



Concepts

Cloud Manager 3.7

NetApp
June 06, 2022

Table of Contents

- Concepts 1
 - Cloud Manager and Cloud Volumes ONTAP overview 1
 - NetApp Cloud Central 2
 - Cloud Central accounts 3
 - Cloud provider accounts 7
 - Storage 13
 - High-availability pairs 21
 - Evaluating 29
 - Licensing 29
 - Security 30
 - Performance 32

Concepts

Cloud Manager and Cloud Volumes ONTAP overview

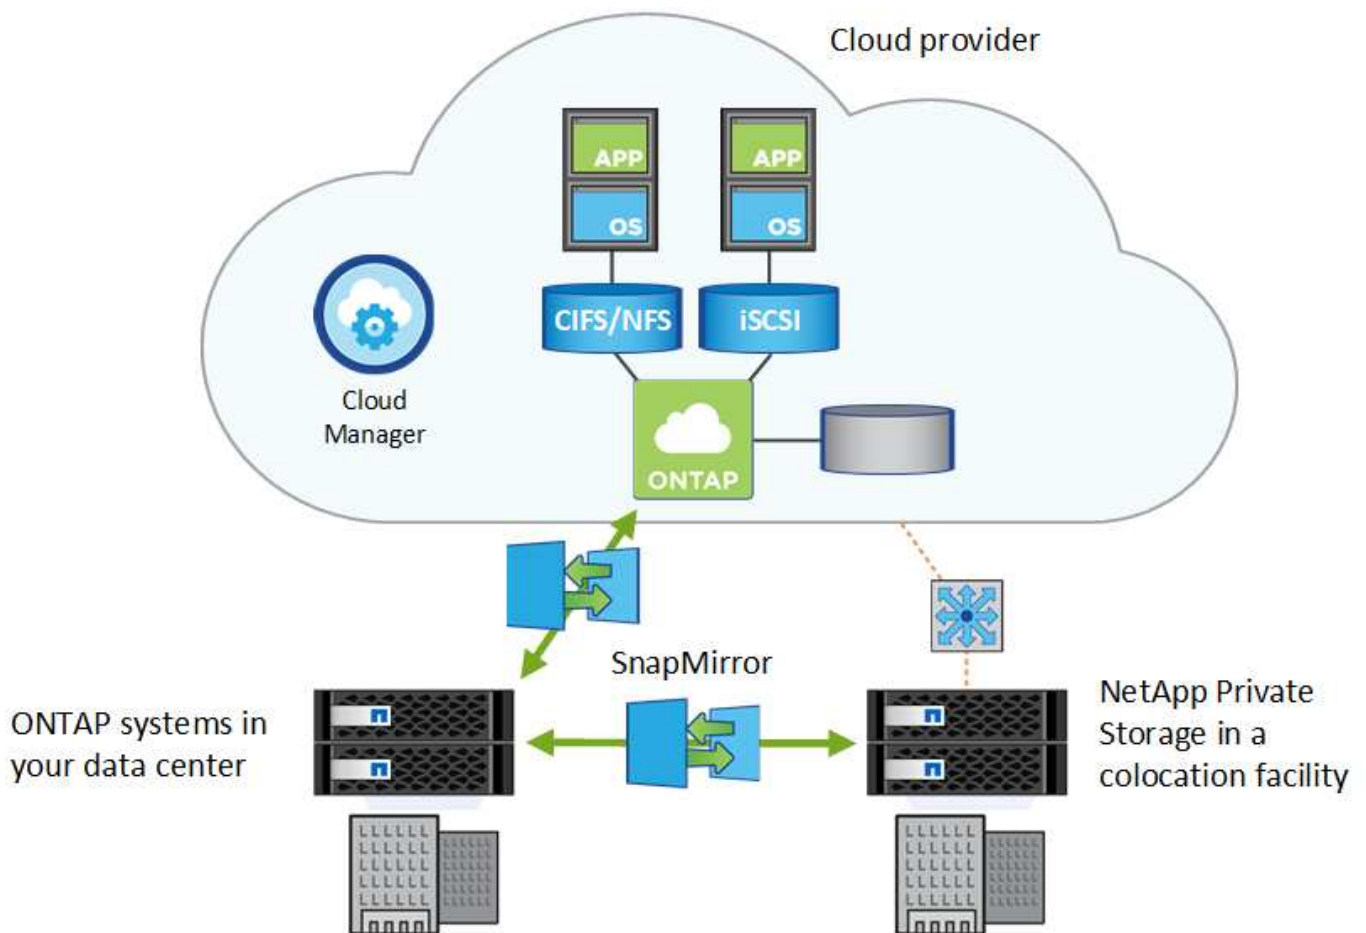
Cloud Manager enables you to deploy Cloud Volumes ONTAP, which provides enterprise-class features for your cloud storage, and to easily replicate data across hybrid clouds built on NetApp.

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 and Cloud Volumes ONTAP running in a cloud provider. 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.

[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.

Fabric View

	Microsoft Azure	Amazon Web Services	Google Cloud Platform	On-Premises
Cloud Sync Go to Cloud Sync				
Cloud Tiering Go to Cloud Tiering				
Cloud Volumes Service Get Started	The industry's leading Network File System (NFS/SMB) service in the cloud			
Cloud Volumes ONTAP Create Cloud Manager	Simple & Fast Enterprise Cloud Storage			
Kubernetes Service Go to	The Universal Control Plane for Managed Kubernetes now available for everyone			
Cloud Insights Go to Cloud Insights	Innovate faster with insights across your application infrastructure stack			
SaaS Backup Go to SaaS Backup	A secure, encrypted cloud-native offering that safeguards your business-critical Microsoft Office 365 and Salesforce data from corruption, malicious or accidental deletion			
Cloud Backup Service Register for Preview	A fully managed Backup and Restore Service for your Cloud Volumes Service data			

Cloud Central accounts

Each Cloud Manager system is associated with a *NetApp Cloud Central account*. A Cloud Central account provides multi-tenancy and enables you to organize users and resources in isolated workspaces.

A Cloud Central account enables multi-tenancy:


- A single Cloud Central account can include multiple Cloud Manager systems that serve different business needs.

Because users are associated with the Cloud Central account, there's no need to configure users for each individual Cloud Manager system.

- Within each Cloud Manager system, multiple users can deploy and manage Cloud Volumes ONTAP systems in isolated environments called workspaces.

These workspaces are invisible to other users, unless they are shared.

When you deploy Cloud Manager, you select the Cloud Central account to associate with the system:



Account Admins can then modify the settings for this account by managing users, workspaces, and service connectors:

For step-by-step instructions, see [Setting up the Cloud Central account](#).



Cloud Manager needs access to <https://cloudmanager.cloud.netapp.com> in order to connect to the Cloud Central account service. Open this URL on your firewall to ensure that Cloud Manager can contact the service.

Users, workspaces, and service connectors

The Account Settings widget in Cloud Manager enables Account Admins to manage a Cloud Central account. If you just created your account, then you'll start from scratch. But if you've already set up an account, then you'll see *all* the users, workspaces, and service connectors that are associated with the account.

Users

These are NetApp Cloud Central users that you associate with your Cloud Central account. Associating a user with an account and one or more workspaces in that account enables those users to create and manage working environments in Cloud Manager.

When you associate a user, you assign them a role:

- *Account Admin*: Can perform any action in Cloud Manager.
- *Workspace Admin*: Can create and manage resources in the assigned workspace.

Workspaces

In Cloud Manager, a workspace isolates any number of *working environments* from other working environments. Workspace Admins can't access the working environments in a workspace unless the Account Admin associates the admin with that workspace.

A working environment represents a storage system:

- A single-node 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

Service connectors

A service connector is part of Cloud Manager. It runs much of the Cloud Manager software (like the user interface), except for a few Cloud Central services that it connects to (auth0 and Cloud Central accounts). The service connector runs on the virtual machine instance that was deployed in your cloud provider, or on an on-prem host that you configured.

You can use a service connector with more than one NetApp cloud data service. For example, if you already have a service connector for Cloud Manager, you can select it when you set up the Cloud Tiering service.


Examples

The following example shows an account that uses two workspaces to create isolated environments for Cloud Volumes ONTAP systems. For example, one workspace might be for a staging environment, while the other is for a production environment.



Cloud Manager and the Cloud Volumes ONTAP systems don't actually reside *in* the NetApp Cloud Central account—they're running in a cloud provider. This is a conceptual representation of the relationship between each component.

NetApp Cloud Central



Here's another example that shows the highest level of multi-tenancy by using two separate Cloud Central accounts. For example, a service provider might use Cloud Manager in one Cloud Central account to provide services for their customers, while using another account to provide disaster recovery for one of their business units.

Note that account 2 includes two separate service connectors. This might happen if you have systems in separate regions or in separate cloud providers.



Again, Cloud Manager and the Cloud Volumes ONTAP systems don't actually reside *in* the NetApp Cloud Central account—they're running in a cloud provider. This is a conceptual representation of the relationship between each component.

FAQ for integration with Cloud Central accounts

Some time after you upgrade to Cloud Manager 3.7, NetApp will choose specific Cloud Manager systems to integrate with Cloud Central accounts. This FAQ can answer questions that you might have about the process.

How long does the process take?

Just a few minutes.

Will Cloud Manager be unavailable?

No, you can still access your Cloud Manager system.

What about Cloud Volumes ONTAP?

There's no disruption to your Cloud Volumes ONTAP systems.

What happens during this process?

NetApp does the following during the integration process:

1. Creates a new Cloud Central account and associates it with your Cloud Manager system.
2. Assigns new roles to each existing user:
 - Cloud Manager Admins become Account Admins
 - Tenant Admins and Working Environment Admins become Workspace Admins
3. Creates workspaces that replace existing tenants.
4. Places your working environments in those workspaces.
5. Associates the service connector with all workspaces.

Does it matter where I installed my Cloud Manager system?

No. NetApp will integrate systems with Cloud Central accounts no matter where they reside, whether that's in AWS, Azure, or on your premises.

Cloud provider accounts

AWS accounts and permissions


Cloud Manager enables you to choose the AWS account in which you want to deploy a Cloud Volumes ONTAP system. 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 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: [REDACTED] | [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

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 in AWS from the [AWS Marketplace](#) and you can [install Cloud Manager on-premises](#).

If you use the Marketplace, permissions are provided in the same way. You just need to manually create and set up the IAM role, and then provide permissions for any additional accounts.

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

Azure accounts and permissions


Cloud Manager enables you to choose the Azure account in which you want to deploy a Cloud Volumes ONTAP system. 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

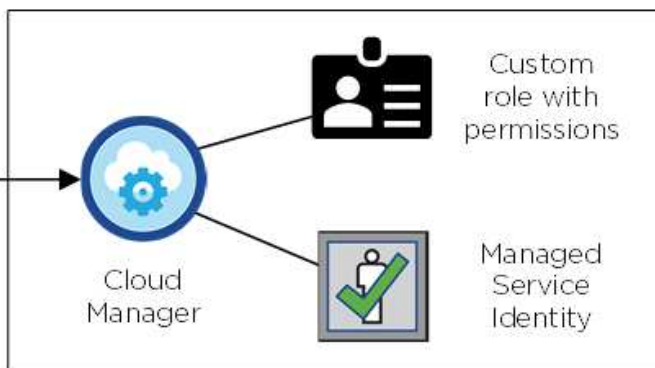
When you deploy Cloud Manager from NetApp Cloud Central, you need to use an Azure account that has permissions to deploy the Cloud Manager virtual machine. The required permissions are listed in the [NetApp Cloud Central policy for Azure](#).

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:

Details & Credentials

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: ...
Dev Keys | Application ID: ...
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 in Azure from the [Azure Marketplace](#), and you can [install Cloud Manager on-premises](#).

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

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

Google Cloud projects, permissions, and accounts

A service account provides Cloud Manager with permissions to deploy and manage Cloud Volumes ONTAP systems in the same project as Cloud Manager, or in different projects. Google Cloud accounts that you add to Cloud Manager are used to enable data tiering.

Project and permissions for Cloud Manager

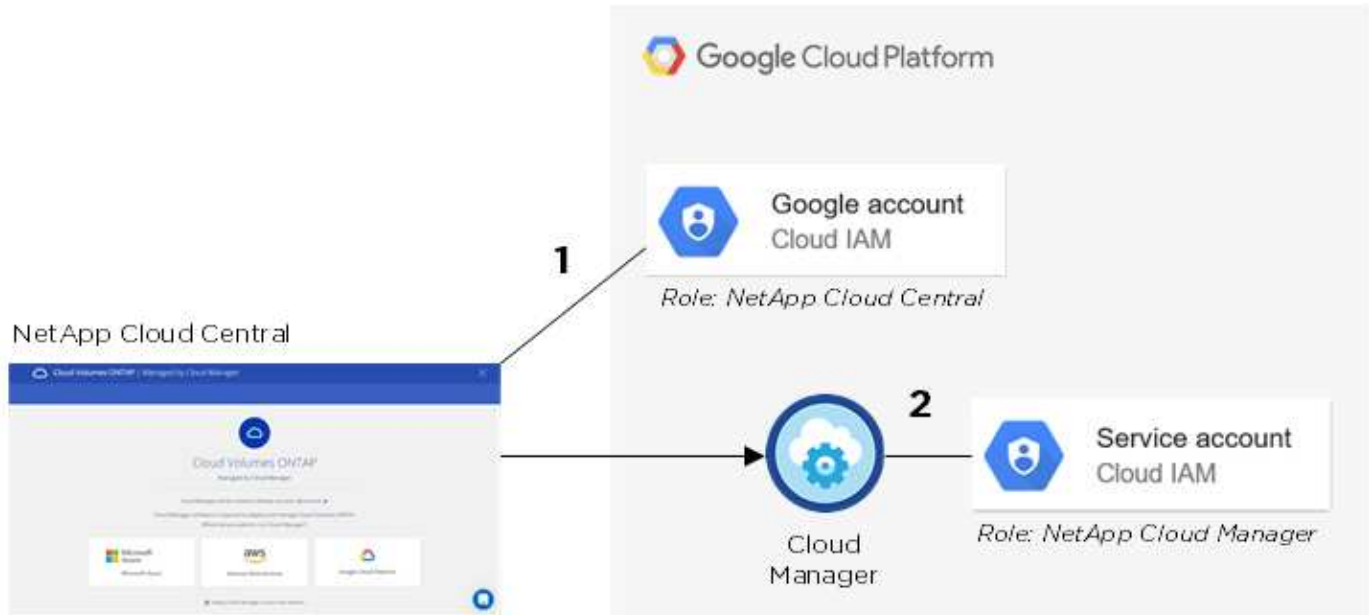
Before you can deploy Cloud Volumes ONTAP in Google Cloud, you must first deploy Cloud Manager in a Google Cloud project. Cloud Manager can't be running on your premises, or in a different cloud provider.

Two sets of permissions must be in place before you deploy Cloud Manager from [NetApp Cloud Central](#):

1. You need to deploy Cloud Manager using a Google account that has permissions to launch the Cloud Manager VM instance from Cloud Central.
2. When deploying Cloud Manager, you are prompted to select a [service account](#) for the VM instance. Cloud Manager gets permissions from the service account to create and manage Cloud Volumes ONTAP systems on your behalf. Permissions are provided by attaching a custom role to the service account.

We have set up two YAML files that include the required permissions for the user and the service account. [Learn how to use the YAML files to set up permissions.](#)

The following image depicts the permission requirements described in numbers 1 and 2 above:



Project for Cloud Volumes ONTAP

Cloud Volumes ONTAP can reside in the same project as Cloud Manager, or in a different project. To deploy Cloud Volumes ONTAP in a different project, you need to first add the Cloud Manager service account and role to that project.

- [Learn how to set up the Cloud Manager service account \(see step 4\).](#)
- [Learn how to deploy Cloud Volumes ONTAP in GCP and select a project.](#)

Account for data tiering

Adding a Google Cloud account to Cloud Manager is required to enable data tiering on a Cloud Volumes ONTAP system. Data tiering automatically tiers cold data to low-cost object storage, enabling you to reclaim space on your primary storage and shrink secondary storage.

When you add the account, you need to provide Cloud Manager with a storage access key for a service account that has Storage Admin permissions. Cloud Manager uses the access keys to set up and manage a Cloud Storage bucket for data tiering.

After you add a Google Cloud account, you can then enable data tiering on individual volumes when you create, modify, or replicate them.

- [Learn how to set up and add GCP accounts to Cloud Manager.](#)
- [Learn how to tier inactive data to low-cost object storage.](#)


Storage

Disks and aggregates

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

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 Cloud Manager creates a 500 GB aggregate, the usable capacity is 442.94 GB.

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. 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 pair an EBS disk with Amazon S3 to [tier inactive data to low-cost object storage](#).

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.

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 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 to [tier inactive data to low-cost object storage](#).

HA pairs

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

Related links

- [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 in Azure](#)

GCP storage

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

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

Related links


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

RAID type

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

Data tiering overview

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.



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



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

Data tiering 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. Changing a system's tiering level enables you to choose a different S3 storage class.

Performance tier

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

Capacity tier

A Cloud Volumes ONTAP system tiers inactive data to a single S3 bucket using the *Standard* storage class. Standard is ideal for frequently accessed data stored across multiple Availability Zones.



Cloud Manager creates a single S3 bucket for each working environment and names it *fabric-pool-cluster unique identifier*. A different S3 bucket is not created for each volume.

Tiering levels

If you don't plan to access the inactive data, you can reduce your storage costs by changing a system's tiering level to one of the following: *Intelligent Tiering*, *One-Zone Infrequent Access*, or *Standard-Infrequent Access*. 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.

The access costs are higher if you do access the data, so take that into consideration before you change the tiering level. [Learn more about Amazon S3 storage classes.](#)

Changing the tiering level is possible after you create the system. For details, see [Tiering inactive data to low-cost object storage.](#)

The tiering level is system wide—it is 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. Changing a system's tiering level enables you to choose a different Azure storage tier.

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 using the Azure *hot* storage tier. The hot tier is ideal for frequently accessed data.



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

Tiering levels

If you don't 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. 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.

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

Changing the tiering level is possible after you create the system. For details, see [Tiering inactive data to low-cost object storage.](#)

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

Data tiering in GCP

When you enable data tiering in GCP, 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 SSDs or HDDs (standard disks).

Capacity tier

A Cloud Volumes ONTAP system tiers inactive data to a single Google Cloud Storage bucket using the *Regional* storage class.



Cloud Manager creates a single bucket for each working environment and names it *fabric-pool-cluster unique identifier*. A different bucket is not created for each volume.

Tiering levels

No other GCP storage classes are supported at this time.

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.

Cloud Manager 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.

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.

When you replicate a volume, you can choose whether to tier the data to object storage. If you do, Cloud Manager applies the **Backup** policy to the data protection volume. Starting with Cloud Volumes ONTAP 9.6, the **All** tiering policy replaces the backup policy.

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 might experience longer cooling times.

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 Account 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 Account 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

The Capacity Management Mode is set to automatic by default. In this mode, 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.

Management of inodes with automatic capacity management

Cloud Manager monitors inode usage on a volume. When 85% of the inodes are used, Cloud Manager increases the size of the volume to increase the number of available inodes. The number of files a volume can contain is determined by how many inodes it has.

Manual capacity management

If the Account 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.

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 i

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.

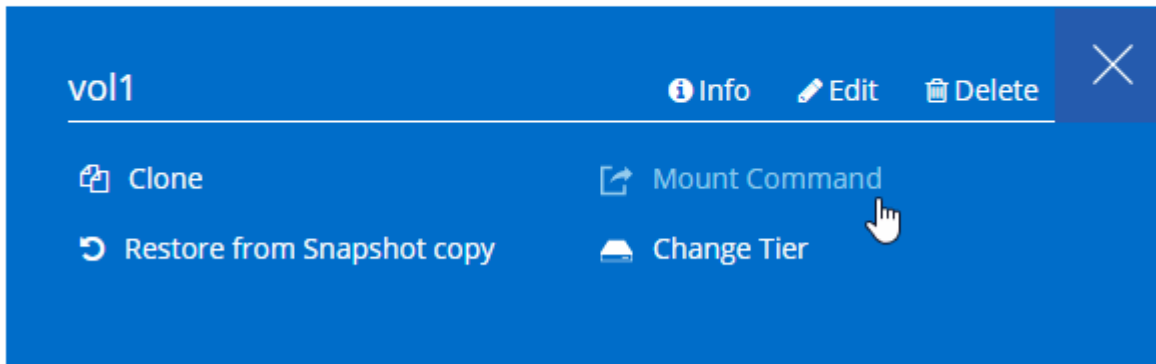


If you move a volume between nodes in an HA pair, you should remount the volume by using the IP address of the other node. Otherwise, you can experience reduced performance. If clients support NFSv4 referrals or folder redirection for CIFS, you can enable those features on the Cloud Volumes ONTAP systems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

You can easily identify the correct IP address from Cloud Manager:

Volumes

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

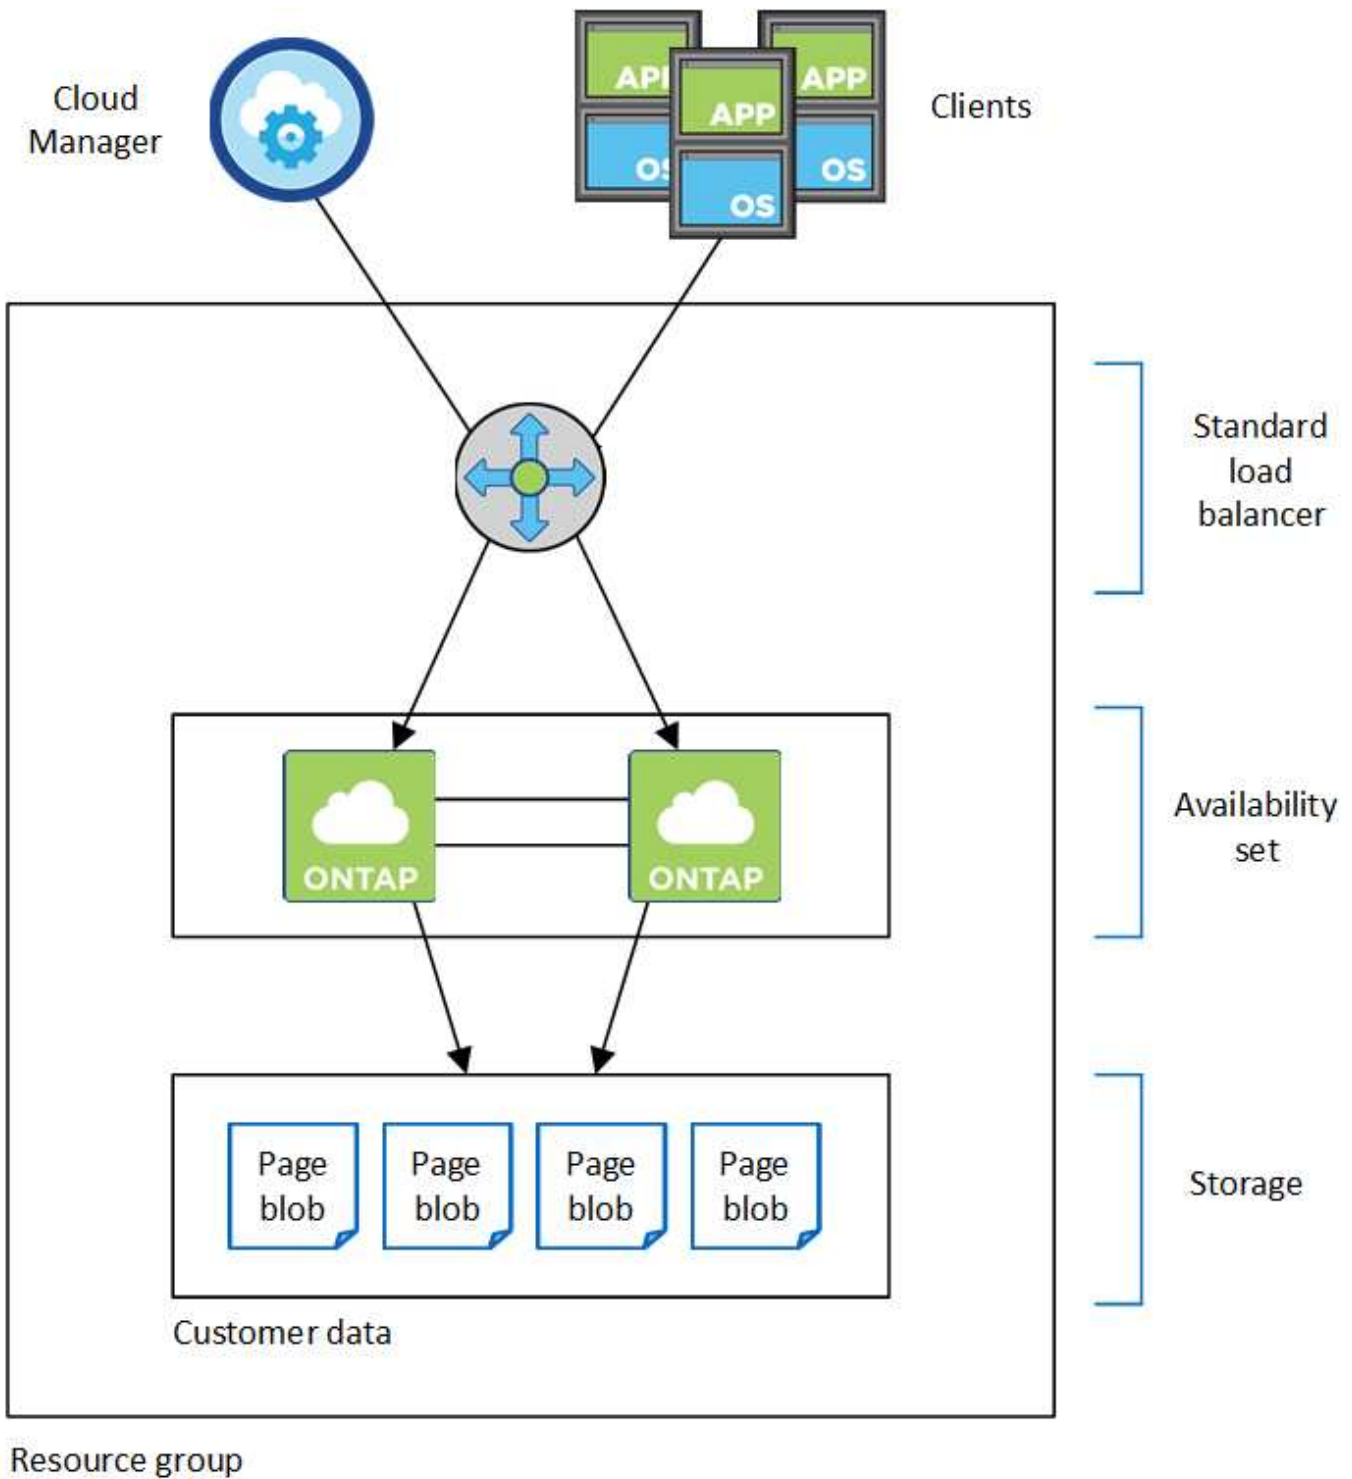


High-availability pairs in Azure

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:



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.

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:

- Two 90 GB Premium SSD disks for the boot volume (one per node)
- Two 140 GB Premium Storage page blobs for the root volume (one per node)
- Two 128 GB Standard HDD disks for saving cores (one per node)

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.

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

- One storage account is required for data tiering to Azure 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 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.
- 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
- Google Cloud Platform default encryption

You can use NetApp Volume Encryption with native AWS, Azure, or GCP 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).



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 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.

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 Cloud Manager 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.

Refer to the [API Developer Guide](#) for details about using the "GcpEncryption" parameters.

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.

1 Enable Snapshot Copy Protection

40 % Protection

3 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.](#)

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 © 2022 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.