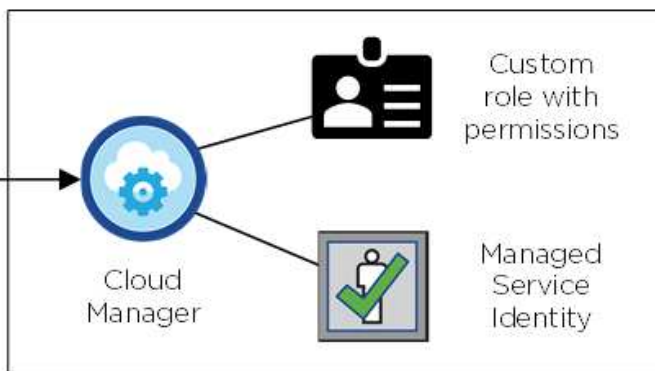
When you deploy Cloud Manager from NetApp Cloud Central, you need to use an Azure account that has permissions to deploy the Cloud Manager virtual machine. The required permissions are listed in the [NetApp Cloud Central policy for Azure](#).

When Cloud Central deploys the Cloud Manager virtual machine in Azure, it enables a [system-assigned managed identity](#) on the Cloud Manager virtual machine, creates a custom role, and assigns it to the virtual machine. The role provides Cloud Manager with permissions to deploy and manage Cloud Volumes ONTAP in that Azure subscription. [Review how Cloud Manager uses the permissions](#).

Cloud Central



Azure account



Cloud Manager selects this cloud provider account by default when you create a new working environment:

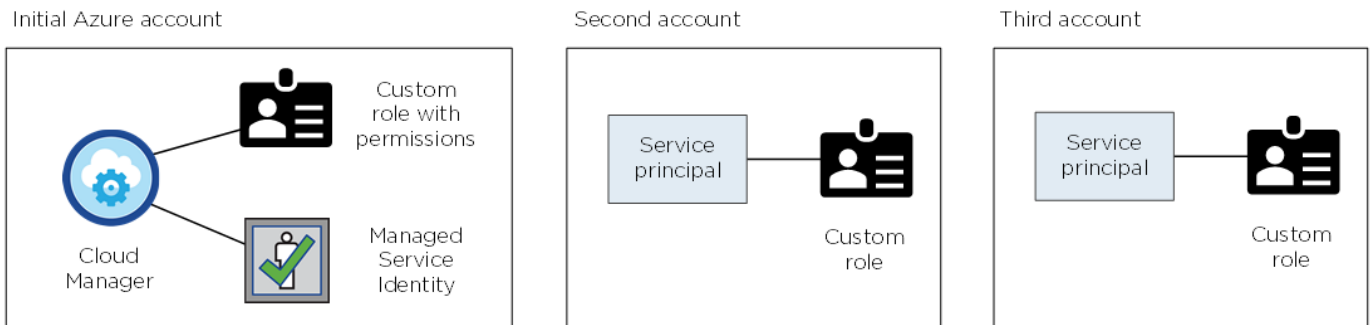
This working environment will be created in Cloud Provider Account: **Managed Service Identity** | Azure Subscription: **OCCM QA1** | [Switch Account](#)

Additional Azure subscriptions for the initial account

The managed identity is associated with the subscription in which you launched Cloud Manager. If you want to select a different Azure subscription, then you need to [associate the managed identity with those subscriptions](#).

Additional Azure accounts

If you want to deploy Cloud Volumes ONTAP in different Azure accounts, then you must grant the required permissions by [creating and setting up a service principal in Azure Active Directory](#) for each Azure account. The following image shows two additional accounts, each set up with a service principal and custom role that provides permissions:



You would then [add the cloud provider accounts to Cloud Manager](#) by providing details about the AD service principal.

After you add another account, you can switch to it when creating a new working environment:



Microsoft Azure Provider Account

Cloud Provider Profile Name

Azure Keys | Application ID: [redacted] ...
Dev Keys | Application ID: [redacted] ...
Managed Service Identity

To add a new Azure cloud provider account, go to the [Cloud Provider Account Settings](#).

Apply

Cancel

What about Marketplace deployments and on-prem deployments?

The sections above describe the recommended deployment method from NetApp Cloud Central. You can also deploy Cloud Manager from the [AWS Marketplace](#), the [Azure Marketplace](#), and you can [install Cloud Manager on-premises](#).

If you use either of the Marketplaces, permissions are provided in the same way. You just need to manually create and set up the IAM role or managed identity for Cloud Manager, and then provide permissions for any additional accounts.

For on-premises deployments, you can't set up an IAM role or managed identity for the Cloud Manager system, but you can provide permissions just like you would for additional accounts.

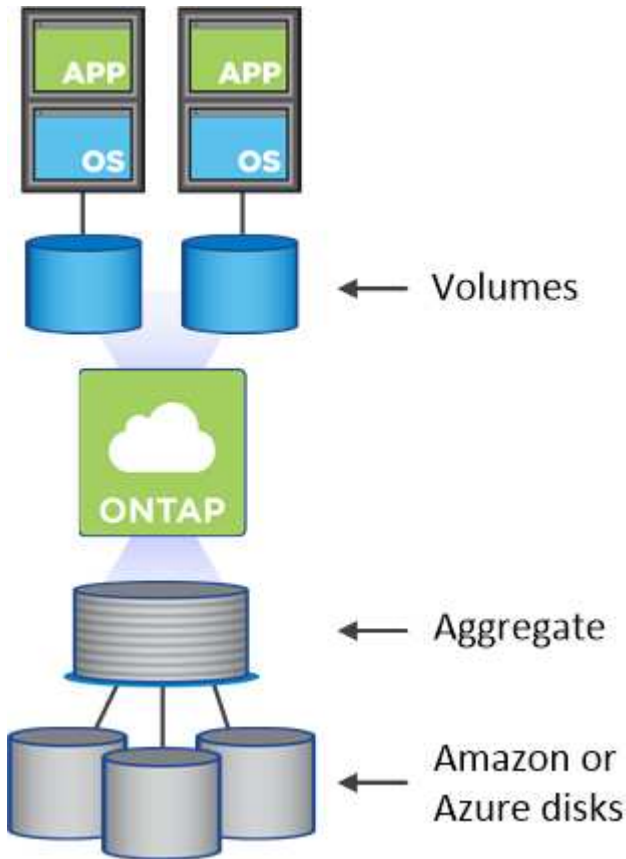
Storage

How Cloud Volumes ONTAP uses cloud storage

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

Overview

Cloud Volumes ONTAP uses AWS and Azure volumes as back-end storage. It sees these volumes 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 creating volumes and the default disk size when you deploy Cloud Volumes ONTAP.



The total amount of storage purchased from AWS or Azure 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 Cloud Manager creates a 500 GB aggregate, the usable capacity is 442.94 GB.

AWS storage

In AWS, an aggregate can contain up to 6 disks that are all the same size. The maximum disk size is 16 TB.

The underlying EBS disk type can be either General Purpose SSD, Provisioned IOPS SSD, Throughput Optimized HDD, or Cold HDD. You can also pair an EBS disk with Amazon S3 for [data tiering](#).

At a high level, the differences between EBS disk types are as follows:

- *General Purpose SSD* disks balance cost and performance for a broad range of workloads. Performance is defined in terms of IOPS.
- *Provisioned IOPS SSD* disks are for critical applications that require the highest performance at a higher cost.
- *Throughput Optimized HDD* disks are for frequently accessed workloads that require fast and consistent throughput at a lower price.
- *Cold HDD* disks are meant for backups, or infrequently accessed data, because the performance is very low. Like Throughput Optimized HDD disks, performance is defined in terms of throughput.



Cold HDD disks are not supported with HA configurations and with data tiering.

For additional details about the use cases for these disks, refer to [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.](#)

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 three types of Azure Managed Disks:

- *Premium SSD Managed Disks* provide high performance for I/O-intensive workloads at a higher cost.
- *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 TB.

You can pair a managed disk with Azure Blob storage for [data tiering](#).

HA pairs

HA pairs use Premium page blobs, which have a maximum disk size of 8 TB.

For additional details about the use cases for these disks, see [Microsoft Azure Documentation: Introduction to Microsoft Azure Storage](#).

[Learn how to choose disk types and disk sizes for your systems in Azure.](#)

[Review storage limits for Cloud Volumes ONTAP.](#)

Data tiering overview

You can 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 (the performance tier), while inactive data is tiered to low-cost object storage (the capacity tier). This enables you to reclaim space on your primary storage and shrink secondary storage.


Cloud Volumes ONTAP supports data tiering in AWS and in Microsoft Azure. Data tiering is powered by FabricPool technology.



You do not need to install a feature license to enable data tiering.

How data tiering works in AWS

When you enable data tiering in AWS, Cloud Volumes ONTAP uses EBS as a performance tier for hot data and AWS S3 as a capacity tier for inactive data:



Performance tier in AWS

The performance tier can be General Purpose SSDs, Provisioned IOPS SSDs, or Throughput Optimized HDDs.

Capacity tier in AWS

By default, Cloud Volumes ONTAP tiers inactive data to the S3 *Standard* storage class. Standard is ideal for frequently accessed data stored across multiple Availability Zones.

If you do not plan to access the inactive data, you can reduce your storage costs by changing a system's tiering level to one of the following, after you deploy Cloud Volumes ONTAP:

Intelligent Tiering

Optimizes storage costs by moving data between two tiers as data access patterns change. One tier is for frequent access and the other is for infrequent access.

One Zone-Infrequent Access

For infrequently accessed data stored in a single Availability Zone.

Standard-Infrequent Access

For infrequently accessed data stored across multiple Availability Zones.

The access costs are higher if you do access the data, so you must take that into consideration before you change the tiering level. For more details about S3 storage classes, refer to [AWS documentation](#).

When you change the tiering level, inactive data starts in the Standard storage class and moves to the storage class that you selected, if the data is not accessed after 30 days. For details about changing the tiering level, see [Tiering inactive data to low-cost object storage](#).


The tiering level is system wide—it is not per volume.



A Cloud Volumes ONTAP working environment uses an S3 bucket for all tiered data from the system. A different S3 bucket is not used for each volume. This includes an HA working environment. Cloud Manager creates an S3 bucket and names it *fabric-pool-cluster unique identifier*.

How data tiering works in Microsoft 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 in Azure

The performance tier can be either Premium Storage (SSD) or Standard Storage (HDD).

Capacity tier in Azure

By default, Cloud Volumes ONTAP tiers inactive data to the Azure *hot* storage tier, which is ideal for frequently accessed data.

If you do not plan to access the inactive data, you can reduce your storage costs by changing a system's tiering level to the Azure *cool* storage tier after you deploy Cloud Volumes ONTAP. The cool tier is ideal for infrequently accessed data that will reside in the tier for at least 30 days.

The access costs are higher if you do access the data, so you must take that into consideration before you change the tiering level. For more details about Azure Blob storage tiers, refer to [Azure documentation](#).

When you change the tiering level, inactive data starts in the hot storage tier and moves to the cool storage tier, if the data is not accessed after 30 days. For details about changing the tiering level, see [Tiering inactive data to low-cost object storage](#).

The tiering level is system wide—it is not per volume.



A Cloud Volumes ONTAP working environment uses an Azure Blob container for all tiered data from the system. A different container is not used for each volume. Cloud Manager creates a new storage account with a container for each Cloud Volumes ONTAP system. The name of the storage account is random.

How data tiering affects 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.

Cloud Volumes ONTAP supports the following tiering policies:

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.

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.

Backup

When you replicate a volume for disaster recovery or long-term retention, data for the destination volume starts in the capacity tier. If you activate the destination volume, the data gradually moves to the performance tier as it is read.

None

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

Setting up data tiering

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

Storage management

Cloud Manager provides simplified and advanced management of Cloud Volumes ONTAP storage.



All disks and aggregates must be created and deleted directly from Cloud Manager. 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

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

Simplified provisioning

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

When you create a volume, Cloud Manager 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.
- It purchases disks for a new aggregate and places the volume on that aggregate.

Cloud Manager 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.



The Cloud Manager Admin can modify free space thresholds from the **Settings** page.

Disk size selection for aggregates in AWS

When Cloud Manager creates new aggregates for Cloud Volumes ONTAP in AWS, it gradually increases the disk size in an aggregate, as the number of aggregates in the system increases. Cloud Manager does this to ensure that you can utilize the system's maximum capacity before it reaches the maximum number of data disks allowed by AWS.

For example, Cloud Manager might choose the following disk sizes for aggregates in a Cloud Volumes ONTAP Premium or BYOL system:

Aggregate number	Disk size	Max aggregate capacity
1	500 MB	3 TB
4	1 TB	6 TB
6	2 TB	12 TB

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

Advanced allocation

Rather than let Cloud Manager manage aggregates for you, you can do it yourself. [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 Cloud Manager Admin can choose whether Cloud Manager notifies you of storage capacity decisions or whether Cloud Manager automatically manages capacity requirements for you. It might help for you to understand how these modes work.

Automatic capacity management

If the Cloud Manager Admin set the Capacity Management Mode to automatic, Cloud Manager automatically purchases new disks for Cloud Volumes ONTAP instances when more capacity is needed, deletes unused collections of disks (aggregates), moves volumes between aggregates when needed, and attempts to unfill disks.

The following examples illustrate how this mode works:

- If an aggregate with 5 or fewer EBS disks reaches the capacity threshold, Cloud Manager automatically purchases new disks for that aggregate so volumes can continue to grow.
- If an aggregate with 12 Azure disks reaches the capacity threshold, Cloud Manager automatically moves a volume from that aggregate to an aggregate with available capacity or to a new aggregate.

If Cloud Manager 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 cannot be returned to AWS or Azure in this scenario.

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

Manual capacity management

If the Cloud Manager Admin set the Capacity Management Mode to manual, Cloud Manager displays Action Required messages when capacity decisions must be made. The same examples described in the automatic mode apply to the manual mode, but it is up to you to accept the actions.

Storage isolation using tenants

Cloud Manager enables you to provision and manage storage in isolated groups called tenants. You need to decide how to organize Cloud Manager users and their working environments across tenants.

Working environments

Cloud Manager represents storage systems as *working environments*. A working environment is any of the following:

- A single Cloud Volumes ONTAP system or an HA pair
- An on-premises ONTAP cluster in your network
- An ONTAP cluster in a NetApp Private Storage configuration

The following image shows a Cloud Volumes ONTAP working environment:

Tenants

A *tenant* isolates working environments in groups. You create one or more working environments within a tenant. The following image shows three tenants defined in Cloud Manager:

User management of tenants and working environments

The tenants and working environments that Cloud Manager users can manage depend on user role and assignments. The three distinct user roles are as follows:

Cloud Manager Admin

Administers the product and can access all tenants and working environments.

Tenant Admin

Administers a single tenant. Can create and manage all working environments and users in the tenant.

Working Environment Admin

Can create and manage one or more working environments in a tenant.

Example of how you can create tenants and users

If your organization has departments that operate independently, it is best to have a tenant for each department.


For example, you might create three tenants for three separate departments. You would then create a Tenant Admin for each tenant. Within each tenant would be one or more Working Environment Admins who manage working environments. The following image depicts this scenario:

Simplified storage management using the Volume View

Cloud Manager provides a separate management view called the *Volume View*, which further simplifies storage management in AWS.

The Volume View enables you to simply specify the NFS volumes that you need in AWS and then Cloud Manager handles the rest: it deploys Cloud Volumes ONTAP systems as needed and it makes capacity allocation decisions as volumes grow. This view gives you the benefits of enterprise-class storage in the cloud with very little storage management.

The following image shows how you interact with Cloud Manager in the Volume View:




1. You create NFS volumes.
2. Cloud Manager launches Cloud Volumes ONTAP instances in AWS for new volumes or it creates volumes on existing instances. It also purchases physical EBS storage for the volumes.
3. You make the volumes available to your hosts and applications.
4. Cloud Manager makes capacity allocation decisions as your volumes grow.

This means that you simply need to interact with volumes (the image on the left), while Cloud Manager interacts with the storage system and its underlying storage (the image on the right).

Allocation of cloud resources for the initial volume

When you create your first volume, Cloud Manager launches a Cloud Volumes ONTAP instance or a Cloud Volumes ONTAP HA pair in AWS and purchases Amazon EBS storage for the volume:



The size of the initial volume determines the EC2 instance type and the number of EBS disks.



Cloud Manager launches a Cloud Volumes ONTAP Explore or Standard instance, depending on the initial volume size. As the volumes grow, Cloud Manager might prompt you to make an AWS instance change which means it needs to upgrade the instance's license to Standard or Premium. Upgrading increases the EBS raw capacity limit, which allows your volumes to grow.



Cloud Manager does not launch Cloud Volumes ONTAP BYOL instances in the Volume View. You should use Cloud Manager in the Storage System View if you purchased a Cloud Volumes ONTAP license.

Allocation of cloud resources for additional volumes

When you create additional volumes, Cloud Manager creates the volumes on existing Cloud Volumes ONTAP instances or on new Cloud Volumes ONTAP instances. Cloud Manager can create a volume on an existing instance if the instance's AWS location and disk type match the requested volume, and if there is enough space.

NetApp storage efficiency features and storage costs

Cloud Manager automatically enables NetApp storage efficiency features on all volumes. These efficiencies can reduce the total amount of storage that you need. You might see a difference between your allocated capacity and the purchased AWS capacity, which can result in storage cost savings.

Capacity allocation decisions that Cloud Manager automatically handles

- Cloud Manager purchases additional EBS disks as capacity thresholds are exceeded. This happens as your volumes grow.
- Cloud Manager deletes unused sets of EBS disks if the disks contain no volumes for 12 hours.
- Cloud Manager moves volumes between sets of disks to avoid capacity issues.

In some cases, this requires purchasing additional EBS disks. It also frees space on the original set of disks for new and existing volumes.

WORM storage

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. WORM storage is powered by SnapLock technology in Enterprise mode, which means WORM files are protected at the file level.

Once a file has been committed to WORM storage, it cannot 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

You can activate WORM storage on a Cloud Volumes ONTAP system when you create a new working environment. This includes specifying an activation code and setting the default retention period for files. You can obtain an activation code by using the chat icon in the lower right of the Cloud Manager interface.



You cannot activate WORM storage on individual volumes—WORM must be activated at the system level.

The following image shows how to activate WORM storage when creating a working environment:

WORM | *Preview*

You can use **write once, read many (WORM)** storage to retain critical files in unmodified form for regulatory and governance purposes and to protect from malware attacks. WORM files are protected at the file level. [Learn More](#)

Disable WORM Activate WORM

Notice: If you enable WORM storage, you cannot enable data tiering to object storage.

WORM Activation Code 

Worm-1111122222aaaaa

Retention Period

15

years 

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 [ONTAP documentation](#).



Cloud Volumes ONTAP support for WORM storage is equivalent to SnapLock Enterprise mode.

Limitations

- If you delete or move a disk directly from AWS or Azure, then a volume can be deleted before its expiry date.
- When WORM storage is activated, data tiering to object storage cannot be enabled.

High-availability pairs

High-availability 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.

Overview

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.



The mediator instance runs the Linux operating system on a t2.micro instance and uses one EBS magnetic disk that is approximately 8 GB.

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.

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 60 seconds.
In the event of an outage, data should be available in 60 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 AZ. You should review more details about each configuration to choose which best fits your needs.

Cloud Volumes ONTAP HA in 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 by Cloud Manager.

For details, see [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.

Storage 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, see the [NetApp Interoperability Matrix Tool](#) and the Host Utilities Installation and Setup Guide for your host operating system.

Storage 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 easily identify the correct IP address from Cloud Manager by selecting the volume and clicking **Mount**

Command.

Cloud Volumes ONTAP HA in a 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.




Cloud Manager creates an [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.

VPC in AWS



Storage 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, see the [NetApp Interoperability Matrix Tool](#) and the Host Utilities Installation and Setup 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.

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, Cloud Manager 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, Cloud Manager 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 Cloud Manager in the Storage System View.

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, see [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, see ONTAP documentation.

You can easily identify the correct IP address from Cloud Manager. The following image shows the Storage System View:

Volumes

2 Volumes | 0.22 TB Allocated | <0.01 TB Used (0 TB in S3)




The following image shows the Volume View:

High-availability 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

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



Resource group

Note the following about the Azure components that Cloud Manager deploys for you:

Azure Standard Load Balancer

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

Availability Set

The Availability Set ensures that the nodes are in different fault and update domains.

Storage

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 and root data:

- A node's boot data resides on a Premium SSD Managed Disk.
- A node's root data resides on a Premium Storage page blob.

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 60 seconds.
In the event of an outage, data should be available in 60 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, see the [NetApp Interoperability Matrix Tool](#) and the Host Utilities Installation and Setup Guide for your host operating system.

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.

HA limitations

The following limitations affect Cloud Volumes ONTAP HA pairs in Azure:

- HA pairs are supported with Cloud Volumes ONTAP Standard, Premium, and BYOL. Explore is not supported.
- Data tiering is not supported.
- NFSv4 is not supported. NFSv3 is supported.
- HA pairs are not supported in some regions.

[See the list of supported Azure regions.](#)

[Learn how to deploy an HA system in Azure.](#)

Evaluating

You can evaluate Cloud Volumes ONTAP before you pay for the software.

A 30-day free trial of a single-node Cloud Volumes ONTAP system is available from [NetApp Cloud Central](#). There are no hourly software charges, but infrastructure charges still apply. A free trial automatically converts to a paid hourly subscription when it expires.

If you need assistance with your proof of concept, contact [the Sales team](#) or reach out through the chat option available from [NetApp Cloud Central](#) and from within Cloud Manager.

Licensing

Each Cloud Volumes ONTAP BYOL system must have a license installed with an active subscription. If an active license is not installed, the Cloud Volumes ONTAP system shuts itself down after 30 days. Cloud Manager simplifies the process by managing licenses for you and by notifying you before they expire.

License management for a new system

When you create a BYOL system, Cloud Manager prompts you for a NetApp Support Site account. Cloud Manager 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 Cloud Manager.](#)

If Cloud Manager cannot access the license file over the secure internet connection, you can obtain the file yourself and then manually upload the file to Cloud Manager. For instructions, see [Installing license files on Cloud Volumes ONTAP BYOL systems](#).

License expiration

Cloud Manager warns you 30 days before a license is due to expire and again when the license expires. The following image shows a 30-day expiration warning:



You can select the working environment to review the message.

If you do not renew the license in time, the Cloud Volumes ONTAP system shuts itself down. If you restart it, it shuts itself down again.



Cloud Volumes ONTAP can also notify you through email, an SNMP trap host, or syslog server using EMS (Event Management System) event notifications. For instructions, see the [ONTAP 9 EMS Configuration Express Guide](#).

License renewal

When you renew a BYOL subscription by contacting a NetApp representative, Cloud Manager automatically obtains the new license from NetApp and installs it on the Cloud Volumes ONTAP system.

If Cloud Manager cannot access the license file over the secure internet connection, you can obtain the file yourself and then manually upload the file to Cloud Manager. For instructions, see [Installing license files on Cloud Volumes ONTAP BYOL systems](#).

Security

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 Volume Encryption (starting with Cloud Volumes ONTAP 9.5)
- AWS Key Management Service
- Azure Storage Service Encryption

You can use NetApp Volume Encryption with native AWS and Azure encryption, which encrypt data at the hypervisor level.

NetApp Volume Encryption

NetApp Volume Encryption (NVE) is a software-based technology for encrypting data at rest one volume at a time. Data, Snapshot copies, and metadata are encrypted. Access to the data is given by a unique XTS-AES-256 key, one per volume.

Cloud Volumes ONTAP supports NetApp Volume Encryption with an external key management server. An Onboard Key Manager is not supported. You can find the supported key managers in the [NetApp Interoperability Matrix Tool](#) under the **Key Managers** solution.

You can enable NetApp Volume Encryption on a new or existing volume by using the CLI or System Manager. Cloud Manager does not support NetApp Volume Encryption. For instructions, see [Encrypting volumes with NetApp Volume Encryption](#).

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 \(KMS\)](#). Cloud Manager requests data keys using a customer master key (CMK).

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

Azure Storage Service Encryption

[Azure Storage Service Encryption](#) for data at rest is enabled by default for Cloud Volumes ONTAP data in Azure. No setup is required.



Customer-managed keys are not supported with Cloud Volumes ONTAP.

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, see the [NetApp Interoperability Matrix](#).

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

Ransomware protection

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

- Cloud Manager 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.

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

The screenshot displays two panels from the Cloud Manager interface. The left panel, titled '1 Enable Snapshot Copy Protection', features a circular progress indicator showing '40 % Protection' and a red warning icon with the text '3 Volumes without a Snapshot Policy'. Below this, it instructs the user to 'activate the default Snapshot policy for these volumes' and provides an 'Activate Snapshot Policy' button. The right panel, titled '2 Block Ransomware File Extensions', shows a shield icon with an 'F' and a document icon. It explains that 'ONTAP's native FPolicy configuration monitors and blocks file operations based on a file's extension' and offers a 'View Denied File Names' link. An 'Activate FPolicy' button is also present.

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

Performance

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

For Cloud Volumes ONTAP for AWS, refer to [NetApp Technical Report 4383: Performance Characterization of Cloud Volumes ONTAP in Amazon Web Services with Application Workloads](#).

For Cloud Volumes ONTAP for Microsoft Azure, refer to [NetApp Technical Report 4671: Performance Characterization of Cloud Volumes ONTAP in Azure with Application Workloads](#).

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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