



# **Cloud Manager and Cloud Volumes ONTAP documentation**

Cloud Manager 3.8

NetApp  
November 16, 2022

# Table of Contents

- Cloud Manager and Cloud Volumes ONTAP documentation . . . . . 1
  - Discover what's new . . . . . 1
  - Get started . . . . . 1
  - Automate with APIs . . . . . 1
  - Connect with peers, get help, and find more information . . . . . 1
- Release notes . . . . . 2
  - Cloud Manager . . . . . 2
- Important changes in Cloud Manager . . . . . 28
  - SaaS changes . . . . . 28
  - Machine type changes . . . . . 28
  - Account settings . . . . . 28
  - New permissions . . . . . 28
  - New endpoints . . . . . 30
- Get started with Cloud Manager . . . . . 32
  - Learn about Cloud Manager . . . . . 32
  - Networking overview . . . . . 33
  - Signing up to NetApp Cloud Central . . . . . 34
  - Logging in to Cloud Manager . . . . . 35
  - Set up a Cloud Central account . . . . . 35
  - Set up a Connector . . . . . 45
  - Where to go next . . . . . 65
- Manage Cloud Volumes ONTAP . . . . . 67
  - Learn . . . . . 67
  - Get started in AWS . . . . . 93
  - Get started in Azure . . . . . 131
  - Get started in GCP . . . . . 150
  - Provision and manage storage . . . . . 169
  - Replicating data between systems . . . . . 195
  - Monitor performance . . . . . 202
  - Improving protection against ransomware . . . . . 209
  - Administer . . . . . 210
- Provision volumes using a file service . . . . . 233
  - Azure NetApp Files . . . . . 233
  - Cloud Volumes Service for AWS . . . . . 243
  - Cloud Volumes Service for GCP . . . . . 267
- Manage ONTAP clusters . . . . . 283
  - Discovering ONTAP clusters . . . . . 283
  - Managing storage for ONTAP clusters . . . . . 284
- Back up to the cloud . . . . . 287
  - Learn about Backup to Cloud . . . . . 287
  - Get started . . . . . 291
  - Managing backups for Cloud Volumes ONTAP and on-premises ONTAP systems . . . . . 304
- Copy and synchronize data . . . . . 310

Cloud Sync overview	310
Get started	313
Tutorials	343
Managing sync relationships	349
Cloud Sync APIs	353
Cloud Sync technical FAQ	357
Gain insight into data privacy	364
Learn about Cloud Compliance	364
Get started	367
Gaining visibility and control of private data	390
Viewing compliance reports	403
Responding to a Data Subject Access Request	408
Disabling Cloud Compliance	410
Frequently asked questions about Cloud Compliance	411
Enable real-time global file sharing	415
Learn about Global File Cache	415
Before you begin to deploy Global File Cache	418
Getting started	422
Before you begin to deploy Global File Cache Edge instances	433
Deploy Global File Cache Edge instances	439
End-user training	442
Additional information	442
Optimize cloud compute costs	444
Learn about the Compute service	444
Start optimizing your cloud compute costs	445
Tier data to the cloud	448
Learn about Cloud Tiering	448
Get started	452
Set up licensing for Cloud Tiering	471
Managing data tiering from your clusters	473
Cloud Tiering technical FAQ	477
Reference	480
Viewing your Amazon S3 buckets	484
Administer Cloud Manager	486
Finding your Cloud Manager system ID	486
Manage Connectors	486
Manage credentials	500
Managing users, workspaces, Connectors, and subscriptions	524
Managing an HTTPS certificate for secure access	530
Removing Cloud Volumes ONTAP working environments	532
Configuring a Connector to use a proxy server	532
Overriding CIFS locks for Cloud Volumes ONTAP HA in Azure	534
Reference	534
Use APIs and automation	545
Automation resources for infrastructure as code	545

- Where to get help and find more information . . . . . 546
- Earlier versions of Cloud Manager documentation . . . . . 548
- Legal notices . . . . . 549
  - Copyright . . . . . 549
  - Trademarks . . . . . 549
  - Patents . . . . . 549
  - Privacy policy . . . . . 549
  - Open source . . . . . 549



# Cloud Manager and Cloud Volumes ONTAP documentation

Cloud Manager enables IT experts and cloud architects to centrally manage their hybrid multi-cloud infrastructure using NetApp's cloud solutions.

## Discover what's new

- [Important changes in Cloud Manager](#)
- [What's new in Cloud Manager](#)
- [What's new in Cloud Volumes ONTAP](#)

## Get started

- [Cloud Manager](#)
- [Account settings](#)
- [Cloud Volumes ONTAP for AWS](#)
- [Cloud Volumes ONTAP for Azure](#)
- [Cloud Volumes ONTAP for Google Cloud](#)
- [Azure NetApp Files](#)
- [Cloud Volumes Service for AWS](#)
- [Cloud Volumes Service for Google Cloud](#)
- [Cloud Compliance](#)
- [Global File Cache](#)
- [Backup to Cloud](#)
- [Cloud Insights](#)

## Automate with APIs

- [API Developer Guide](#)
- [Automation samples](#)

## Connect with peers, get help, and find more information

- [NetApp Community: Cloud Data Services](#)
- [NetApp Cloud Volumes ONTAP Support](#)
- [Where to get help and find more information](#)

# Release notes

## Cloud Manager

### What's new in Cloud Manager 3.8

Cloud Manager typically introduces a new release every month to bring you new features, enhancements, and bug fixes.



Looking for a previous release?

[What's new in 3.7](#)

[What's new in 3.6](#)

[What's new in 3.5](#)

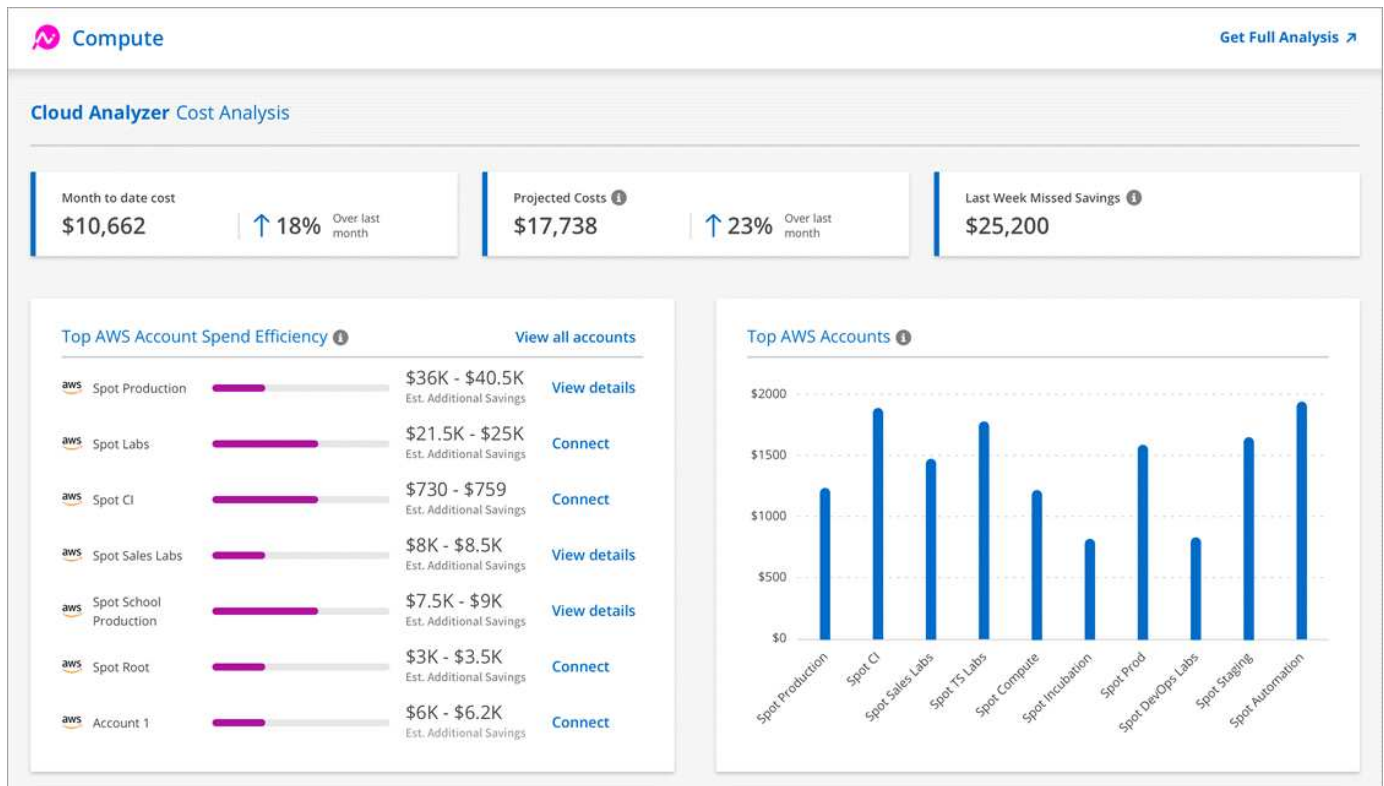
### New Terraform provider (19 Oct 2020)

We have developed a new Terraform provider that DevOps teams can use with Cloud Manager to automate and integrate Cloud Volumes ONTAP with infrastructure as code.

[View the netapp-cloudmanager provider.](#)

### Cloud Manager 3.8.9 update (18 Oct 2020)

By leveraging [Spot's Cloud Analyzer](#), Cloud Manager can now provide a high-level cost analysis of your cloud compute spending and identify potential savings. This information is available from the **Compute** service in Cloud Manager. [Learn more.](#)



## Cloud Manager 3.8.9 update (13 Oct 2020)

We have released two Cloud Tiering updates:

- Licensing for Cloud Tiering is now available from Cloud Manager.

Pay for data tiering from an on-prem ONTAP cluster to the cloud through a pay-as-you-go subscription, an ONTAP tiering license called *FabricPool*, or a combination of both.

- The standalone Cloud Tiering service has been retired. You should now access Cloud Tiering directly from Cloud Manager where all of the same features and functionality are available.

## Cloud Manager 3.8.9 (4 Oct 2020)

- [Cloud Compliance enhancements](#)
- [Cloud Volumes Service for AWS enhancements](#)
- [Cloud Sync integration](#)
- [Account management enhancements](#)
- [Changes for Government regions](#)

### Cloud Compliance enhancements

- A new **Cloud Compliance Viewer** role is available in Cloud Manager.

Users who are assigned this role can only view compliance information and generate reports for workspaces that they have permission to access. They cannot manage Cloud Compliance settings and they cannot access any other Cloud Manager features and services. This may be the perfect role for your legal team so they can monitor Cloud Compliance scan results. See [user roles](#) for details.

- Added support to scan MongoDB and PostgreSQL database schemas. See [scanning database schemas](#) for more information.
- Cloud Compliance pricing is changing as of October 7th.

The first 1 TB of data that Cloud Compliance scans in a Cloud Manager workspace is free. This includes data from Cloud Volumes ONTAP volumes, Azure NetApp Files volumes, Amazon S3 buckets, and database schemas. A subscription is required to scan any additional data after you reach 1 TB. See [pricing](#) for details.

### Cloud Volumes Service for AWS enhancements

When creating a new volume, you can choose to base that volume on an existing Snapshot copy of another volume.

### Cloud Sync integration

NetApp's Cloud Sync service is now available from within Cloud Manager. Cloud Sync offers a simple, secure, and automated way to migrate your data from any source destination to any target destination, in the cloud or on your premises. [Learn more](#).

### Account management enhancements

We've added more ways to manage your account.

- An overview of your account resources is now available.


You can quickly view the number of users, workspaces, Connectors, and subscriptions in your account.

- You can change the name of your account.
- You can copy your account ID, workspace ID, or Connector ID.

Copying these IDs will help with automation features that we're planning.

- You can disable use of the SaaS platform.

We don't recommend disabling the SaaS platform unless you need to in order to comply with your company's security policies. Disabling the SaaS platform limits your ability to use NetApp's integrated cloud services. [Learn more.](#)



### Changes for Government regions

If you deploy a Connector in an AWS GovCloud region, an Azure Gov region, or an Azure DoD region, access to Cloud Manager is now available only through a Connector's host IP address. Access to the SaaS platform is disabled for the entire account.

This means that only privileged users who can access the end-user internal VPC/VNet can use Cloud Manager's UI or API.

[Learn more about this limitation.](#)

### Cloud Manager 3.8.8 update (22 Sept 2020)

We enhanced the Kubernetes service to make it easier to use and to provide additional capabilities:

- We've made it easier to discover the Kubernetes clusters running in your cloud provider's managed Kubernetes service.

Just click **Discover Clusters** and Cloud Manager will discover your managed clusters using the cloud provider permissions that you've already provided.

- You can now view more information about a discovered Kubernetes cluster, including its state, the number of volumes, storage classes, and more.
- We've added resource and error checking to ensure that communication is available between the cluster and Cloud Volumes ONTAP. And if it isn't, then we'll let you know.

[Learn how to get started.](#)

Note that the service account for a Connector requires the following permissions to discover and manage Kubernetes clusters running in Google Kubernetes Engine (GKE):

```
- container.*
```

### Cloud Manager 3.8.8 update (10 Sept 2020)

The following enhancements are available when deploying Global File Cache through Cloud Manager:

- A Cloud Volumes ONTAP HA pair in AWS is now supported as the backend storage platform for your central storage.
- Multiple Global File Cache Core instances can be deployed in a Load Distributed design.

[Learn more about Global File Cache.](#)

### Cloud Manager 3.8.8 (9 Sept 2020)

- [Support for Cloud Volumes Service for Google Cloud](#)
- [Backup to Cloud now supports on-premises ONTAP clusters](#)
- [Backup to Cloud enhancements](#)
- [Cloud Compliance enhancements](#)
- [Refreshed navigation](#)
- [Administration improvements](#)

### Support for Cloud Volumes Service for Google Cloud

- Add a working environment to manage existing Cloud Volumes Service for GCP volumes and to create new volumes. [Learn how.](#)
- Create and manage NFSv3 and NFSv4.1 volumes for Linux and UNIX clients, and SMB 3.x volumes for Windows clients.
- Create, delete, and restore volume snapshots.

### Backup to Cloud now supports on-premises ONTAP clusters

Start backing up data from your on-premises ONTAP systems to the cloud. Enable Backup to Cloud on your on-prem working environments to back up volumes to Azure Blob storage. [Learn more.](#)

## Backup to Cloud enhancements

We revised the user interface for better usability:

- Volume list page to easily see the volumes being backed up along with the available backups
- Backup settings page to view backup settings for each working environment

## Cloud Compliance enhancements

- Ability to scan data from databases

Scan your databases to identify the personal and sensitive data that resides in each schema. Supported databases include Oracle, SAP HANA, and SQL Server (MSSQL). [Learn more about scanning databases.](#)

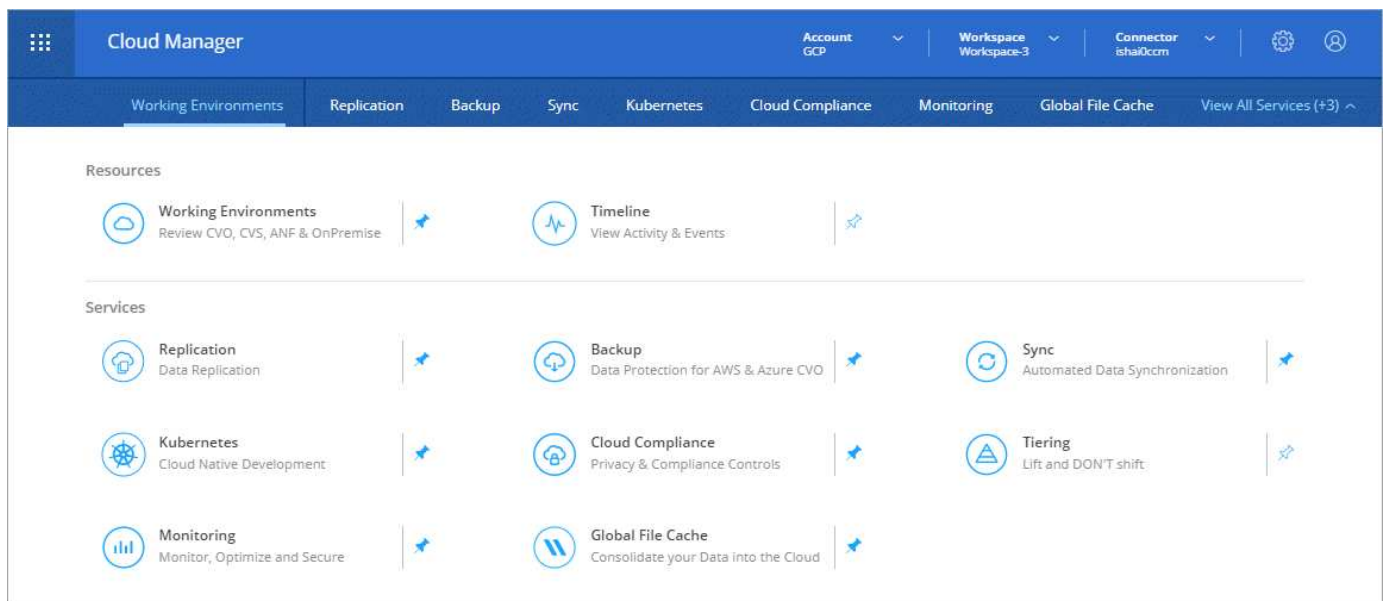
- Ability to scan data protection (DP) volumes

DP volumes are destination volumes from SnapMirror operations typically from on-premises ONTAP clusters. Now you can easily identify the personal and sensitive data that resides in those on-prem files. [See how.](#)

## Refreshed navigation

We've refreshed the header in Cloud Manager to make it easier for you to navigate between NetApp cloud services.


Click **View All Services** and you can pin and unpin the services that you want to see in the navigation.



As you can see, we've also refreshed the Account, Workspace, and Connector drop-downs, so it's easier to view your current selections.

## Administration improvements

- You can now remove inactive Connectors from Cloud Manager. [Learn how.](#)



- You can now replace the Marketplace subscription that's currently associated with your cloud provider credentials. If you ever need to change how you're charged, this change can help you ensure that you're being charged through the right Marketplace subscription.

Learn how [in AWS](#), [in Azure](#), and [in GCP](#).

### Update on required Azure permissions (6 Aug 2020)

To avoid Azure deployment failures, make sure that your Cloud Manager policy in Azure includes the following permission:

```
"Microsoft.Resources/deployments/operationStatuses/read"
```

Azure now requires this permission for some virtual machine deployments (it depends on the underlying physical hardware that's used during deployment).

[View the latest Cloud Manager policy for Azure.](#)

### Cloud Manager 3.8.7 (3 Aug 2020)

- [New software-as-a-service experience](#)
- [Cloud Volumes ONTAP enhancements](#)
- [Azure NetApp Files enhancements](#)
- [Cloud Volumes Service for AWS enhancements](#)
- [Cloud Compliance enhancements](#)
- [Backup to Cloud enhancements](#)
- [Support for Global File Cache](#)

#### New software-as-a-service experience

We have fully introduced a software-as-a-service experience for Cloud Manager. This new experience makes it easier for you to use Cloud Manager and enables us to provide additional features to manage your hybrid cloud infrastructure.

Cloud Manager includes a [SaaS-based interface](#) that is integrated with NetApp Cloud Central, and Connectors that enable Cloud Manager to manage resources and processes within your public cloud environment. (The

Connector is actually the same as the existing Cloud Manager software that you have installed.)



A Connector is required in most cases, but it's not required to use Azure NetApp Files, Cloud Volumes Service, or Cloud Sync from Cloud Manager.

As previously mentioned in these release notes, you'll need to upgrade the machine type for your Connectors to access the new capabilities that we're offering. Cloud Manager will prompt you with instructions to change the machine type. [Learn more.](#)

### Cloud Volumes ONTAP enhancements

Two enhancements are available for Cloud Volumes ONTAP.

- **Multiple BYOL licenses to allocate additional capacity**

You can now purchase multiple licenses for a Cloud Volumes ONTAP BYOL system to allocate more than 368 TB of capacity. For example, you might purchase two licenses to allocate up to 736 TB of capacity to Cloud Volumes ONTAP. Or you could purchase four licenses to get up to 1.4 PB.

The number of licenses that you can purchase for a single node system or HA pair is unlimited.

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

[Learn how to add a new system license.](#)

- **Encrypt Azure managed disks using external keys**

You can now encrypt Azure managed disks on single node Cloud Volumes ONTAP systems using external keys from another account. This feature is supported using APIs.

You just need to add the following to the API request when creating the single node system:

```
"azureEncryptionParameters": {  
  "key": <azure id of encryptionset>  
}
```

This feature requires new permissions as shown in the latest [Cloud Manager policy for Azure](#).

```
"Microsoft.Compute/diskEncryptionSets/read"
```

### Azure NetApp Files enhancements

This release includes several enhancements in support of Azure NetApp Files.

- **Azure NetApp Files setup**

You can now set up and manage Azure NetApp Files directly from Cloud Manager. [Learn how.](#)



- **New protocol support**

You can now create NFSv4.1 volumes and SMB volumes.

- **Capacity pool and volume snapshot management**

Cloud Manager enables you to create, delete, and restore volume snapshots. You can also create new capacity pools and specify their service levels.

- **Ability to edit volumes**

You can edit a volume by changing its size and managing tags.

#### Cloud Volumes Service for AWS enhancements

There are many enhancements in Cloud Manager in support of Cloud Volumes Service for AWS.

- **New protocol support**

Now you can create NFSv4.1 volumes, SMB volumes, and dual protocol volumes. Previously you could only create and discover NFSv3 volumes within Cloud Manager.

- **Snapshot support**

You can create snapshot policies to automate the creation of volume snapshots, create an on-demand snapshot, restore a volume from a snapshot, create a new volume based on an existing snapshot, and more. See [Managing cloud volumes snapshots](#) for more information.

- **Create the initial volume in a region from Cloud Manager**

Before this release the first volume in each region had to be created in the Cloud Volumes Service for AWS interface. Now you can subscribe to [one of the NetApp Cloud Volumes Service offerings on the AWS Marketplace](#) and then create the first volume from Cloud Manager.

#### Cloud Compliance enhancements

The following enhancements are now available for Cloud Compliance.

- **Revised deployment process for your Cloud Compliance instance**

The Cloud Compliance instance is set up and deployed using a new wizard in Cloud Manager. After deployment is complete you enable the service for each working environment you want to scan.

- **Ability to select the volumes to be scanned within a working environment**

Now you can enable and disable scanning for individual volumes in a Cloud Volumes ONTAP or Azure NetApp Files working environment. If you don't need to scan certain volumes for compliance, turn them off.

[Learn more about disabling scanning for volumes.](#)

- **Navigation tabs to quickly jump to your area of interest**

New tabs for Dashboard, Investigation, and Configuration enable you to get to these sections more easily.

- **HIPAA Report**

A new Health Insurance Portability and Accountability Act (HIPAA) Report is now available. This report is designed to aid in your organization's requirement to comply with HIPAA data privacy laws.

[Learn more about the HIPAA report.](#)

- **New sensitive personal data type**

Cloud Compliance can now find ICD-9-CM Medical Codes in files.

- **New personal data type**

Cloud Compliance can now find two new national identifiers in files: Croatian ID (OIB) and Greek ID.

### **Backup to Cloud enhancements**

The following enhancements are now available for Backup to Cloud.

- **Bring Your Own License (BYOL) is now available**

Backup to Cloud has been available only with a Pay As You Go (PAYGO) license. A BYOL license allows you to purchase a license from NetApp to use Backup to Cloud for a certain period of time and for a maximum amount backup space. When either limit is reached you will need to renew the license.

[Learn more about the new Backup to Cloud BYOL license.](#)

- **Support for data protection (DP) volumes**

Data protection volumes can be backed up and restored now.

### **Support for Global File Cache**

NetApp Global File Cache enables you to consolidate silos of distributed file servers into one cohesive global storage footprint in the public cloud. This creates a globally accessible file system in the cloud that all distributed locations can use as if they were local.

Starting with this release, the Global File Cache Management instance and Core instance can be deployed and managed through Cloud Manager. This saves many hours during your initial deployment process and provides a single pane of glass through Cloud Manager for this and other deployed systems. Global File Cache Edge instances are still deployed locally at your remote offices.

See [Global File Cache overview](#) for more information.

The initial configuration that can be deployed using Cloud Manager must meet the following requirements. Other configurations like Cloud Volumes Service, Azure NetApp Files, and Cloud Volumes Service for AWS and GCP continue to be deployed using the legacy procedures. [Learn more.](#)

- The backend storage platform that is used as your central storage must be a working environment where you have deployed a Cloud Volumes ONTAP HA pair in Azure.

Other storage platforms and other cloud providers are not supported at this time using Cloud Manager but can be deployed using legacy deployment procedures.

- The GFC Core can be deployed only as a stand-alone instance.

If you need to use a Load Distributed design that includes multiple Core instances you must use the legacy

procedures.

This feature requires new permissions as shown in the latest [Cloud Manager policy for Azure](#).

```
"Microsoft.Resources/deployments/operationStatuses/read",  
"Microsoft.Insights/Metrics/Read",  
"Microsoft.Compute/virtualMachines/extensions/write",  
"Microsoft.Compute/virtualMachines/extensions/read",  
"Microsoft.Compute/virtualMachines/extensions/delete",  
"Microsoft.Compute/virtualMachines/delete",  
"Microsoft.Network/networkInterfaces/delete",  
"Microsoft.Network/networkSecurityGroups/delete",  
"Microsoft.Resources/deployments/delete",
```

### Improved experience requires stronger machine type (15 July 2020)

As we improve the Cloud Manager experience, you'll need to upgrade your machine type to access the new capabilities that we'll be offering. The improvements will include a [software-as-a-service experience for Cloud Manager](#) and new and enhanced cloud service integrations.

Cloud Manager will prompt you with instructions to change the machine type.

Here are some details:

1. To ensure adequate resources are available for proper functionality of the new features in Cloud Manager, we've changed the default instance, VM, and machine type as follows:
  - AWS: t3.xlarge
  - Azure: DS3 v2
  - GCP: n1-standard-4

These default sizes are the minimum supported [based on CPU and RAM requirements](#).

2. As part of this transition, Cloud Manager requires access to the following endpoint so it can obtain software images of container components for a Docker infrastructure:

<https://cloudmanagerinfraprod.azurecr.io>

Ensure that your firewall enables access to this endpoint from Cloud Manager.

### Cloud Manager 3.8.6 (6 July 2020)

- [Support for iSCSI volumes](#)
- [Support for the All tiering policy](#)

#### Support for iSCSI volumes

Cloud Manager now enables you to create iSCSI volumes for Cloud Volumes ONTAP and on-prem ONTAP clusters directly from the user interface.

When you create an iSCSI volume, Cloud Manager automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, [use the IQN to connect to the LUN from your hosts](#).



You can create additional LUNs from System Manager or the CLI.

### Support for the All tiering policy

You can now choose the All tiering policy when you create or modify a volume for Cloud Volumes ONTAP. When you use the All tiering policy, data is immediately marked as cold and tiered to object storage as soon as possible. [Learn more about data tiering](#).

### Cloud Manager transition to SaaS (22 June 2020)

We're introducing a software-as-a-service experience for Cloud Manager. This new experience makes it easier for you to use Cloud Manager and enables us to provide additional features to manage your hybrid cloud infrastructure. [Learn more](#).

### Cloud Manager 3.8.5 (31 May 2020)

- [New subscription required in the Azure Marketplace](#)
- [Backup to Cloud enhancements](#)
- [Cloud Compliance enhancements](#)

### New subscription required in the Azure Marketplace

A new subscription is available in the Azure Marketplace. This one-time subscription is required to deploy Cloud Volumes ONTAP 9.7 PAYGO (except for your 30-day free trial system). The subscription also enables us to offer add-on features for Cloud Volumes ONTAP PAYGO and BYOL. You'll be charged from this subscription for every Cloud Volumes ONTAP PAYGO system that you create and each add-on feature that you enable.

Cloud Manager will prompt you to subscribe to this offering when you deploy a new Cloud Volumes ONTAP system (9.7 P1 or later).

The screenshot shows the 'Details & Credentials' page for a Cloud Manager system. At the top, there are three sections: 'MyAzureCredentials' (Credentials), 'AzureSubscription1222aaaa' (Azure Subscription), and 'Marketplace Subscription'. The 'Marketplace Subscription' section displays a yellow warning icon and the text 'No subscription is associated'. A red arrow points to this text. To the right of this section is an 'Edit Credentials' button. Below these sections are two columns: 'Details' and 'Credentials'. The 'Details' column contains 'Working Environment Name (Cluster Name)' (input field), 'Resource Group Name' (input field with a dropdown showing '[Working Environment Name]-rg') and a checked 'Use Default' checkbox. The 'Credentials' column contains 'User Name' (input field) and 'Password' (input field).

## Backup to Cloud enhancements

The following enhancements are now available for Backup to Cloud.

- In Azure, you can now create a new resource group or select an existing resource group instead of having Cloud Manager create one for you. The resource group can't be changed after you enable Backup to Cloud.
- In AWS, you can now back up Cloud Volumes ONTAP instances that reside on a different AWS account than your Cloud Manager AWS account.
- Additional options are now available when selecting the backup schedule for volumes. In addition to daily, weekly, and monthly backup options, you can now select one of the system-defined policies that provide combination policies such as 30 daily, 13 weekly, and 12 monthly backups.
- After deleting all backups for a volume, you can now start creating backups again for that volume. This was a known limitation in the previous release.

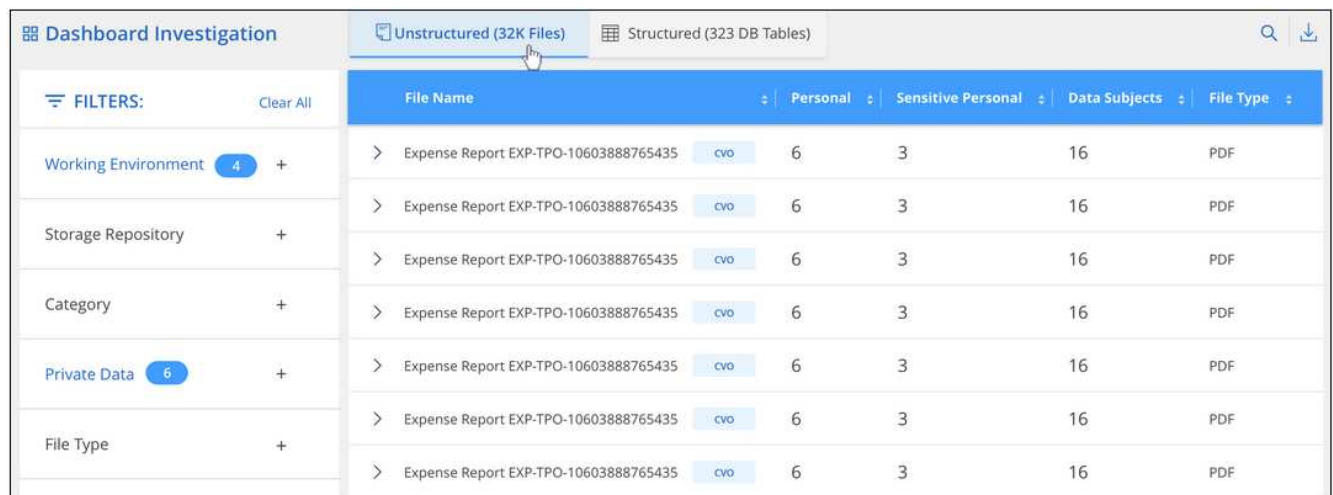
## Cloud Compliance enhancements

The following enhancements are available for Cloud Compliance.

- You can now scan S3 buckets that are in different AWS accounts than the Cloud Compliance instance. You just need to create a role on that new account so that the existing Cloud Compliance instance can connect to those buckets. [Learn more](#).

If you configured Cloud Compliance before release 3.8.5, you will need to modify the existing [IAM role for the Cloud Compliance instance](#) to use this functionality.

- You can now filter the contents of the investigation page to display only the results you want to see. Filters include working environment, category, private data, file type, last modified date, and whether the S3 object's permissions are open to public access.



The screenshot shows the 'Dashboard Investigation' interface. At the top, there are two tabs: 'Unstructured (32K Files)' and 'Structured (323 DB Tables)'. Below the tabs is a search bar and a download icon. On the left side, there is a 'FILTERS' section with a 'Clear All' button. The filters include: 'Working Environment' (4 items), 'Storage Repository', 'Category', 'Private Data' (6 items), and 'File Type'. The main table displays a list of files with columns for 'File Name', 'Personal', 'Sensitive Personal', 'Data Subjects', and 'File Type'. The table contains six rows of data, all showing 'Expense Report EXP-TPO-10603888765435' as the file name and 'PDF' as the file type. The 'Personal' column shows 6 items, 'Sensitive Personal' shows 3 items, and 'Data Subjects' shows 16 items for each row.

File Name	Personal	Sensitive Personal	Data Subjects	File Type
> Expense Report EXP-TPO-10603888765435	6	3	16	PDF
> Expense Report EXP-TPO-10603888765435	6	3	16	PDF
> Expense Report EXP-TPO-10603888765435	6	3	16	PDF
> Expense Report EXP-TPO-10603888765435	6	3	16	PDF
> Expense Report EXP-TPO-10603888765435	6	3	16	PDF
> Expense Report EXP-TPO-10603888765435	6	3	16	PDF

- You can now activate and deactivate Cloud Compliance on a working environment directly from the Cloud Compliance tab.

## Cloud Manager 3.8.4 update (10 May 2020)

We released an enhancement to Cloud Manager 3.8.4.

## Cloud Insights integration

By leveraging NetApp's Cloud Insights service, Cloud Manager gives you insights into the health and performance of your Cloud Volumes ONTAP instances and helps you troubleshoot and optimize the performance of your cloud storage environment. [Learn more.](#)

## Cloud Manager 3.8.4 (3 May 2020)

Cloud Manager 3.8.4 includes the following improvement.

### Backup to Cloud enhancements

The following enhancements are now available for Backup to Cloud (previously called *Backup to S3* for AWS):

- **Backing up to Azure Blob storage**

Backup to Cloud is now available for Cloud Volumes ONTAP in Azure. Backup to Cloud provides backup and restore capabilities for protection, and long-term archive of your cloud data. [Learn more.](#)

- **Deleting backups**

You can now delete all backups for a specific volume directly from the Cloud Manager interface. [Learn more.](#)

## Cloud Manager 3.8.3 (5 Apr 2020)

- [Cloud Tiering integration](#)
- [Data migration to Azure NetApp Files](#)
- [Cloud Compliance enhancements](#)
- [Backup to S3 enhancements](#)
- [iSCSI volumes using APIs](#)

### Cloud Tiering integration

NetApp's Cloud Tiering service is now available from within Cloud Manager. Cloud Tiering enables you to tier data from an on-premises ONTAP cluster to lower-cost object storage in the cloud. This frees up high-performance storage space on the cluster for more workloads.

[Learn more.](#)

### Data migration to Azure NetApp Files

You can now migrate NFS or SMB data to Azure NetApp Files directly from Cloud Manager. Data syncs are powered by NetApp's Cloud Sync service.

[Learn how to migrate data to Azure NetApp Files.](#)

### Cloud Compliance enhancements

The following enhancements are now available for Cloud Compliance.

- **30-day free trial for Amazon S3**

A 30-day free trial is now available to scan Amazon S3 data with Cloud Compliance. If you previously enabled Cloud Compliance on Amazon S3, your 30-day free trial is active starting today (5 Apr 2020).

A subscription to the AWS Marketplace is required to continue scanning Amazon S3 after the free trial ends. [Learn how to subscribe.](#)

[Learn about pricing to scan Amazon S3.](#)

- **New personal data type**

Cloud Compliance can now find a new national identifier in files: Brazilian ID (CPF).

[Learn more about personal data types.](#)

- **Support for additional metadata categories**

Cloud Compliance can now categorize your data into nine additional metadata categories. [See the full list of supported metadata categories.](#)

### Backup to S3 enhancements

The following enhancements are now available for the Backup to S3 service.

- **S3 lifecycle policy for backups**

Backups start in the *Standard* storage class and transition to the *Standard-Infrequent Access* storage class after 30 days.

- **Deleting backups**

You can now delete backups using a Cloud Manager API. [Learn more.](#)

- **Block public access**

Cloud Manager now enables the [Amazon S3 Block Public Access feature](#) on the S3 bucket where backups are stored.

### iSCSI volumes using APIs

The Cloud Manager APIs now enable you to create iSCSI volumes. [View an example here.](#)

### Cloud Manager 3.8.2 (1 Mar 2020)

- [Amazon S3 working environments](#)
- [Cloud Compliance enhancements](#)
- [NFS version for volumes](#)
- [Support for Azure US Gov regions](#)

### Amazon S3 working environments

Cloud Manager now automatically discovers information about the Amazon S3 buckets that reside in the AWS account where it's installed. This enables you to easily see details about your S3 buckets, including the region, access level, storage class, and whether the bucket is used with Cloud Volumes ONTAP for backups or data

tiering. And you can scan the S3 buckets with Cloud Compliance, as described below.

The screenshot shows the Amazon S3 console interface. At the top, there's a header for 'Amazon S3'. Below it, there are two summary boxes: 'S3 Information' showing 242 Total Buckets and 15 Regions, and 'Number of buckets with active services' showing 144 Backup Targets and 23 Tiering Targets. Below these is a pagination control for '1 - 50 of 242' items. The main part of the screenshot is a table listing buckets with columns for Bucket Name, Region, Backup, Tiering, Access, and Storage Class. Three buckets are visible: 'appsinstall', 'automationbucketeran', and 'aws-athena-query-results-64...'. The 'Access' column shows 'Objects can be public' for the first and third buckets, and 'Public' for the second.

Bucket Name	Region	Backup	Tiering	Access	Storage Class
appsinstall	US West (Oregon)			Objects can be public	normal
automationbucketeran	US West (Oregon)			Public	normal
aws-athena-query-results-64...	US West (Oregon)			Objects can be public	normal

### Cloud Compliance enhancements

The following enhancements are now available for Cloud Compliance.

- **Support for Amazon S3**

Cloud Compliance can now scan your Amazon S3 buckets to identify the personal and sensitive data that resides in S3 object storage. Cloud Compliance can scan any bucket in the account, regardless if it was created for a NetApp solution.

[Learn how to get started.](#)

- **Investigation page**

A new Investigation page is now available for each type of personal file, sensitive personal file, category, and file type. The page shows details about the affected files and enables you to sort by the files that include the most personal data, sensitive personal data, and names of data subjects. This page replaces the CSV report that was previously available.

Here's a sample:

The screenshot shows the Cloud Compliance investigation page. At the top, there's a header for 'Cloud Compliance'. Below it, there's a navigation bar with a '< Back' button. The main heading is 'Dashboard Investigation for 'German Tax Identification Number (Steuerliche Identifikationsnummer)'' with search and download icons. Below the heading, it says '1034 results found in 3 Working Environments'. The main part of the screenshot is a table with columns for File Name, Personal, Sensitive Personal, Data Subjects, and File Type. Three rows are visible, all with 'Expense Report EXP-TPO-1060388' as the file name and 'PDF' as the file type. The counts for Personal, Sensitive Personal, and Data Subjects are 6, 3, 16 for the first row; 9, 2, 11 for the second row; and 4, 1, 7 for the third row.

File Name	Personal	Sensitive Personal	Data Subjects	File Type
> Expense Report EXP-TPO-1060388	6	3	16	PDF
> Expense Report EXP-TPO-1060388	9	2	11	PDF
> Expense Report EXP-TPO-1060388	4	1	7	PDF



[Learn more about the Investigation page.](#)

- **PCI DSS Report**

A new Payment Card Industry Data Security Standard (PCI DSS) Report is now available. This report can help you identify the distribution of credit card information across your files. You can view how many files contain credit card information, whether the working environments are protected by encryption or ransomware protection, retention details, and more.

[Learn more about the PCI DSS report.](#)

- **New sensitive personal data type**

Cloud Compliance can now find ICD-10-CM Medical Codes, which are used in the medical and health industry.

### NFS version for volumes

You can now select the NFS version to enable on a volume when you create or edit a volume for Cloud Volumes ONTAP.

**Volume Details, Protection & Protocol**

**Details & Protection**

Volume Name:  Size (GB):

Snapshot Policy:

*i* Default Policy

**Protocol**

NFS Protocol  CIFS Protocol

Access Control:

Custom export policy *i*

**Advanced options** ^

Select NFS Version:  NFSv3  NFSv4

### Support for Azure US Gov regions

Cloud Volumes ONTAP HA pairs are now supported in Azure US Gov regions.

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

### Cloud Manager 3.8.1 update (16 Feb 2020)

We released a few enhancements to Cloud Manager 3.8.1.

#### Backup to S3 enhancements

- Backup copies are now stored in an S3 bucket that Cloud Manager creates in your AWS account, with one bucket per Cloud Volumes ONTAP working environment.

- Backup to S3 is now supported in all AWS regions [where Cloud Volumes ONTAP is supported](#).
- You can set the backup schedule to daily, weekly, or monthly.
- Cloud Manager no longer needs to set up *private links* to the Backup to S3 service.

Additional S3 permissions are required for these enhancements. The IAM role that provides Cloud Manager with permissions must include permissions from the latest [Cloud Manager policy](#).

[Learn more about Backup to S3.](#)

### AWS updates

We've introduced support for new EC2 instances and a change in the number of supported data disks for Cloud Volumes ONTAP 9.6 and 9.7. Check out the changes in the Cloud Volumes ONTAP Release Notes.

- [Cloud Volumes ONTAP 9.7 Release Notes](#)
- [Cloud Volumes ONTAP 9.6 Release Notes](#)

### Cloud Manager 3.8.1 (2 Feb 2020)

- [Cloud Compliance enhancements](#)
- [Enhancements to accounts and subscriptions](#)
- [Timeline enhancements](#)

### Cloud Compliance enhancements

The following enhancements are now available for Cloud Compliance.

- **Support for Azure NetApp Files**

We're pleased to announce that Cloud Compliance can now scan Azure NetApp Files to identify personal and sensitive data that resides on volumes.

[Learn how to get started.](#)

- **Scan status**

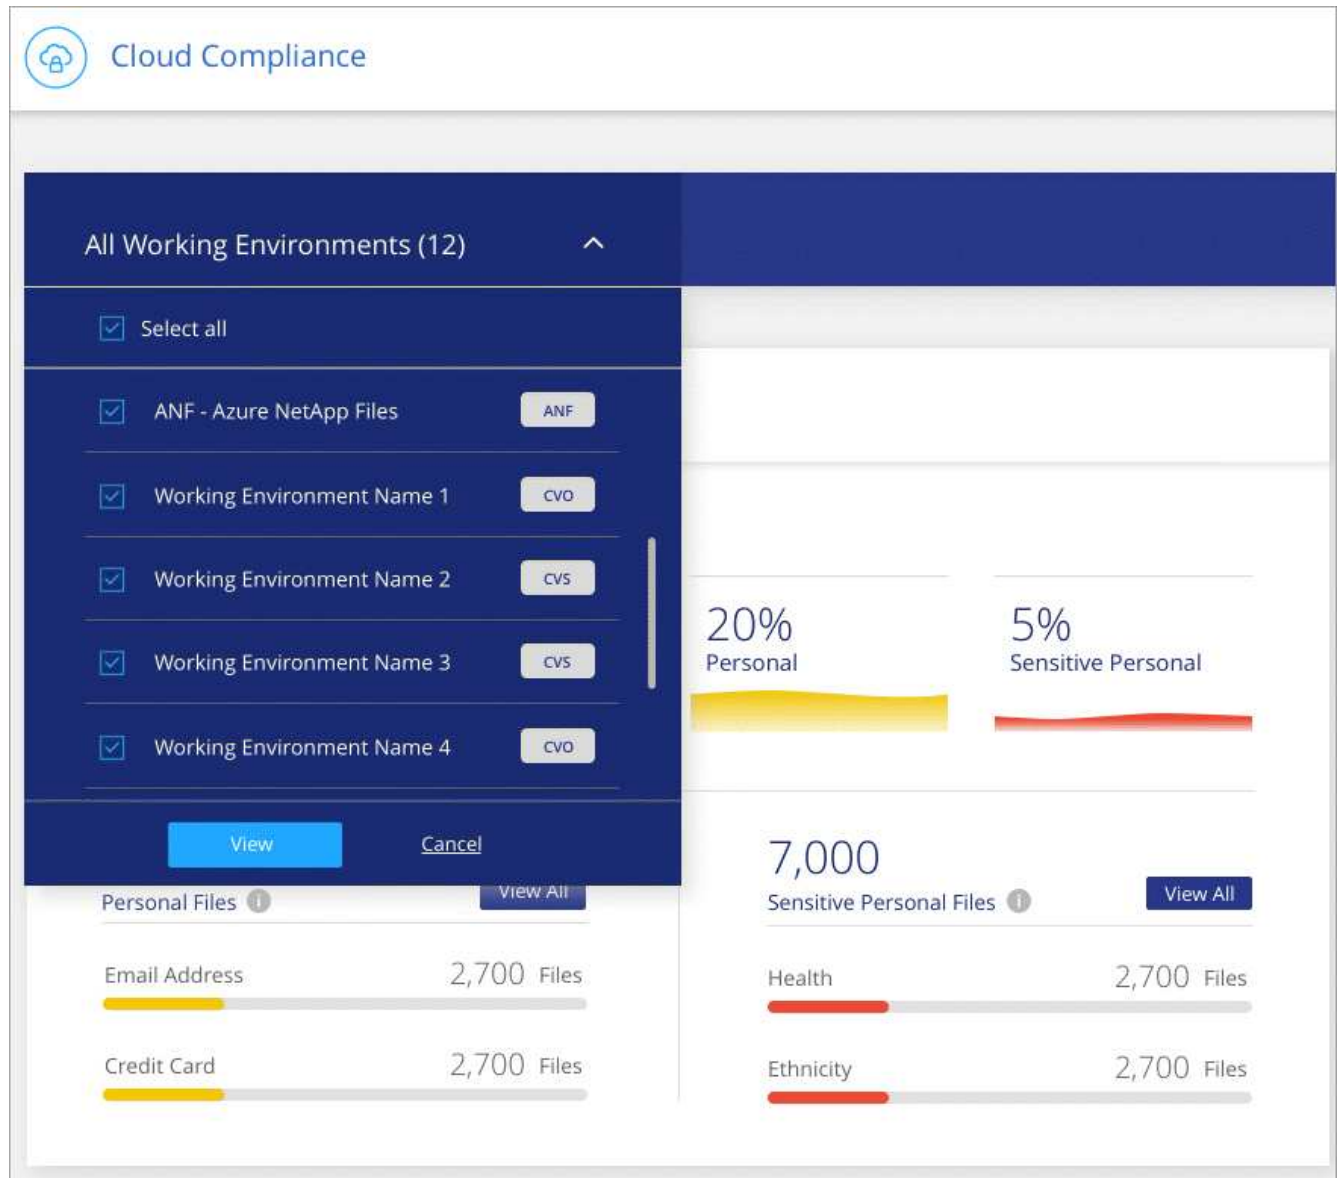
Cloud Compliance now shows you a scan status for each CIFS and NFS volume, including error messages that you can use to correct any issues.

The screenshot shows a mobile interface for 'Volumes Scan Status for cognigoWE'. It displays '2 Volumes found' and a table with the following data:

Name ↑↑	Protocol ↓↑	Status ↓↑	Details ↓↑
\\172.31.134.172\cifs_vol_share	CIFS	Not Scanning	The CIFS credentials that you provided don't have sufficient per...
172.31.134.172:/parallel_tests	NFS	Continuously Scanning	

- **Filter dashboard by working environment**

You can now filter the contents of the Cloud Compliance dashboard to see compliance data for specific working environments.



- **New personal data type**

Cloud Compliance can now identify a California Driver's License when scanning data.

- **Support for additional categories**

Three additional categories are supported: Application data, logs, and database and index files.

[Learn more about categories.](#)

### Enhancements to accounts and subscriptions

We've made it easier to select an AWS account or GCP project and an associated marketplace subscription for a pay-as-you-go Cloud Volumes ONTAP system. These enhancements help to ensure that you're paying from the right account or project.

For example, when you create a system in AWS, click **Edit Credentials** if you don't want to use the default account and subscription:

### Details & Credentials

Instance Profile Credentials	Account ID	QA Subscription Marketplace Subscription	<a href="#" style="border: 1px solid #0070c0; padding: 2px 5px; color: #0070c0; text-decoration: none;">Edit Credentials</a>
---------------------------------	------------	---	--

From there, you can choose the account credentials that you want to use and the associated AWS marketplace subscription. You can even add a marketplace subscription, if you need to.

### Edit Account & Add Subscription

---

Credentials

Instance Profile | Account ID: [REDACTED]▼

Associated Subscription

● QA Subscription▼

#### Associate Subscription to Credentials


To create a pay-as-you-go Cloud Volumes ONTAP system, you need to select AWS credentials that are associated with a subscription to Cloud Volumes ONTAP from the AWS Marketplace.

[+ Add Subscription](#)

Apply

Cancel

And if you manage multiple AWS subscriptions, you can assign each one of them to different AWS credentials from the Credentials page in the settings:



[Learn how to manage AWS credentials in Cloud Manager.](#)

### Timeline enhancements

The Timeline was enhanced to provide you with more information about the NetApp cloud services that you use.

- The Timeline now shows actions for all Cloud Manager systems within the same Cloud Central account
- You can now find information more easily by filtering, searching, and adding and removing columns
- You can now download the timeline data in CSV format
- In the future, the Timeline will show actions for each NetApp cloud service that you use (but you can filter the information down to a single service)

Time	Action	Service	Agent	Resource	User	Status
Jan 23 2020, 10:00:19 am	Check Connectivity	Cloud Manager	Ben_23Jan2020	CloudVolumesONTAP1	Ben	Success
Jan 23 2020, 10:00:02 am	Create Vsa Working Environment	Cloud Manager	Ben_23Jan2020		Ben	Pending
Jan 23 2020, 9:59:49 am	Update Cloud Ontap Metadata	Cloud Manager	Ben_23Jan2020		System	Success
Jan 23 2020, 9:58:43 am	Attach Subscription To Cloud Account	Cloud Manager	Ben_23Jan2020		Ben	Success
Jan 23 2020, 9:57:46 am	Initial Setup With Portal	Cloud Manager	Ben_23Jan2020		Ben	Success

### Cloud Manager 3.8 (8 Jan 2020)

- [HA enhancements in Azure](#)
- [Data tiering enhancements in GCP](#)

### HA enhancements in Azure

The following enhancements are now available for Cloud Volumes ONTAP HA pairs in Azure.

- **Override CIFS locks for Cloud Volumes ONTAP HA in Azure**

You can now enable a setting in Cloud Manager that prevents issues with Cloud Volumes ONTAP storage failover during Azure maintenance events. When you enable this setting, Cloud Volumes ONTAP vetoes CIFS locks and resets active CIFS sessions. [Learn more](#).

- **HTTPS connection from Cloud Volumes ONTAP to storage accounts**

You can now enable an HTTPS connection from a Cloud Volumes ONTAP 9.7 HA pair to Azure storage accounts when creating a working environment. Note that enabling this option can impact write performance. You can't change the setting after you create the working environment.

- **Support for Azure general-purpose v2 storage accounts**

The storage accounts that Cloud Manager creates for Cloud Volumes ONTAP 9.7 HA pairs are now general-purpose v2 storage accounts.

### Data tiering enhancements in GCP

The following enhancements are available for Cloud Volumes ONTAP data tiering in GCP.

- **Google Cloud storage classes for data tiering**

You can now choose a storage class for data that Cloud Volumes ONTAP tiers to Google Cloud Storage:

- Standard Storage (default)
- Nearline Storage
- Coldline Storage

[Learn more about Google Cloud storage classes](#).

[Learn how to change the storage class for Cloud Volumes ONTAP](#).

- **Data tiering using a service account**

Starting with the 9.7 release, Cloud Manager now sets a service account on the Cloud Volumes ONTAP instance. This service account provides permissions for data tiering to a Google Cloud Storage bucket. This change provides more security and requires less setup. For step-by-step instructions when deploying a new system, [see step 4 on this page](#).

The following image shows the Working Environment wizard where you can select a storage class and service account:

## Data Tiering in Google Cloud Platform

Data tiering can reduce your storage costs by automatically tiering cold data to a Google Cloud Storage bucket.

---

<a href="#">Tiering data to object storage</a>	Data Tiering Tiering Enabled	<a href="#">Edit</a>	Storage Class Standard Storage	<a href="#">Edit</a>
--	---------------------------------	----------------------	-----------------------------------	----------------------

Select a GCP service account to enable data tiering.  
[Learn more about data tiering in GCP.](#)

Service Account  
tiering-cloud-volumes-ontap

Cloud Manager requires the following GCP permissions for these enhancements, as shown in the latest [Cloud Manager policy for GCP](#).

- `storage.buckets.update`
- `compute.instances.setServiceAccount`
- `iam.serviceAccounts.getIamPolicy`
- `iam.serviceAccounts.list`

## Cloud Manager transition to SaaS

We've introduced a software-as-a-service experience for Cloud Manager. This new experience makes it easier for you to use Cloud Manager and enables us to provide additional features to manage your hybrid cloud infrastructure.

### The previous Cloud Manager experience

Cloud Manager software was previously comprised of a user interface and a management layer that sent requests to cloud providers. To get started, you would deploy Cloud Manager in your cloud network or on-premises network and then access the user interface that runs on that instance.

That experience has changed.

### The new SaaS experience

The Cloud Manager interface is now accessible through a SaaS-based user interface that you log in to from NetApp Cloud Central. You no longer need to access a user interface from software that runs in your network.


In most cases, you need to deploy a *Connector* in your cloud or on-premises network. The Connector is software that's needed to manage Cloud Volumes ONTAP and other cloud data services. (The Connector is actually the same as the existing Cloud Manager software that you have installed.)

## Benefits

This SaaS-based approach provides several benefits:


- It enables us to offer additional management capabilities for Azure NetApp Files and Cloud Volumes Service without needing to deploy software in your environment.
- You can easily switch between your Cloud Central accounts.

If a user is associated with multiple Cloud Central accounts, they can change to a different account at any time from the User Settings menu. They can then see the Connectors and working environments that are associated with that account.



- You can easily switch between Connectors (what you know today as the Cloud Manager software) that are installed in different networks or different cloud providers.






## The local user interface

While you should perform almost all tasks from the SaaS user interface, a local user interface is still available on the Connector. This interface is needed for a few tasks that need to be performed from the Connector itself:

- Setting a proxy server
- Installing a patch
- Downloading AutoSupport messages

You can access the local user interface directly from the SaaS user interface:



### Instance, VM, and machine type changes

To ensure that adequate resources are available for new and upcoming features in Cloud Manager, we've changed the minimum required instance, VM, and machine type as follows:

- AWS: t3.xlarge
- Azure: DS3 v2
- GCP: n1-standard-4

When you upgrade the machine type, you'll get access to features like a new Kubernetes experience, Global File Cache, Monitoring, and more.

These default sizes are the minimum supported [based on CPU and RAM requirements](#).

Cloud Manager will prompt you with instructions to change the machine type of the Connector.

### Known issues

Known issues identify problems that might prevent you from using this release of the product successfully.

There are no known issues in this release of Cloud Manager.

You can find known issues for Cloud Volumes ONTAP in the [Cloud Volumes ONTAP Release Notes](#) and for ONTAP software in general in the [ONTAP Release Notes](#).

### Known limitations

Known limitations identify platforms, devices, or functions that are not supported by this release of the product, or that do not interoperate correctly with it. Review these limitations carefully.

## **Connectors should remain running**

A Connector should remain running at all times. It's important for the continued health and operation of the services that you enable.

For example, a Connector is a key component in the health and operation of Cloud Volumes ONTAP PAYGO systems. If a Connector is powered down, Cloud Volumes ONTAP PAYGO systems will shut down after losing communication with a Connector for longer than 14 days.

## **SaaS platform is disabled for Government regions**

If you deploy a Connector in an AWS GovCloud region, an Azure Gov region, or an Azure DoD region, access to Cloud Manager is available only through a Connector's host IP address. Access to the SaaS platform is disabled for the entire account.

This means that only privileged users who can access the end-user internal VPC/VNet can use Cloud Manager's UI or API.

It also means that the following services aren't available from Cloud Manager:

- Cloud Compliance
- Kubernetes
- Cloud Tiering
- Global File Cache
- Monitoring (Cloud Insights)

The SaaS platform is required to use these services.

## **Shared Linux hosts are not supported**

The Connector isn't supported on a host that is shared with other applications. The host must be a dedicated host.

## **Cloud Manager does not support FlexGroup volumes**

While Cloud Volumes ONTAP supports FlexGroup volumes, Cloud Manager does not. If you create a FlexGroup volume from System Manager or from the CLI, then you should set Cloud Manager's Capacity Management mode to Manual. Automatic mode might not work properly with FlexGroup volumes.

# Important changes in Cloud Manager

This page highlights important changes in Cloud Manager that can help you use the service as we introduce new enhancements. You should continue to read the [What's new](#) page to learn about all new features and enhancements.

## SaaS changes

We have introduced a software-as-a-service experience for Cloud Manager. This new experience makes it easier for you to use Cloud Manager and enables us to provide additional features to manage your hybrid cloud infrastructure.

- [Cloud Manager transition to SaaS](#)
- [Learn how Cloud Manager works](#)

## Machine type changes

To ensure that adequate resources are available for new and upcoming features in Cloud Manager, we've changed the minimum required instance, VM, and machine type as follows:

- AWS: t3.xlarge
- Azure: DS3 v2
- GCP: n1-standard-4

When you upgrade the machine type, you'll get access to features like a new Kubernetes experience, Global File Cache, Monitoring, and more.

These default sizes are the minimum supported [based on CPU and RAM requirements](#).

Cloud Manager will prompt you with instructions to change the machine type of the Connector.

## Account settings

We introduced Cloud Central accounts to provide multi-tenancy, to help you organize users and resources in isolated workspaces, and to manage access to Connectors and subscriptions.

- [Learn about Cloud Central accounts: users, workspaces, Connectors, and subscriptions](#)
- [Learn how to get started with your account](#)
- [Learn how to manage your account after you set it up](#)

## New permissions

Cloud Manager occasionally requires additional cloud provider permissions as we introduce new features and enhancements. This section identifies new permissions that are now required.

You can find the latest list of permissions on the [Cloud Manager policies page](#).

## AWS

Starting with the 3.8.1 release, the following permissions are required to use Backup to Cloud with Cloud Volumes ONTAP. [Learn more](#).

```
{
  "Sid": "backupPolicy",
  "Effect": "Allow",
  "Action": [
    "s3:DeleteBucket",
    "s3:GetLifecycleConfiguration",
    "s3:PutLifecycleConfiguration",
    "s3:PutBucketTagging",
    "s3:ListBucketVersions",
    "s3:GetObject",
    "s3:ListBucket",
    "s3:ListAllMyBuckets",
    "s3:GetBucketTagging",
    "s3:GetBucketLocation",
    "s3:GetBucketPolicyStatus",
    "s3:GetBucketPublicAccessBlock",
    "s3:GetBucketAcl",
    "s3:GetBucketPolicy",
    "s3:PutBucketPublicAccessBlock"
  ],
  "Resource": [
    "arn:aws:s3:::netapp-backup-*"
  ]
}
```

## Azure

- To avoid Azure deployment failures, make sure that your Cloud Manager policy in Azure includes the following permission:

```
"Microsoft.Resources/deployments/operationStatuses/read"
```

- Starting with the 3.8.7 release, the following permission is required to encrypt Azure managed disks on single node Cloud Volumes ONTAP systems using external keys from another account. [Learn more](#).

```
"Microsoft.Compute/diskEncryptionSets/read"
```

- The following permissions are required to enable Global File Cache on Cloud Volumes ONTAP. [Learn more](#).

```
"Microsoft.Resources/deployments/operationStatuses/read",  
"Microsoft.Insights/Metrics/Read",  
"Microsoft.Compute/virtualMachines/extensions/write",  
"Microsoft.Compute/virtualMachines/extensions/read",  
"Microsoft.Compute/virtualMachines/extensions/delete",  
"Microsoft.Compute/virtualMachines/delete",  
"Microsoft.Network/networkInterfaces/delete",  
"Microsoft.Network/networkSecurityGroups/delete",  
"Microsoft.Resources/deployments/delete",
```

## GCP

### New permissions for Kubernetes management

Starting with the 3.8.8 release, the service account for a Connector requires the following permissions to discover and manage Kubernetes clusters running in Google Kubernetes Engine (GKE):

```
- container.*
```

### New permissions for data tiering

Starting with the 3.8 release, the following permissions are required to use a service account for data tiering. [Learn more.](#)

```
- storage.buckets.update  
- compute.instances.setServiceAccount  
- iam.serviceAccounts.getIamPolicy  
- iam.serviceAccounts.list
```

## New endpoints

The Connector requires outbound internet access to manage resources and processes within your public cloud environment. This section identifies new endpoints that are now required.

You can find the [full list of endpoints accessed from your web browser here](#) and the [full list of endpoints accessed by the Connector here](#).

- Users need to access Cloud Manager from a web browser by contacting the following endpoint:

```
https://cloudmanager.netapp.com
```

- Connectors require access to the following endpoint to obtain software images of container components for a Docker infrastructure:

```
https://cloudmanagerinfraproduct.azurecr.io
```

Ensure that your firewall enables access to this endpoint from the Connector.

# Get started with Cloud Manager

## Learn about Cloud Manager

Cloud Manager enables IT experts and cloud architects to centrally manage their hybrid multi-cloud infrastructure using NetApp's cloud solutions.

### Features

Cloud Manager is an enterprise-class, SaaS-based management platform that keeps you in control of your data no matter where it is.

- Set up and use [Cloud Volumes ONTAP](#) for efficient, multi-protocol data management across clouds.
- Set up and use file-storage services: [Azure NetApp Files](#), [Cloud Volumes Service for AWS](#), and [Cloud Volumes Service for Google Cloud](#).
- Discover and manage your on-prem ONTAP clusters by creating volumes, backing up to the cloud, replicating data across your hybrid cloud, and tiering cold data to the cloud.
- Enable integrated cloud services and software like [Cloud Compliance](#), [Cloud Insights](#), [Cloud Backup Service](#), [Trident](#), and more.

[Learn more about Cloud Manager.](#)

### Supported object storage providers

Cloud Manager enables you to manage cloud storage and use cloud services in Amazon Web Services, Microsoft Azure, and Google Cloud.

### Cost

Cloud Manager software is free of charge from NetApp.

For most tasks, Cloud Manager prompts you to deploy a Connector in your cloud network, which results in charges from your cloud provider for the compute instance and associated storage. You do have the option to run the Connector software on your premises.

### How Cloud Manager works

Cloud Manager includes a SaaS-based interface that is integrated with NetApp Cloud Central, and Connectors that manage Cloud Volumes ONTAP and other cloud services.

### Software-as-a-service

Cloud Manager is accessible through a [SaaS-based user interface](#) and APIs. This SaaS experience enables you to automatically access the latest features as they're released and to easily switch between your Cloud Central accounts and Connectors.

### NetApp Cloud Central

[NetApp Cloud Central](#) provides a centralized location to access and manage [NetApp cloud services](#). With centralized user authentication, you can use the same set of credentials to access Cloud Manager and other cloud services like Cloud Insights.



When you log in to Cloud Manager for the first time, you're prompted to create a *Cloud Central account*. This account provides multi-tenancy and enables you to organize users and resources in isolated *workspaces*.

## Connectors

In most cases, an Account Admin will need to deploy a *Connector* in your cloud or on-premises network. The Connector enables Cloud Manager to manage resources and processes within your public cloud environment.

A Connector should remain running at all times. It's important for the continued health and operation of the services that you enable.

For example, a Connector is a key component in the health and operation of Cloud Volumes ONTAP PAYGO systems. If a Connector is powered down, Cloud Volumes ONTAP PAYGO systems will shut down after losing communication with a Connector for longer than 14 days.

[Learn more about when Connectors are required and how they work.](#)

## Networking overview

Before users log in to Cloud Manager, you'll need to ensure that their web browsers can access specific endpoints. After that, you need to verify networking requirements for the specific type of working environment and services that will be used.

### Endpoints accessed from your web browser

Users must access Cloud Manager from a web browser. The machine running the web browser must have connections to the following endpoints:

Endpoints	Purpose
<a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a>	To connect you to the Cloud Manager SaaS interface.
<a href="https://api.services.cloud.netapp.com">https://api.services.cloud.netapp.com</a>	To contact Cloud Central APIs.
<a href="https://auth0.com">https://auth0.com</a> <a href="https://cdn.auth0.com">https://cdn.auth0.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://services.cloud.netapp.com">https://services.cloud.netapp.com</a>	Your web browser connects to these endpoints for centralized user authentication through NetApp Cloud Central.
<a href="https://widget.intercom.io">https://widget.intercom.io</a>	For in-product chat that enables you to talk to NetApp cloud experts.

### Index of networking requirements

- [Connectors](#)
- [Cloud Volumes ONTAP for AWS](#)
- [Cloud Volumes ONTAP for Azure](#)
- [Cloud Volumes ONTAP for GCP](#)
- [Data replication between ONTAP systems](#)
- [Cloud Compliance for Cloud Volumes ONTAP or Azure NetApp Files](#)
- [Cloud Compliance for Amazon S3](#)

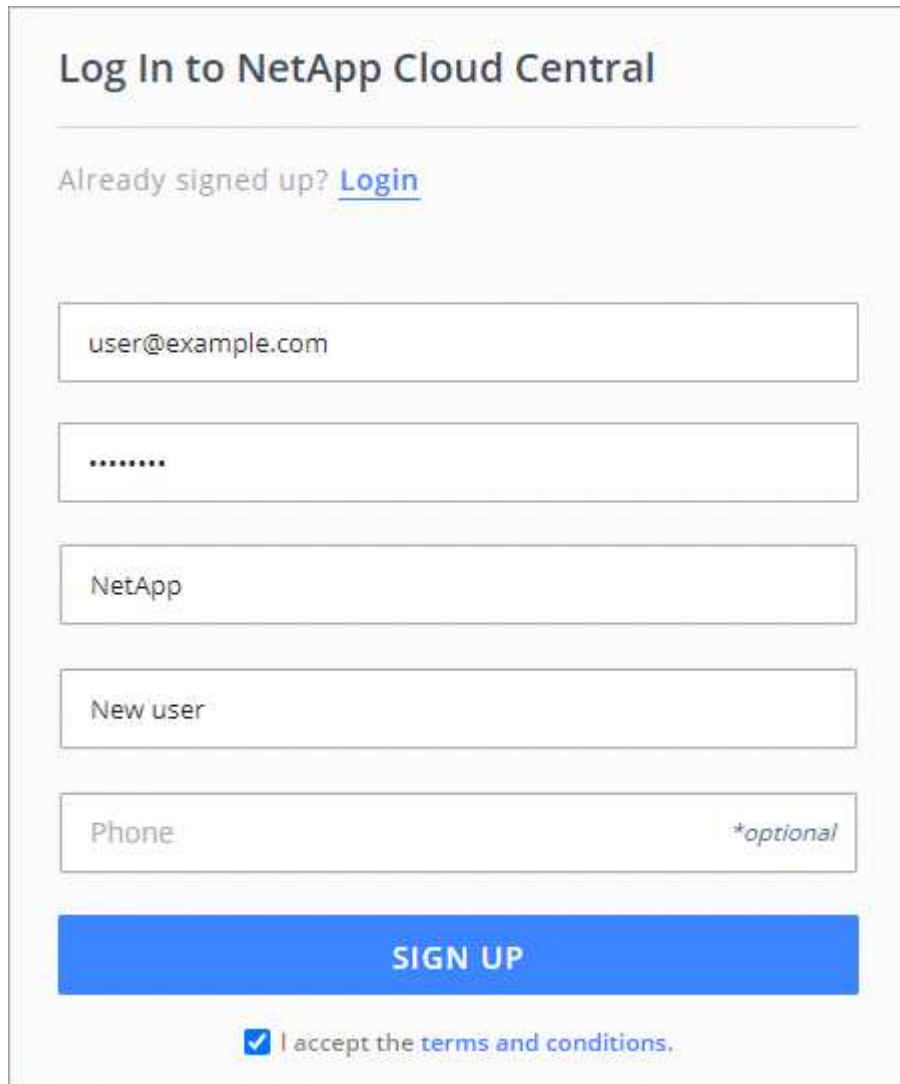
- On-prem ONTAP clusters
  - Data tiering from ONTAP clusters to Amazon S3
  - Data tiering from ONTAP clusters to Azure Blob storage
  - Data tiering from ONTAP clusters to Google Cloud Storage
  - Data tiering from ONTAP clusters to StorageGRID

## Signing up to NetApp Cloud Central

Sign up to NetApp Cloud Central so you can access NetApp's cloud services.

### Steps

1. Open a web browser and go to [NetApp Cloud Central](#).
2. Click **Sign Up**.
3. Fill out the form and click **Sign Up**.



The screenshot shows the 'Log In to NetApp Cloud Central' page. At the top, it says 'Log In to NetApp Cloud Central'. Below that, it asks 'Already signed up?' with a link to 'Login'. The form contains several input fields: 'user@example.com', a password field with dots, 'NetApp', 'New user', and 'Phone' (with '\*optional' next to it). A large blue 'SIGN UP' button is at the bottom. Below the button, there is a checkbox that is checked, followed by the text 'I accept the terms and conditions.'

4. Wait for an email from NetApp Cloud Central.
5. Click the link in the email to verify your email address.

## Result

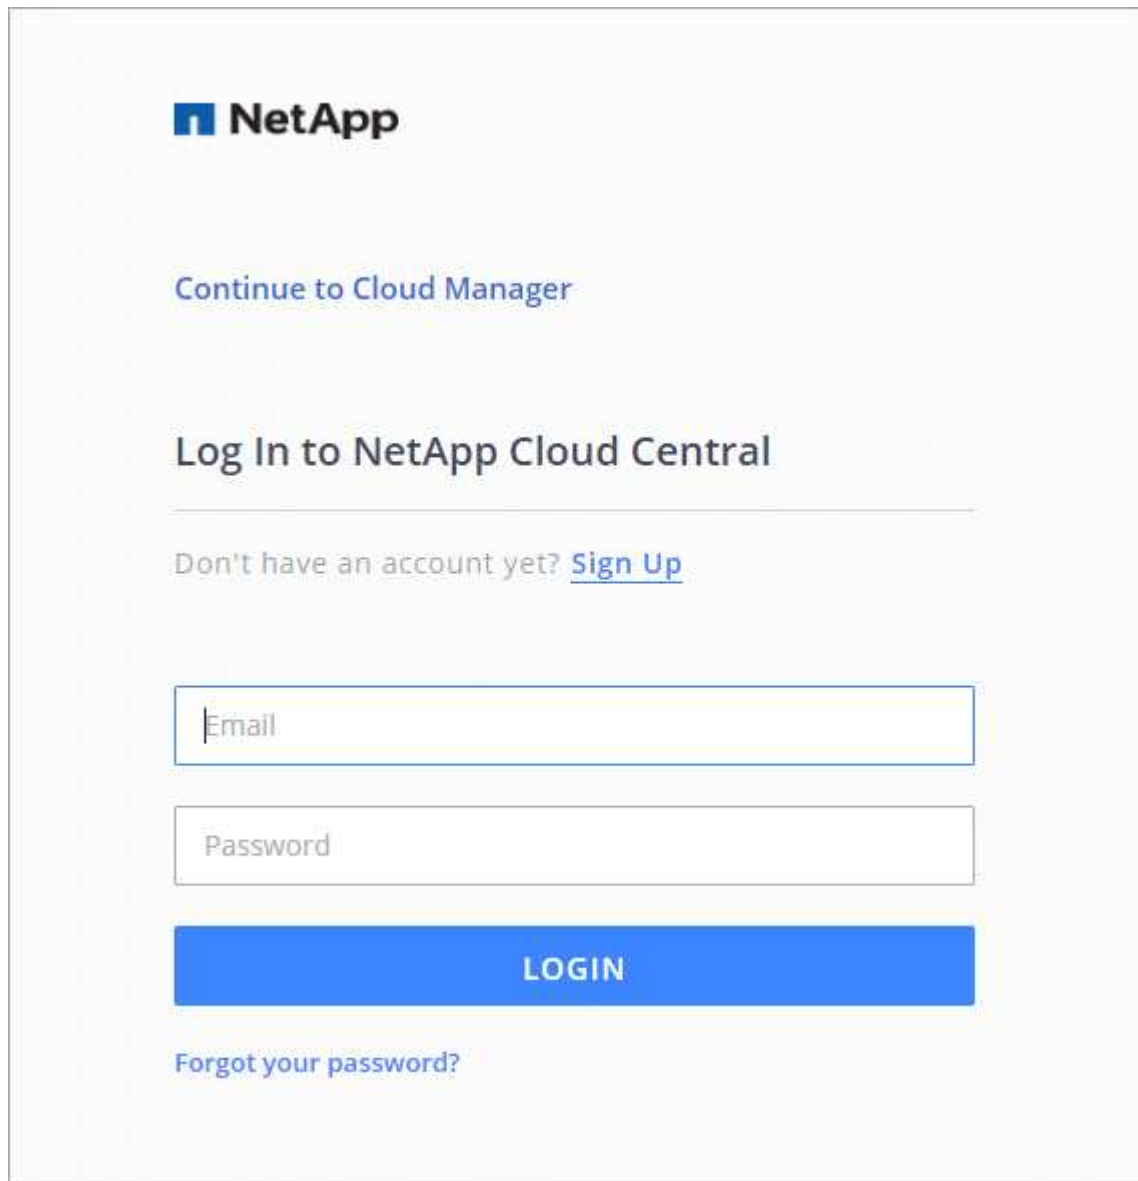
You now have an active Cloud Central user login.

## Logging in to Cloud Manager

The Cloud Manager interface is accessible through a SaaS-based user interface by going to <https://cloudmanager.netapp.com>.

### Steps

1. Open a web browser and go to <https://cloudmanager.netapp.com>.
2. Log in using your NetApp Cloud Central credentials.



The screenshot shows the NetApp Cloud Central login interface. At the top left is the NetApp logo. Below it is a link that says "Continue to Cloud Manager". The main heading is "Log In to NetApp Cloud Central". Underneath this heading is a horizontal line, followed by the text "Don't have an account yet?" and a link for "Sign Up". There are two input fields: the first is labeled "Email" and the second is labeled "Password". Below these fields is a prominent blue button with the text "LOGIN" in white. At the bottom left of the form area is a link that says "Forgot your password?".

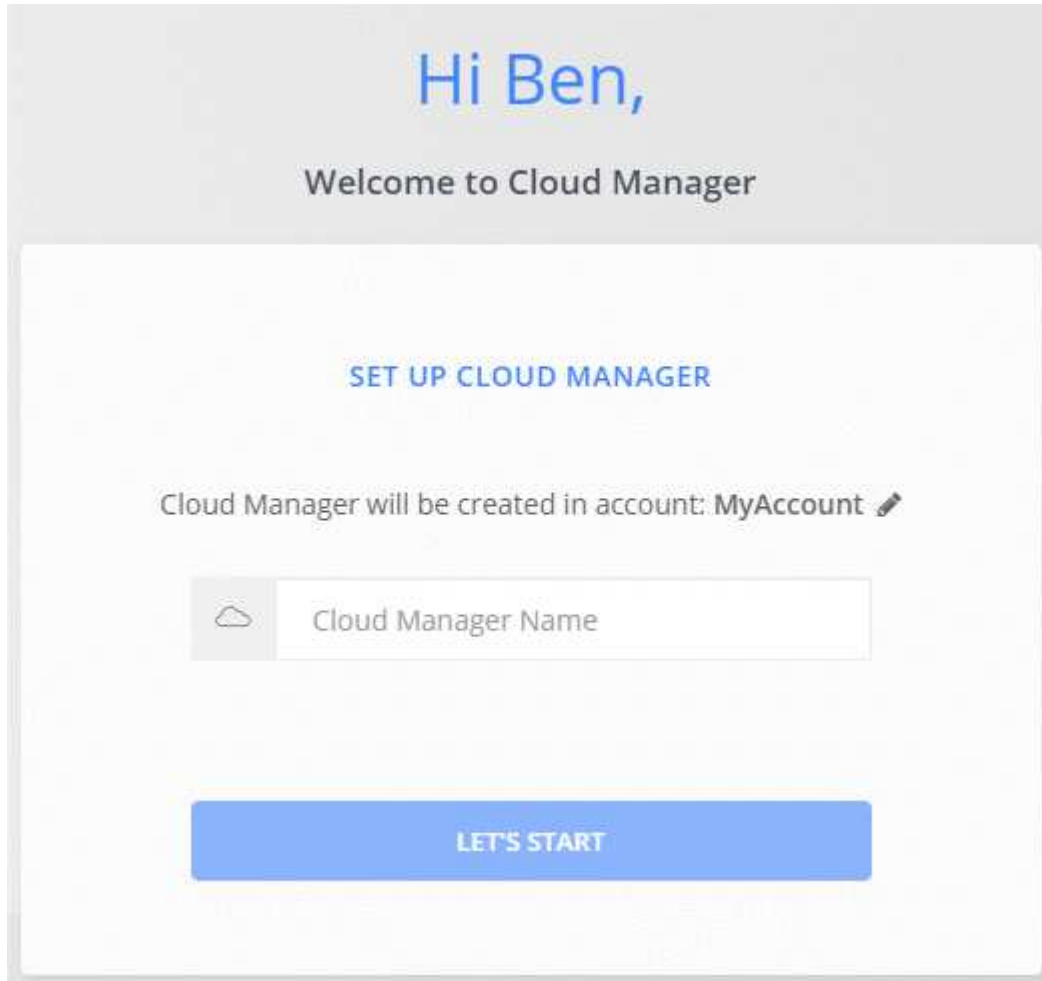
## Set up a Cloud Central account

## Account settings: users, workspaces, Connectors, and subscriptions

A *Cloud Central account* provides multi-tenancy and enables you to organize users and resources in isolated workspaces from within Cloud Manager.

For example, 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 first access Cloud Manager, you're prompted to select or create a Cloud Central account:




Hi Ben,

Welcome to Cloud Manager

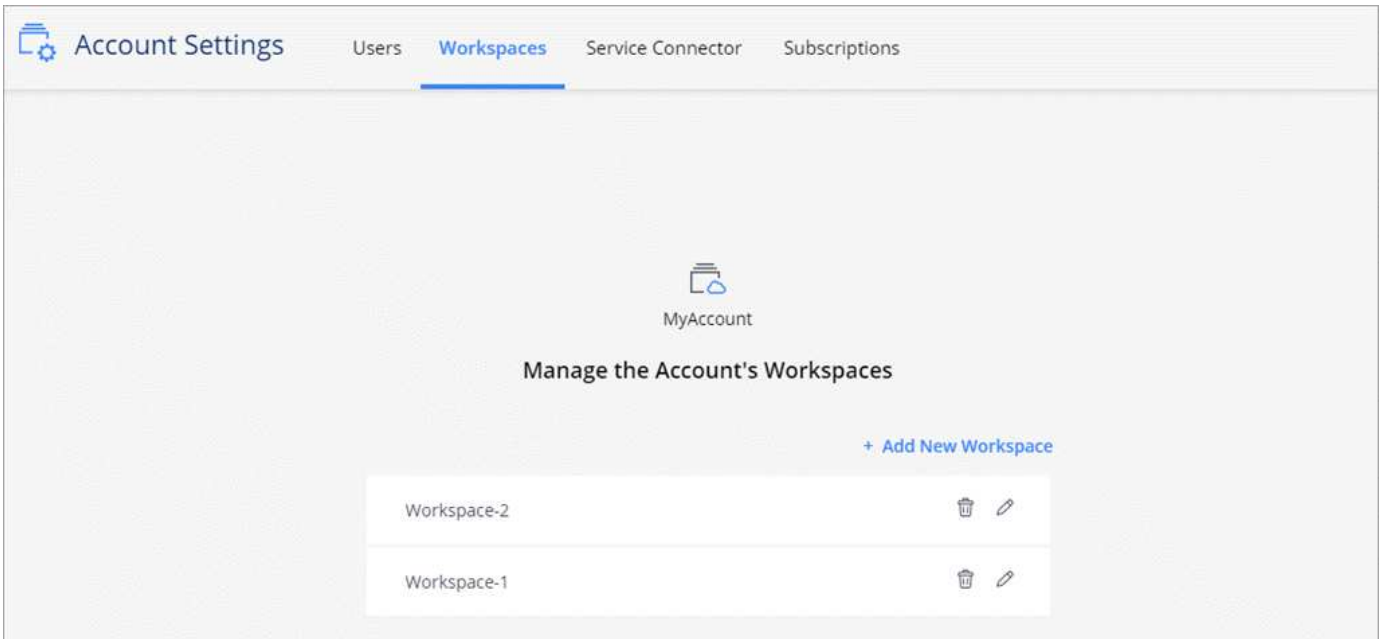
SET UP CLOUD MANAGER

Cloud Manager will be created in account: MyAccount 

 Cloud Manager Name

LET'S START

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



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

## Account Settings

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, Connectors, and subscriptions that are associated with the account.

## Users

The users that display in the Account Settings 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.
- *Cloud Compliance Viewer*: Can only view compliance information and generate reports for systems that they have permission to access.

## 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

## Connectors

A Connector enables Cloud Manager to manage resources and processes within your public cloud environment. The Connector runs on a virtual machine instance that you deploy in your cloud provider, or on an on-prem host that you configured.

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

## Subscriptions

The Account Settings widget shows the NetApp subscriptions associated with the selected account.

When you subscribe to Cloud Manager from a cloud provider's marketplace, you're redirected to Cloud Central where you need to save your subscription and associate it with specific accounts.

After you've subscribed, each subscription is available from the Account Settings widget. You'll only see the subscriptions that are associated with the account that you're currently viewing.

You have the option to rename a subscription and to disassociate the subscription from one or more accounts.

For example, let's say that you have two accounts and each is billed through separate subscriptions. You might disassociate a subscription from one of the accounts so the users in that account don't accidentally choose the wrong subscription when creating a Cloud Volume ONTAP working environment.

## Examples

The following examples depict how you might set up your accounts.




In both example images that follow, the Connector 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.

### Example 1

The following example shows an account that uses two workspaces to create isolated environments. The first workspace is for a production environment and the second is for a dev environment.

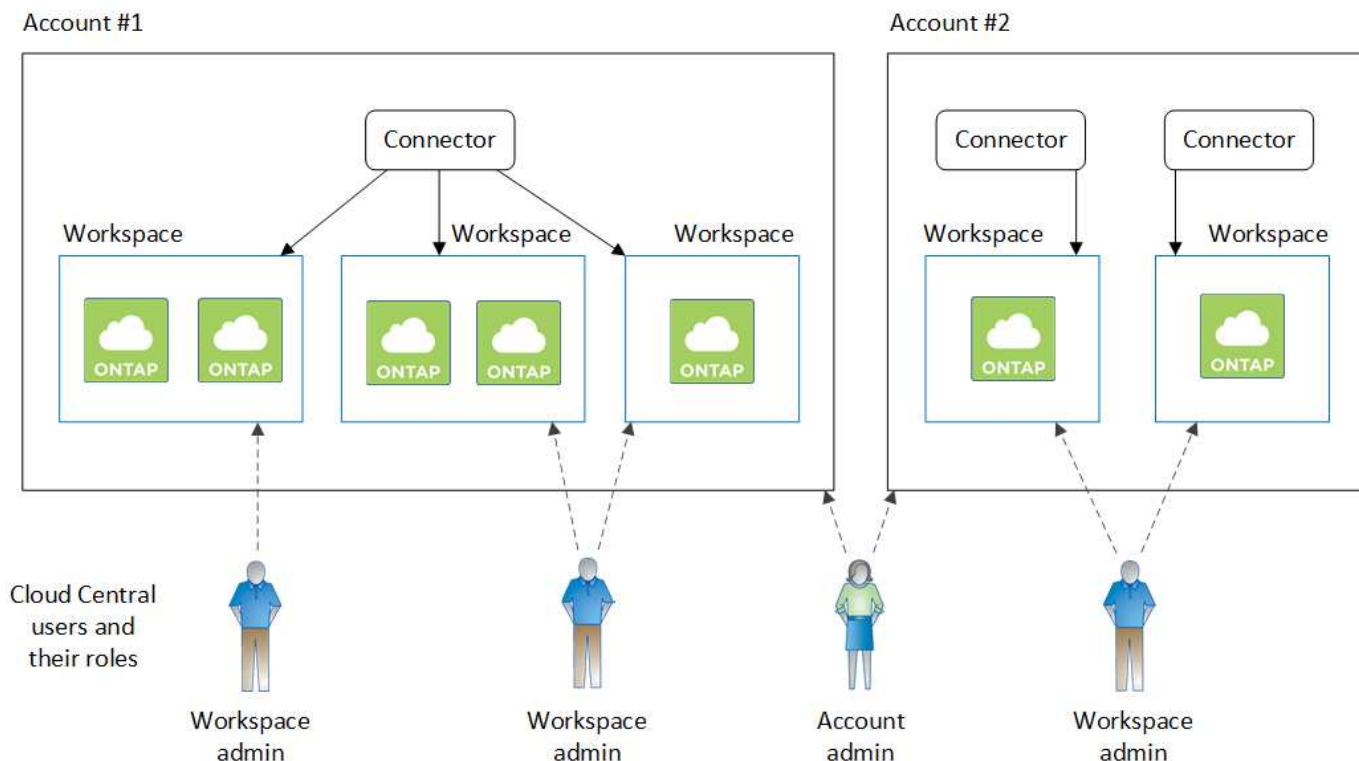
## Account



### Example 2

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 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 Connectors. This might happen if you have systems in separate regions or in separate cloud providers.



## Setting up workspaces and users in the Cloud Central account

When you log in to Cloud Manager for the first time, you're prompted to create a *NetApp Cloud Central account*. This account provides multi-tenancy and enables you to organize users and resources in isolated *workspaces*.

[Learn more about how Cloud Central accounts work.](#)

Set up your Cloud Central account so users can access Cloud Manager and access the working environments in a workspace. Just add a single user or add multiple users and workspaces.


### Adding workspaces

In Cloud Manager, workspaces enable you to isolate a set of working environments from other working environments and from other users. For example, you can create two workspaces and associate separate users with each workspace.


### Steps

1. From the top of Cloud Manager, click the **Account** drop-down.





2. Click **Manage Account** next to the currently selected account.



3. Click **Workspaces**.
4. Click **Add New Workspace**.
5. Enter a name for the workspace and click **Add**.

### After you finish


If a Workspace Admin needs access to this workspace, then you'll need to associate the user. You'll also need to associate Connectors with the workspace so Workspace Admins can use those Connectors.

### Adding users


Associate Cloud Central users with the Cloud Central account so those users can create and manage working environments in Cloud Manager.

### Steps

1. If the user hasn't already done so, ask the user to go to [NetApp Cloud Central](#) and sign up.
2. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.



3. From the Users tab, click **Associate User**.
4. Enter the user's email address and select a role for the user:
  - **Account Admin**: Can perform any action in Cloud Manager.
  - **Workspace Admin**: Can create and manage resources in assigned workspaces.
  - **Compliance Viewer**: Can only view compliance information and generate reports for workspaces that they have permission to access.
5. If you selected Workspace Admin or Compliance Viewer, select one or more workspaces to associate with that user.



## Associate User

To add a user to your NetApp Cloud Account, that user must already have signed up at [NetApp Cloud Central](#). Enter the email address that they used when signing up with Cloud Central.

User's Email

Role

Associate User to Workspaces

6. Click **Associate User**.

### Result


The user should receive an email from NetApp Cloud Central titled "Account Association." The email includes the information needed to access Cloud Manager.

### Associating Workspace Admins with workspaces


You can associate Workspace Admins with additional workspaces at any time. Associating the user enables them to create and view the working environments in that workspace.

### Steps

1. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.



2. From the Users tab, click the action menu in the row that corresponds to the user.



3. Click **Manage Workspaces**.

4. Select one or more workspaces and click **Apply**.

### Result

The user can now access those workspaces from Cloud Manager, as long as the Connector was also associated with the workspaces.

### Associating Connectors with workspaces


You need to associate a Connector with workspaces so Workspace Admins can use those Connectors to create Cloud Volumes ONTAP systems.

If you only have Account Admins, then associating the Connector with workspaces isn't required. Account Admins have the ability to access all workspaces in Cloud Manager by default.

[Learn more about users, workspaces, and Connectors.](#)

### Steps

1. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.



2. Click **Connector**.
3. Click **Manage Workspaces** for the Connector that you want to associate.
4. Select one or more workspaces and click **Apply**.

### Result

Workspace Admins can now use those Connectors to create Cloud Volumes ONTAP systems.

### What's next?

Now that you've set up your account, you can manage it any time by removing users, managing workspaces, Connectors, and subscriptions. [Learn more](#).

## Set up a Connector

### Learn about Connectors

In most cases, an Account Admin will need to deploy a *Connector* in your cloud or on-premises network. The Connector enables Cloud Manager to manage resources and processes within your public cloud environment.

### When a Connector is required

A Connector is required to use any of the following features within Cloud Manager:

- Cloud Volumes ONTAP
- On-premises ONTAP clusters
- Cloud Compliance
- Kubernetes
- Backup to Cloud
- Monitoring

- On-prem tiering
- Global File Cache
- Amazon S3 bucket discovery

A Connector is **not** required for Azure NetApp Files, Cloud Volumes Service, or Cloud Sync.



While a Connector isn't required to set up and manage Azure NetApp Files, a Connector is required if you want to use Cloud Compliance to scan Azure NetApp Files data.

## Supported locations

A Connector is supported in the following locations:

- Amazon Web Services
- Microsoft Azure
- Google Cloud
- On your premises



If you want to create a Cloud Volumes ONTAP system in Google Cloud, then you must have a Connector running in Google Cloud, as well. You can't use a Connector that's running in another location.

## Connectors should remain running

A Connector should remain running at all times. It's important for the continued health and operation of the services that you enable.

For example, a Connector is a key component in the health and operation of Cloud Volumes ONTAP PAYGO systems. If a Connector is powered down, Cloud Volumes ONTAP PAYGO systems will shut down after losing communication with a Connector for longer than 14 days.

## How to create a Connector

An Account Admin needs to create a Connector before a Workspace Admin can create a Cloud Volumes ONTAP working environment and use any of the other features listed above.

An Account Admin can create a Connector in a number of ways:

- Directly from Cloud Manager (recommended)
  - [Create in AWS](#)
  - [Create in Azure](#)
  - [Create in GCP](#)
- [From the AWS Marketplace](#)
- [From the Azure Marketplace](#)
- [By downloading and installing the software on an existing Linux host](#)

When you create your first Cloud Volumes ONTAP working environment, Cloud Manager will prompt you to create a Connector if you don't have one yet.

## Permissions

Specific permissions are needed to create the Connector and another set of permissions are needed for the Connector instance itself.

### Permissions to create a Connector

The user who creates a Connector from Cloud Manager needs specific permissions to deploy the instance in your cloud provider of choice. Cloud Manager will remind you of the permissions requirements when you create a Connector.

[View policies for each cloud provider.](#)

### Permissions for the Connector instance

The Connector needs specific cloud provider permissions to perform operations on your behalf. For example, to deploy and manage Cloud Volumes ONTAP.

When you create a Connector directly from Cloud Manager, Cloud Manager creates the Connector with the permissions that it needs. There's nothing that you need to do.

If you create the Connector yourself from the AWS Marketplace, the Azure Marketplace, or by manually installing the software, then you'll need to make sure that the right permissions are in place.

[View policies for each cloud provider.](#)

## When to use multiple Connectors

In some cases, you might only need one Connector, but you might find yourself needing two or more Connectors.

Here are a few examples:

- You're using a multi-cloud environment (AWS and Azure), so you have one Connector in AWS and another in Azure. Each manages the Cloud Volumes ONTAP systems running in those environments.
- A service provider might use one Cloud Central account to provide services for their customers, while using another account to provide disaster recovery for one of their business units. Each account would have separate Connectors.

## When to switch between Connectors

When you create your first Connector, Cloud Manager automatically uses that Connector for each additional working environment that you create. Once you create an additional Connector, you'll need to switch between them to see the working environments that are specific to each Connector.

[Learn how to switch between Connectors.](#)

## The local user interface

While you should perform almost all tasks from the [SaaS user interface](#), a local user interface is still available on the Connector. This interface is needed for a few tasks that need to be performed from the Connector itself:

- [Setting a proxy server](#)
- Installing a patch (you'll typically work with NetApp personnel to install a patch)

- Downloading AutoSupport messages (usually directed by NetApp personnel when you have issues)

[Learn how to access the local UI.](#)

## Connector upgrades

The Connector automatically updates its software to the latest version, as long as it has [outbound internet access](#) to obtain the software update.

## Networking requirements for the Connector

Set up your networking so the Connector can manage resources and processes within your public cloud environment. The most important step is ensuring outbound internet access to various endpoints.



If your network uses a proxy server for all communication to the internet, you can specify the proxy server from the Settings page. Refer to [Configuring the Connector to use a proxy server](#).

## Connection to target networks

A Connector requires a network connection to the type of working environment that you're creating and the services that you're planning to enable.

For example, if you install a Connector in your corporate network, then you must set up a VPN connection to the VPC or VNet in which you launch Cloud Volumes ONTAP.

## Outbound internet access

The Connector requires outbound internet access to manage resources and processes within your public cloud environment. Outbound internet access is also required if you want to manually install the Connector on a Linux host or access the local UI running on the Connector.

The following sections identify the specific endpoints.

## Endpoints to manage resources in AWS

A Connector contacts the following endpoints when managing resources in AWS:



Endpoints	Purpose
<p>AWS services (amazonaws.com):</p> <ul style="list-style-type: none"> <li>• CloudFormation</li> <li>• Elastic Compute Cloud (EC2)</li> <li>• Key Management Service (KMS)</li> <li>• Security Token Service (STS)</li> <li>• Simple Storage Service (S3)</li> </ul> <p>The exact endpoint depends on the region in which you deploy Cloud Volumes ONTAP.  <a href="#">Refer to AWS documentation for details.</a></p>	<p>Enables the Connector to deploy and manage Cloud Volumes ONTAP in AWS.</p>
<p><a href="https://api.services.cloud.netapp.com:443">https://api.services.cloud.netapp.com:443</a></p>	<p>API requests to NetApp Cloud Central.</p>
<p><a href="https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com">https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com</a></p>	<p>Provides access to software images, manifests, and templates.</p>
<p><a href="https://repo.cloud.support.netapp.com">https://repo.cloud.support.netapp.com</a></p>	<p>Used to download Cloud Manager dependencies.</p>
<p><a href="http://repo.mysql.com/">http://repo.mysql.com/</a></p>	<p>Used to download MySQL.</p>
<p><a href="https://cognito-idp.us-east-1.amazonaws.com">https://cognito-idp.us-east-1.amazonaws.com</a>  <a href="https://cognito-identity.us-east-1.amazonaws.com">https://cognito-identity.us-east-1.amazonaws.com</a>  <a href="https://sts.amazonaws.com">https://sts.amazonaws.com</a>  <a href="https://cloud-support-netapp-com-accelerated.s3.amazonaws.com">https://cloud-support-netapp-com-accelerated.s3.amazonaws.com</a></p>	<p>Enables the Connector to access and download manifests, templates, and Cloud Volumes ONTAP upgrade images.</p>
<p><a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a></p>	<p>Access to software images of container components for an infrastructure that's running Docker and provides a solution for service integrations with Cloud Manager.</p>
<p><a href="https://kinesis.us-east-1.amazonaws.com">https://kinesis.us-east-1.amazonaws.com</a></p>	<p>Enables NetApp to stream data from audit records.</p>
<p><a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a></p>	<p>Communication with the Cloud Manager service, which includes Cloud Central accounts.</p>
<p><a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a></p>	<p>Communication with NetApp Cloud Central for centralized user authentication.</p>
<p><a href="https://w86yt021u5.execute-api.us-east-1.amazonaws.com/production/whitelist">https://w86yt021u5.execute-api.us-east-1.amazonaws.com/production/whitelist</a></p>	<p>Used to add your AWS account ID to the list of allowed users for Backup to S3.</p>
<p><a href="https://support.netapp.com/aods/asupmessage">https://support.netapp.com/aods/asupmessage</a>  <a href="https://support.netapp.com/asupprod/post/1.0/postAsup">https://support.netapp.com/asupprod/post/1.0/postAsup</a></p>	<p>Communication with NetApp AutoSupport.</p>

Endpoints	Purpose
<a href="https://support.netapp.com/svcgw">https://support.netapp.com/svcgw</a> <a href="https://support.netapp.com/ServiceGW/entitlement">https://support.netapp.com/ServiceGW/entitlement</a> <a href="https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com">https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com</a> <a href="https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com">https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com</a>	Communication with NetApp for system licensing and support registration.
<a href="https://client.infra.support.netapp.com.s3.us-west-1.amazonaws.com">https://client.infra.support.netapp.com.s3.us-west-1.amazonaws.com</a> <a href="https://cloud-support-netapp-com-accelerated.s3.us-west-1.amazonaws.com">https://cloud-support-netapp-com-accelerated.s3.us-west-1.amazonaws.com</a> <a href="https://trigger.asup.netapp.com.s3.us-west-1.amazonaws.com">https://trigger.asup.netapp.com.s3.us-west-1.amazonaws.com</a>	Enables NetApp to collect information needed to troubleshoot support issues.
<a href="https://ipa-signer.cloudmanager.netapp.com">https://ipa-signer.cloudmanager.netapp.com</a>	Enables Cloud Manager to generate licenses (for example, a FlexCache license for Cloud Volumes ONTAP)
<a href="https://packages.cloud.google.com/yum">https://packages.cloud.google.com/yum</a> <a href="https://github.com/NetApp/trident/releases/download/">https://github.com/NetApp/trident/releases/download/</a>	Required to connect Cloud Volumes ONTAP systems with a Kubernetes cluster. The endpoints enable installation of NetApp Trident.
<p>Various third-party locations, for example:</p> <ul style="list-style-type: none"> <li>• <a href="https://repo1.maven.org/maven2">https://repo1.maven.org/maven2</a></li> <li>• <a href="https://oss.sonatype.org/content/repositories">https://oss.sonatype.org/content/repositories</a></li> <li>• <a href="https://repo.typesafe.com">https://repo.typesafe.com</a></li> </ul> <p>Third-party locations are subject to change.</p>	During upgrades, Cloud Manager downloads the latest packages for third-party dependencies.

### Endpoints to manage resources in Azure

A Connector contacts the following endpoints when managing resources in Azure:

Endpoints	Purpose
<a href="https://management.azure.com">https://management.azure.com</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a>	Enables Cloud Manager to deploy and manage Cloud Volumes ONTAP in most Azure regions.
<a href="https://management.microsoftazure.de">https://management.microsoftazure.de</a> <a href="https://login.microsoftonline.de">https://login.microsoftonline.de</a>	Enables Cloud Manager to deploy and manage Cloud Volumes ONTAP in the Azure Germany regions.
<a href="https://management.usgovcloudapi.net">https://management.usgovcloudapi.net</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a>	Enables Cloud Manager to deploy and manage Cloud Volumes ONTAP in the Azure US Gov regions.
<a href="https://api.services.cloud.netapp.com:443">https://api.services.cloud.netapp.com:443</a>	API requests to NetApp Cloud Central.
<a href="https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com">https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com</a>	Provides access to software images, manifests, and templates.
<a href="https://repo.cloud.support.netapp.com">https://repo.cloud.support.netapp.com</a>	Used to download Cloud Manager dependencies.

Endpoints	Purpose
http://repo.mysql.com/	Used to download MySQL.
<a href="https://cognito-idp.us-east-1.amazonaws.com">https://cognito-idp.us-east-1.amazonaws.com</a> <a href="https://cognito-identity.us-east-1.amazonaws.com">https://cognito-identity.us-east-1.amazonaws.com</a> <a href="https://sts.amazonaws.com">https://sts.amazonaws.com</a> <a href="https://cloud-support-netapp-com-accelerated.s3.amazonaws.com">https://cloud-support-netapp-com-accelerated.s3.amazonaws.com</a>	Enables the Connector to access and download manifests, templates, and Cloud Volumes ONTAP upgrade images.
https://cloudmanagerinfraprod.azurecr.io	Access to software images of container components for an infrastructure that's running Docker and provides a solution for service integrations with Cloud Manager.
https://kinesis.us-east-1.amazonaws.com	Enables NetApp to stream data from audit records.
https://cloudmanager.cloud.netapp.com	Communication with the Cloud Manager service, which includes Cloud Central accounts.
https://netapp-cloud-account.auth0.com	Communication with NetApp Cloud Central for centralized user authentication.
https://mysupport.netapp.com	Communication with NetApp AutoSupport.
<a href="https://support.netapp.com/svcgw">https://support.netapp.com/svcgw</a> <a href="https://support.netapp.com/ServiceGW/entitlement">https://support.netapp.com/ServiceGW/entitlement</a> <a href="https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com">https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com</a> <a href="https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com">https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com</a>	Communication with NetApp for system licensing and support registration.
<a href="https://client.infra.support.netapp.com.s3.us-west-1.amazonaws.com">https://client.infra.support.netapp.com.s3.us-west-1.amazonaws.com</a> <a href="https://cloud-support-netapp-com-accelerated.s3.us-west-1.amazonaws.com">https://cloud-support-netapp-com-accelerated.s3.us-west-1.amazonaws.com</a> <a href="https://trigger.asup.netapp.com.s3.us-west-1.amazonaws.com">https://trigger.asup.netapp.com.s3.us-west-1.amazonaws.com</a>	Enables NetApp to collect information needed to troubleshoot support issues.
https://ipa-signer.cloudmanager.netapp.com	Enables Cloud Manager to generate licenses (for example, a FlexCache license for Cloud Volumes ONTAP)
<a href="https://packages.cloud.google.com/yum">https://packages.cloud.google.com/yum</a> <a href="https://github.com/NetApp/trident/releases/download/">https://github.com/NetApp/trident/releases/download/</a>	Required to connect Cloud Volumes ONTAP systems with a Kubernetes cluster. The endpoints enable installation of NetApp Trident.
*.blob.core.windows.net	Required for HA pairs when using a proxy.

Endpoints	Purpose
Various third-party locations, for example: <ul style="list-style-type: none"> <li>• <a href="https://repo1.maven.org/maven2">https://repo1.maven.org/maven2</a></li> <li>• <a href="https://oss.sonatype.org/content/repositories">https://oss.sonatype.org/content/repositories</a></li> <li>• <a href="https://repo.typesafe.com">https://repo.typesafe.com</a></li> </ul> Third-party locations are subject to change.	During upgrades, Cloud Manager downloads the latest packages for third-party dependencies.

### Endpoints to manage resources in GCP

A Connector contacts the following endpoints when managing resources in GCP:

Endpoints	Purpose
<a href="https://www.googleapis.com">https://www.googleapis.com</a>	Enables the Connector to contact Google APIs for deploying and managing Cloud Volumes ONTAP in GCP.
<a href="https://api.services.cloud.netapp.com:443">https://api.services.cloud.netapp.com:443</a>	API requests to NetApp Cloud Central.
<a href="https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com">https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com</a>	Provides access to software images, manifests, and templates.
<a href="https://repo.cloud.support.netapp.com">https://repo.cloud.support.netapp.com</a>	Used to download Cloud Manager dependencies.
<a href="http://repo.mysql.com/">http://repo.mysql.com/</a>	Used to download MySQL.
<a href="https://cognito-idp.us-east-1.amazonaws.com">https://cognito-idp.us-east-1.amazonaws.com</a> <a href="https://cognito-identity.us-east-1.amazonaws.com">https://cognito-identity.us-east-1.amazonaws.com</a> <a href="https://sts.amazonaws.com">https://sts.amazonaws.com</a> <a href="https://cloud-support-netapp-com-accelerated.s3.amazonaws.com">https://cloud-support-netapp-com-accelerated.s3.amazonaws.com</a>	Enables the Connector to access and download manifests, templates, and Cloud Volumes ONTAP upgrade images.
<a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a>	Access to software images of container components for an infrastructure that's running Docker and provides a solution for service integrations with Cloud Manager.
<a href="https://kinesis.us-east-1.amazonaws.com">https://kinesis.us-east-1.amazonaws.com</a>	Enables NetApp to stream data from audit records.
<a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a>	Communication with the Cloud Manager service, which includes Cloud Central accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	Communication with NetApp Cloud Central for centralized user authentication.
<a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	Communication with NetApp AutoSupport.

Endpoints	Purpose
<a href="https://support.netapp.com/svcgw">https://support.netapp.com/svcgw</a> <a href="https://support.netapp.com/ServiceGW/entitlement">https://support.netapp.com/ServiceGW/entitlement</a> <a href="https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com">https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com</a> <a href="https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com">https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com</a>	Communication with NetApp for system licensing and support registration.
<a href="https://client.infra.support.netapp.com.s3.us-west-1.amazonaws.com">https://client.infra.support.netapp.com.s3.us-west-1.amazonaws.com</a> <a href="https://cloud-support-netapp-com-accelerated.s3.us-west-1.amazonaws.com">https://cloud-support-netapp-com-accelerated.s3.us-west-1.amazonaws.com</a> <a href="https://trigger.asup.netapp.com.s3.us-west-1.amazonaws.com">https://trigger.asup.netapp.com.s3.us-west-1.amazonaws.com</a>	Enables NetApp to collect information needed to troubleshoot support issues.
<a href="https://ipa-signer.cloudmanager.netapp.com">https://ipa-signer.cloudmanager.netapp.com</a>	Enables Cloud Manager to generate licenses (for example, a FlexCache license for Cloud Volumes ONTAP)
<a href="https://packages.cloud.google.com/yum">https://packages.cloud.google.com/yum</a> <a href="https://github.com/NetApp/trident/releases/download/">https://github.com/NetApp/trident/releases/download/</a>	Required to connect Cloud Volumes ONTAP systems with a Kubernetes cluster. The endpoints enable installation of NetApp Trident.
<p>Various third-party locations, for example:</p> <ul style="list-style-type: none"> <li>• <a href="https://repo1.maven.org/maven2">https://repo1.maven.org/maven2</a></li> <li>• <a href="https://oss.sonatype.org/content/repositories">https://oss.sonatype.org/content/repositories</a></li> <li>• <a href="https://repo.typesafe.com">https://repo.typesafe.com</a></li> </ul> <p>Third-party locations are subject to change.</p>	During upgrades, Cloud Manager downloads the latest packages for third-party dependencies.

### Endpoints to install the Connector on a Linux host

You have the option to manually install the Connector software on your own Linux host. If you do, the installer for the Connector must access the following URLs during the installation process:

- <http://dev.mysql.com/get/mysql-community-release-el7-5.noarch.rpm>
- <https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm>
- <https://s3.amazonaws.com/aws-cli/awscli-bundle.zip>

The host might try to update operating system packages during installation. The host can contact different mirroring sites for these OS packages.

### Endpoints accessed from your web browser when using the local UI

While you should perform almost all tasks from the SaaS user interface, a local user interface is still available on the Connector. The machine running the web browser must have connections to the following endpoints:

Endpoints	Purpose
The Connector host	<p>You must enter the host's IP address from a web browser to load the Cloud Manager console.</p> <p>Depending on your connectivity to your cloud provider, you can use the private IP or a public IP assigned to the host:</p> <ul style="list-style-type: none"> <li>• A private IP works if you have a VPN and direct connect access to your virtual network</li> <li>• A public IP works in any networking scenario</li> </ul> <p>In any case, you should secure network access by ensuring that security group rules allow access from only authorized IPs or subnets.</p>
<a href="https://auth0.com">https://auth0.com</a> <a href="https://cdn.auth0.com">https://cdn.auth0.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://services.cloud.netapp.com">https://services.cloud.netapp.com</a>	Your web browser connects to these endpoints for centralized user authentication through NetApp Cloud Central.
<a href="https://widget.intercom.io">https://widget.intercom.io</a>	For in-product chat that enables you to talk to NetApp cloud experts.

## Ports and security groups

There's no incoming traffic to the Connector, unless you initiate it. HTTP and HTTPS provide access to the [local UI](#), which you'll use in rare circumstances. SSH is only needed if you need to connect to the host for troubleshooting.

### Rules for the Connector in AWS

The security group for the Connector requires both inbound and outbound rules.

#### Inbound rules

The source for inbound rules in the predefined security group is 0.0.0.0/0.

Protocol	Port	Purpose
SSH	22	Provides SSH access to the Connector host
HTTP	80	Provides HTTP access from client web browsers to the local user interface and connections from Cloud Compliance
HTTPS	443	Provides HTTPS access from client web browsers to the local user interface
TCP	3128	Provides the Cloud Compliance instance with internet access, if your AWS network doesn't use a NAT or proxy

#### Outbound rules

The predefined security group for the Connector opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

## Basic outbound rules

The predefined security group for the Connector includes the following outbound rules.

Protocol	Port	Purpose
All TCP	All	All outbound traffic
All UDP	All	All outbound traffic

## Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by the Connector.



The source IP address is the Connector host.

Service	Protocol	Port	Destination	Purpose
Active Directory	TCP	88	Active Directory forest	Kerberos V authentication
	TCP	139	Active Directory forest	NetBIOS service session
	TCP	389	Active Directory forest	LDAP
	TCP	445	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	TCP	464	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	TCP	749	Active Directory forest	Active Directory Kerberos V change & set password (RPCSEC_GSS)
	UDP	137	Active Directory forest	NetBIOS name service
	UDP	138	Active Directory forest	NetBIOS datagram service
	UDP	464	Active Directory forest	Kerberos key administration
API calls and AutoSupport	HTTPS	443	Outbound internet and ONTAP cluster management LIF	API calls to AWS and ONTAP, and sending AutoSupport messages to NetApp
API calls	TCP	3000	ONTAP cluster management LIF	API calls to ONTAP
	TCP	8088	Backup to S3	API calls to Backup to S3
DNS	UDP	53	DNS	Used for DNS resolve by Cloud Manager
Cloud Compliance	HTTP	80	Cloud Compliance instance	Cloud Compliance for Cloud Volumes ONTAP

## Rules for the Connector in Azure

The security group for the Connector requires both inbound and outbound rules.

## Inbound rules

The source for inbound rules in the predefined security group is 0.0.0.0/0.

Port	Protocol	Purpose
22	SSH	Provides SSH access to the Connector host
80	HTTP	Provides HTTP access from client web browsers to the local user interface
443	HTTPS	Provides HTTPS access from client web browsers to the local user interface

## Outbound rules

The predefined security group for the Connector opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

### Basic outbound rules

The predefined security group for the Connector includes the following outbound rules.

Port	Protocol	Purpose
All	All TCP	All outbound traffic
All	All UDP	All outbound traffic

### Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by the Connector.



The source IP address is the Connector host.



Service	Port	Protocol	Destination	Purpose
Active Directory	88	TCP	Active Directory forest	Kerberos V authentication
	139	TCP	Active Directory forest	NetBIOS service session
	389	TCP	Active Directory forest	LDAP
	445	TCP	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	464	TCP	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	749	TCP	Active Directory forest	Active Directory Kerberos V change & set password (RPCSEC_GSS)
	137	UDP	Active Directory forest	NetBIOS name service
	138	UDP	Active Directory forest	NetBIOS datagram service
	464	UDP	Active Directory forest	Kerberos key administration
API calls and AutoSupport	443	HTTPS	Outbound internet and ONTAP cluster management LIF	API calls to AWS and ONTAP, and sending AutoSupport messages to NetApp
API calls	3000	TCP	ONTAP cluster management LIF	API calls to ONTAP
DNS	53	UDP	DNS	Used for DNS resolve by Cloud Manager

### Rules for the Connector in GCP

The firewall rules for the Connector requires both inbound and outbound rules.

#### Inbound rules

The source for inbound rules in the predefined firewall rules is 0.0.0.0/0.

Protocol	Port	Purpose
SSH	22	Provides SSH access to the Connector host
HTTP	80	Provides HTTP access from client web browsers to the local user interface
HTTPS	443	Provides HTTPS access from client web browsers to the local user interface

#### Outbound rules

The predefined firewall rules for the Connector opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

#### Basic outbound rules

The predefined firewall rules for the Connector includes the following outbound rules.

Protocol	Port	Purpose
All TCP	All	All outbound traffic
All UDP	All	All outbound traffic

### Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by the Connector.



The source IP address is the Connector host.

Service	Protocol	Port	Destination	Purpose
Active Directory	TCP	88	Active Directory forest	Kerberos V authentication
	TCP	139	Active Directory forest	NetBIOS service session
	TCP	389	Active Directory forest	LDAP
	TCP	445	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	TCP	464	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	TCP	749	Active Directory forest	Active Directory Kerberos V change & set password (RPCSEC_GSS)
	UDP	137	Active Directory forest	NetBIOS name service
	UDP	138	Active Directory forest	NetBIOS datagram service
	UDP	464	Active Directory forest	Kerberos key administration
API calls and AutoSupport	HTTPS	443	Outbound internet and ONTAP cluster management LIF	API calls to GCP and ONTAP, and sending AutoSupport messages to NetApp
API calls	TCP	3000	ONTAP cluster management LIF	API calls to ONTAP
DNS	UDP	53	DNS	Used for DNS resolve by Cloud Manager

### Creating a Connector in AWS from Cloud Manager

An Account Admin needs to deploy a *Connector* before you can use most Cloud Manager features. [Learn when a Connector is required](#). The Connector enables Cloud Manager to manage resources and processes within your public cloud environment.

This page describes how to create a Connector in AWS directly from Cloud Manager. You also have the option to [create the Connector from the AWS Marketplace](#), or to [download the software and install it on your own host](#).

These steps must be completed by a user who has the Account Admin role. A Workspace Admin can't create a Connector.



When you create your first Cloud Volumes ONTAP working environment, Cloud Manager will prompt you to create a Connector if you don't have one yet.

## Setting up AWS permissions to create a Connector

Before you can deploy a Connector from Cloud Manager, you need to ensure that your AWS account has the correct permissions.

### Steps

1. Download the Connector IAM policy from the following location:

[NetApp Cloud Manager: AWS, Azure, and GCP Policies](#)

2. From the AWS IAM console, create your own policy by copying and pasting the text from the Connector IAM policy.
3. Attach the policy that you created in the previous step to the IAM user who will create the Connector from Cloud Manager.

### Result

The AWS user now has the permissions required to create the Connector from Cloud Manager. You'll need to specify AWS access keys for this user when you're prompted by Cloud Manager.

## Creating a Connector in AWS


Cloud Manager enables you to create a Connector in AWS directly from its user interface.

### What you'll need

- An AWS access key and secret key for an IAM user who has the [required permissions](#).
- A VPC, subnet, and keypair in your AWS region of choice.

### Steps

1. If you're creating your first Working Environment, click **Add Working Environment** and follow the prompts. Otherwise, click the **Connector** drop-down and select **Add Connector**.



2. Click **Let's Start**.
3. Choose **Amazon Web Services** as your cloud provider.

Remember that the Connector must have a network connection to the type of working environment that you're creating and the services that you're planning to enable.

[Learn more about networking requirements for the Connector.](#)

4. Review what you'll need and click **Continue**.

5. Provide the required information:

- **AWS Credentials:** Enter a name for the instance and specify the AWS access key and secret key that meet permissions requirements.
- **Location:** Specify an AWS region, VPC, and subnet for the instance.
- **Network:** Select the key pair to use with the instance, whether to enable a public IP address, and optionally specify a proxy configuration.
- **Security Group:** Choose whether to create a new security group or whether to select an existing security group that allows inbound HTTP, HTTPS, and SSH access.



There's no incoming traffic to the Connector, unless you initiate it. HTTP and HTTPS provide access to the [local UI](#), which you'll use in rare circumstances. SSH is only needed if you need to connect to the host for troubleshooting.

6. Click **Create**.

The instance should be ready in about 7 minutes. You should stay on the page until the process is complete.

### After you finish

You need to associate a Connector with workspaces so Workspace Admins can use those Connectors to create Cloud Volumes ONTAP systems. If you only have Account Admins, then associating the Connector with workspaces isn't required. Account Admins have the ability to access all workspaces in Cloud Manager by default. [Learn more](#).

## Creating a Connector in Azure from Cloud Manager

An Account Admin needs to deploy a *Connector* before you can use most Cloud Manager features. [Learn when a Connector is required](#). The Connector enables Cloud Manager to manage resources and processes within your public cloud environment.

This page describes how to create a Connector in Azure directly from Cloud Manager. You also have the option to [create the Connector from the Azure Marketplace](#), or to [download the software and install it on your own host](#).

These steps must be completed by a user who has the Account Admin role. A Workspace Admin can't create a Connector.



When you create your first Cloud Volumes ONTAP working environment, Cloud Manager will prompt you to create a Connector if you don't have one yet.

### Setting up Azure permissions to create a Connector

Before you can deploy a Connector from Cloud Manager, you need to ensure that your Azure account has the correct permissions.

## Steps

1. Create a custom role using the Azure policy for the Connector:
  - a. Download the [Azure policy for the Connector](#).



Right-click the link and click **Save link as...** to download the file.

- b. Modify the JSON file by adding your Azure subscription ID to the assignable scope.

### Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
]
```

- c. Use the JSON file to create a custom role in Azure.

The following example shows how to create a custom role using the Azure CLI 2.0:

```
az role definition create --role-definition  
C:\Policy_for_Setup_As_Service_Azure.json
```

You should now have a custom role called *Azure SetupAsService*.

2. Assign the role to the user who will deploy the Connector from Cloud Manager:
  - a. Open the **Subscriptions** service and select the user's subscription.
  - b. Click **Access control (IAM)**.
  - c. Click **Add > Add role assignment** and then add the permissions:
    - Select the **Azure SetupAsService** role.



Azure SetupAsService is the default name provided in the [Connector deployment policy for Azure](#). If you chose a different name for the role, then select that name instead.

- Assign access to an **Azure AD user, group, or application**.
- Select the user account.
- Click **Save**.

## Result

The Azure user now has the permissions required to deploy the Connector from Cloud Manager.

## Creating a Connector in Azure

Cloud Manager enables you to create a Connector in Azure directly from its user interface.


### What you'll need

- The [required permissions](#) for your Azure account.
- An Azure subscription.

- A VNet and subnet in your Azure region of choice.

## Steps

1. If you're creating your first Working Environment, click **Add Working Environment** and follow the prompts. Otherwise, click the **Connector** drop-down and select **Add Connector**.



2. Click **Let's Start**.
3. Choose **Microsoft Azure** as your cloud provider.

Remember that the Connector must have a network connection to the type of working environment that you're creating and the services that you're planning to enable.

[Learn more about networking requirements for the Connector.](#)

4. Review what you'll need and click **Continue**.
5. If you're prompted, log in to your Microsoft account, which should have the required permissions to create the virtual machine.

The form is owned and hosted by Microsoft. Your credentials are not provided to NetApp.



If you're already logged in to an Azure account, then Cloud Manager will automatically use that account. If you have multiple accounts, then you might need to log out first to ensure that you're using the right account.

6. Provide the required information:
  - **VM Authentication:** Enter a name for the virtual machine and a user name and password or public key.
  - **Basic Settings:** Choose an Azure subscription, an Azure region, and whether to create a new resource group or to use an existing resource group.
  - **Network:** Choose a VNet and subnet, whether to enable a public IP address, and optionally specify a proxy configuration.
  - **Security Group:** Choose whether to create a new security group or whether to select an existing security group that allows inbound HTTP, HTTPS, and SSH access.



There's no incoming traffic to the Connector, unless you initiate it. HTTP and HTTPS provide access to the [local UI](#), which you'll use in rare circumstances. SSH is only needed if you need to connect to the host for troubleshooting.

#### 7. Click **Create**.

The virtual machine should be ready in about 7 minutes. You should stay on the page until the process is complete.

#### **After you finish**

You need to associate a Connector with workspaces so Workspace Admins can use those Connectors to create Cloud Volumes ONTAP systems. If you only have Account Admins, then associating the Connector with workspaces isn't required. Account Admins have the ability to access all workspaces in Cloud Manager by default. [Learn more](#).

## **Creating a Connector in GCP from Cloud Manager**

An Account Admin needs to deploy a *Connector* before you can use most Cloud Manager features. [Learn when a Connector is required](#). The Connector enables Cloud Manager to manage resources and processes within your public cloud environment.

This page describes how to create a Connector in GCP directly from Cloud Manager. You also have the option to [download the software and install it on your own host](#).

These steps must be completed by a user who has the Account Admin role. A Workspace Admin can't create a Connector.



When you create your first Cloud Volumes ONTAP working environment, Cloud Manager will prompt you to create a Connector if you don't have one yet.

### **Setting up GCP permissions to create a Connector**

Before you can deploy a Connector from Cloud Manager, you need to ensure that your GCP account has the correct permissions and that a service account is set up for the Connector VM.

#### **Steps**

1. Ensure that the GCP user who deploys Cloud Manager from NetApp Cloud Central has the permissions in the [Connector deployment policy for GCP](#).

[You can create a custom role using the YAML file](#) and then attach it to the user. You'll need to use the `gcloud` command line to create the role.

2. Set up a service account that has the permissions that Cloud Manager needs to create and manage Cloud Volumes ONTAP systems in projects.

You'll associate this service account with the Connector VM when you create it from Cloud Manager.

- a. [Create a role in GCP](#) that includes the permissions defined in the [Cloud Manager policy for GCP](#). Again, you'll need to use the `gcloud` command line.

The permissions contained in this YAML file are different than the permissions in step 2a.

- b. [Create a GCP service account and apply the custom role that you just created.](#)
- c. If you want to deploy Cloud Volumes ONTAP in other projects, [grant access by adding the service account with the Cloud Manager role to that project.](#) You'll need to repeat this step for each project.

## Result

The GCP user now has the permissions required to create the Connector from Cloud Manager and the service account for the Connector VM is set up.

## Enabling Google Cloud APIs

Several APIs are required to deploy the Connector and Cloud Volumes ONTAP.

### Step

1. [Enable the following Google Cloud APIs in your project.](#)
  - Cloud Deployment Manager V2 API
  - Cloud Logging API
  - Cloud Resource Manager API
  - Compute Engine API
  - Identity and Access Management (IAM) API

## Creating a Connector in GCP


Cloud Manager enables you to create a Connector in GCP directly from its user interface.

### What you'll need

- The [required permissions](#) for your Google Cloud account.
- A Google Cloud project.
- A service account that has the required permissions to create and manage Cloud Volumes ONTAP.
- A VPC and subnet in your Google Cloud region of choice.

### Steps

1. If you're creating your first Working Environment, click **Add Working Environment** and follow the prompts. Otherwise, click the **Connector** drop-down and select **Add Connector**.



2. Click **Let's Start**.



3. Choose **Google Cloud Platform** as your cloud provider.

Remember that the Connector must have a network connection to the type of working environment that you're creating and the services that you're planning to enable.

[Learn more about networking requirements for the Connector.](#)

4. Review what you'll need and click **Continue**.
5. If you're prompted, log in to your Google account, which should have the required permissions to create the virtual machine instance.

The form is owned and hosted by Google. Your credentials are not provided to NetApp.

6. Provide the required information:
  - **Basic Settings:** Enter a name for the virtual machine instance and specify a project and service account that has the required permissions.
  - **Location:** Specify a region, zone, VPC, and subnet for the instance.
  - **Network:** Choose whether to enable a public IP address and optionally specify a proxy configuration.
  - **Firewall Policy:** Choose whether to create a new firewall policy or whether to select an existing firewall policy that allows inbound HTTP, HTTPS, and SSH access.



There's no incoming traffic to the Connector, unless you initiate it. HTTP and HTTPS provide access to the [local UI](#), which you'll use in rare circumstances. SSH is only needed if you need to connect to the host for troubleshooting.

7. Click **Create**.

The instance should be ready in about 7 minutes. You should stay on the page until the process is complete.

### After you finish

You need to associate a Connector with workspaces so Workspace Admins can use those Connectors to create Cloud Volumes ONTAP systems. If you only have Account Admins, then associating the Connector with workspaces isn't required. Account Admins have the ability to access all workspaces in Cloud Manager by default. [Learn more.](#)

## Where to go next

Now that you've logged in and set up Cloud Manager, users can start creating and discovering working environments.

- [Get started with Cloud Volumes ONTAP for AWS](#)
- [Get started with Cloud Volumes ONTAP for Azure](#)
- [Get started with Cloud Volumes ONTAP for Google Cloud](#)
- [Set up Azure NetApp Files](#)
- [Set up Cloud Volumes Service for AWS](#)
- [Discover an on-premises ONTAP cluster](#)

- [Discover your Amazon S3 buckets](#)

If you're an administrator, you can manage Cloud Manager settings after you create your first Connector.

- [Learn about Connectors](#)
- [Manage an HTTPS certificate for secure access](#)
- [Configure proxy settings](#)

# Manage Cloud Volumes ONTAP

## Learn

### Learn about Cloud Volumes ONTAP

Cloud Volumes ONTAP enables you to optimize your cloud storage costs and performance while enhancing data protection, security, and compliance.

Cloud Volumes ONTAP is a software-only storage appliance that runs ONTAP data management software in the cloud. It provides enterprise-grade storage 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 protection

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.

Cloud Volumes ONTAP also integrates with Cloud Backup Service to deliver backup and restore capabilities for protection, and long-term archive of your cloud data.

- 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.

- Data security

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

- Privacy compliance controls

Integration with Cloud Compliance helps you understand data context and identify sensitive data.



Licenses for ONTAP features are included with Cloud Volumes ONTAP.

[View supported Cloud Volumes ONTAP configurations](#)

[Learn more about Cloud Volumes ONTAP](#)

## Storage

### Disks and aggregates


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



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.

### 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

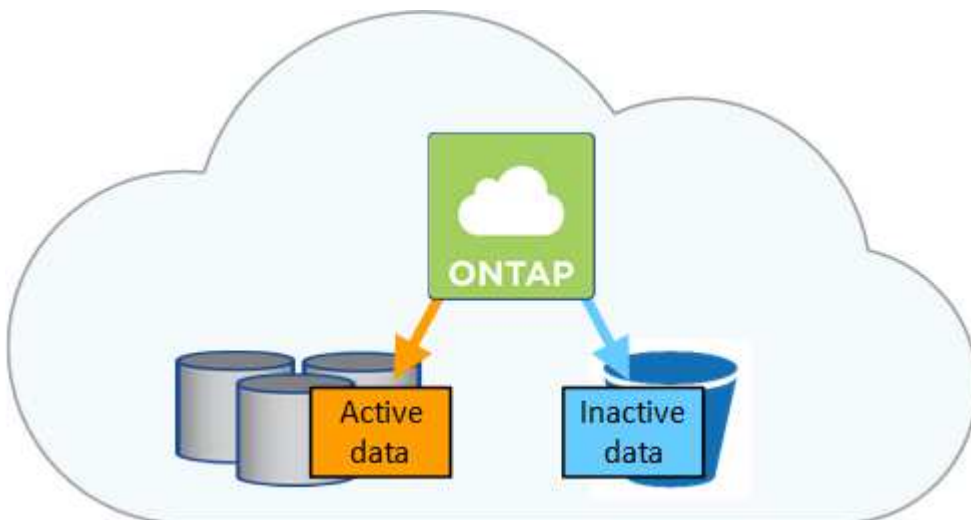
- [Google Cloud Platform documentation: Storage Options](#)
- [Review storage limits for Cloud Volumes ONTAP in GCP](#)

## 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 don't 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.

#### 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.

### Storage classes

The default storage class for tiered data in AWS is *Standard*. If you don't plan to access the inactive data, you can reduce your storage costs by changing the storage class to one of the following: *Intelligent Tiering*, *One-Zone Infrequent Access*, or *Standard-Infrequent Access*. When you change the storage class, inactive data starts in the Standard storage class and transitions to the storage class that you selected, if the data is not accessed after 30 days.

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

You can select a storage class when you create the working environment and you can change it any time after. For details about changing the storage class, see [Tiering inactive data to low-cost object storage](#).

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

### Data tiering in Azure

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

#### Performance tier

The performance tier can be either SSDs or HDDs.

#### Capacity tier

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

## Storage access tiers

The default storage access tier for tiered data in Azure is the *hot* tier. If you don't plan to access the inactive data, you can reduce your storage costs by changing to the *cool* storage tier. When you change the storage tier, inactive data starts in the hot storage tier and transitions 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 storage tier. [Learn more about Azure Blob storage access tiers.](#)

You can select a storage tier when you create the working environment and you can change it any time after. For details about changing the storage tier, see [Tiering inactive data to low-cost object storage.](#)

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

## Data tiering in 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.

## Storage classes

The default storage class for tiered data is the *Standard Storage* class. If the data is infrequently accessed, you can reduce your storage costs by changing to *Nearline Storage* or *Coldline Storage*. When you change the storage class, inactive data starts in the Standard Storage class and transitions to the storage class that you selected, if the data is not accessed after 30 days.

The access costs are higher if you do access the data, so take that into consideration before you change the storage class. [Learn more about storage classes for Google Cloud Storage.](#)

You can select a storage tier when you create the working environment and you can change it any time after. For details about changing the storage class, see [Tiering inactive data to low-cost object storage.](#)

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

## Data tiering and capacity limits

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

## Volume tiering policies

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

Some tiering policies have an associated minimum cooling period, which sets the time that user data in a



volume must remain inactive for the data to be considered "cold" and moved to the capacity tier.

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.

### All

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

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

### Auto

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

This policy is supported starting with Cloud Volumes ONTAP 9.4.

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

### None

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

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 can't be returned to AWS, Azure, or GCP in this scenario.

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

## Management of LUNs with automatic capacity management

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

## 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.

## Flash Cache

Some Cloud Volumes ONTAP configurations in AWS and Azure include local NVMe

storage, which Cloud Volumes ONTAP uses as *Flash Cache* for better performance.

### What's Flash Cache?

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

### Supported instances in AWS

Select one of the following EC2 instance types with a new or existing Cloud Volumes ONTAP Premium or BYOL system:

- c5d.4xlarge
- c5d.9xlarge
- c5d.18xlarge
- m5d.8xlarge
- m5d.12xlarge
- r5d.2xlarge

### Supported VM type in Azure

Select the Standard\_L8s\_v2 VM type with a single node Cloud Volumes ONTAP BYOL system in Azure.

### Limitations

- Compression must be disabled on all volumes to take advantage of the Flash Cache performance improvements.

Choose no storage efficiency when creating a volume from Cloud Manager, or create a volume and then [disable data compression by using the CLI](#).

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

### 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

Retention Period         ▼

### 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 can't be enabled.
- Backup to Cloud must be disabled in order to enable WORM storage.

## 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.





### 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:



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.

## Disks

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

## Storage accounts

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

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

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

## 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. The most common way is to launch the PAYGO version of your first Cloud Volumes ONTAP system to get a 30-day free trial. An evaluation BYOL license is also an option.

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.

### 30-day free trials for PAYGO

A 30-day free trial is available if you plan to pay for Cloud Volumes ONTAP as you go. You can start a 30-day free trial of Cloud Volumes ONTAP from Cloud Manager by creating your first Cloud Volumes ONTAP system in a payer's account.

There are no hourly software license charges for the instance, but infrastructure charges from your cloud provider still apply.

A free trial automatically converts to a paid hourly subscription when it expires. If you terminate the instance within the time limit, the next instance that you deploy is not part of the free trial (even if it's deployed within those 30 days).

Pay-as-you-go trials are awarded through a cloud provider and are not extendable by any means.

### Evaluation licenses for BYOL

An evaluation BYOL license is an option for customers who expect to pay for Cloud Volumes ONTAP by purchasing a termed license from NetApp. You can obtain an evaluation license from your account team, your Sales Engineer, or your partner.

The evaluation key is good for 30 days, and can be used multiple times, each for 30 days (regardless of the creation day).

At the end of 30 days, daily shutdowns will occur, so it's best to plan ahead. You can apply a new BYOL license on top of the evaluation license for an in-place upgrade (this requires a restart of single node systems). Your hosted data is **not** deleted at the end of the trial period.



You can't upgrade Cloud Volumes ONTAP software when using an evaluation license.

## Licensing

Each Cloud Volumes ONTAP BYOL system must have a system license installed with an

active subscription. Cloud Manager simplifies the process by managing licenses for you and by notifying you before they expire. BYOL licenses are also available for Backup to Cloud.

### BYOL system licenses

You can purchase multiple licenses for a Cloud Volumes ONTAP BYOL system to allocate more than 368 TB of capacity. For example, you might purchase two licenses to allocate up to 736 TB of capacity to Cloud Volumes ONTAP. Or you could purchase four licenses to get up to 1.4 PB.

The number of licenses that you can purchase for a single node system or HA pair is unlimited.

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

### License management for a new system

When you create a BYOL system, Cloud Manager prompts you for the serial number of your license and your 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 can't 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 [Managing BYOL licenses for Cloud Volumes ONTAP](#).

### License expiration warning

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 don't 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 can't 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 [Managing BYOL licenses for Cloud Volumes ONTAP](#).

## BYOL backup licenses

A BYOL backup license allows you to purchase a license from NetApp to use Backup to Cloud for a certain period of time and for a maximum amount backup space. When either limit is reached you will need to renew the license.

[Learn more about the Backup to Cloud BYOL license.](#)

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

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

### NetApp encryption solutions (NVE and NAE)

Cloud Volumes ONTAP supports both NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) with an external key manager. NVE and NAE are software-based solutions that enable (FIPS) 140-2-compliant data-at-rest encryption of volumes.

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

Both NVE and NAE use AES 256-bit encryption.

[Learn more about NetApp Volume Encryption and NetApp Aggregate Encryption.](#)

Starting with Cloud Volumes ONTAP 9.7, new aggregates will have NetApp Aggregate Encryption (NAE) enabled by default after you set up an external key manager. New volumes that aren't part of an NAE aggregate will have NetApp Volume Encryption (NVE) enabled by default (for example, if you have existing aggregates that were created before setting up an external key manager).

Setting up a supported key manager is the only required step. For set up instructions, see [Encrypting volumes](#)



with [NetApp encryption solutions](#).

### AWS Key Management Service

When you launch a Cloud Volumes ONTAP system in AWS, you can enable data encryption using the [AWS Key Management Service \(KMS\)](#). 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.

You can encrypt Azure managed disks on single node Cloud Volumes ONTAP systems using external keys from another account. This feature is supported using Cloud Manager APIs.

You just need to add the following to the API request when creating the single node system:

```
"azureEncryptionParameters": {  
  "key": <azure id of encryptionset>  
}
```



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

### 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. [Learn more](#).

### ONTAP virus scanning

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

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

For information about the vendors, software, and versions supported by Vscan, 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.

### Ransomware Protection

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

#### 1 Enable Snapshot Copy Protection




50 % Protection

1 Volumes without a Snapshot Policy

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

Activate Snapshot Policy

#### 2 Block Ransomware File Extensions



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

View Denied File Names

Activate FPolicy

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

## Performance

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

- Cloud Volumes ONTAP for AWS

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

- Cloud Volumes ONTAP for Microsoft Azure

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

- Cloud Volumes ONTAP for Google Cloud

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

## Default configuration for Cloud Volumes ONTAP

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

### Defaults

- Cloud Volumes ONTAP is available as a single-node system in AWS, Azure, and GCP, and as an HA pair in AWS and Azure.
- Cloud Manager creates one data-serving storage VM when it deploys Cloud Volumes ONTAP. Some configurations support additional storage VMs. [Learn more about managing storage VMs.](#)
- Cloud Manager automatically installs the following ONTAP feature licenses on Cloud Volumes ONTAP:
  - CIFS
  - FlexCache
  - FlexClone
  - iSCSI
  - NetApp Volume Encryption (only for BYOL or registered PAYGO systems)
  - NFS
  - SnapMirror
  - SnapRestore
  - SnapVault
- Several network interfaces are created by default:
  - A cluster management LIF
  - An intercluster LIF
  - An SVM management LIF on HA systems in Azure, single node systems in AWS, and optionally on HA systems in multiple AWS Availability Zones
  - A node management LIF
  - An iSCSI data LIF
  - A CIFS and NFS data LIF




LIF failover is disabled by default for Cloud Volumes ONTAP due to EC2 requirements. Migrating a LIF to a different port breaks the external mapping between IP addresses and network interfaces on the instance, making the LIF inaccessible.

- Cloud Volumes ONTAP sends configuration backups to the Connector using HTTPS.

The backups are accessible from <https://ipaddress/occm/offboxconfig/> where *ipaddress* is the IP address of the Connector host.

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

The following table lists the volume attributes that Cloud Manager sets differently from the defaults:

Attribute	Value set by Cloud Manager
Autosize mode	grow
Maximum autosize	1,000 percent  The Account Admin can modify this value from the Settings page.
Security style	NTFS for CIFS volumes UNIX for NFS volumes
Space guarantee style	none
UNIX permissions (NFS only)	777

See the *volume create* man page for information about these attributes.

## Boot and root data for Cloud Volumes ONTAP

In addition to the storage for user data, Cloud Manager also purchases cloud storage for boot and root data on each Cloud Volumes ONTAP system.

### AWS

- Two disks per node for boot and root data:
  - 9.7: 160 GB io1 disk for boot data and a 220 GB gp2 disk for root data
  - 9.6: 93 GB io1 disk for boot data and a 140 GB gp2 disk for root data
  - 9.5: 45 GB io1 disk for boot data and a 140 GB gp2 disk for root data
- One EBS snapshot for each boot disk and root disk
- For HA pairs, one EBS volume for the Mediator instance, which is approximately 8 GB

### Azure (single node)

- Three Premium SSD disks:
  - One 10 GB disk for boot data
  - One 140 GB disk for root data
  - One 128 GB disk for NVRAM

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

- One 1024 GB Standard HDD disk for saving cores
- One Azure snapshot for each boot disk and root disk

## Azure (HA pairs)

- Two 10 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 1024 GB Standard HDD disks for saving cores (one per node)
- Two 128 GB Premium SSD disks for NVRAM (one per node)
- One Azure snapshot for each boot disk and root disk

## GCP

- One 10 GB Standard persistent disk for boot data
- One 64 GB Standard persistent disk for root data
- One 500 GB Standard persistent disk for NVRAM
- One 216 GB Standard persistent disk for saving cores
- One GCP snapshot each for the boot disk and root disk

## Where the disks reside

Cloud Manager lays out the storage as follows:

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

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

- Root data, which contains the system configuration and logs, resides in aggr0.
- The storage virtual machine (SVM) root volume resides in aggr1.
- Data volumes also reside in aggr1.

## Encryption

Boot and root disks are always encrypted in Azure and Google Cloud Platform because encryption is enabled by default in those cloud providers.

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

# Get started in AWS

## Getting started with Cloud Volumes ONTAP for AWS

Get started with Cloud Volumes ONTAP for AWS in a few steps.



### Create a Connector

If you don't have a [Connector](#) yet, an Account Admin needs to create one. [Learn how to create a Connector in AWS.](#)

When you create your first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to deploy a Connector if you don't have one yet.



## Plan your configuration

Cloud Manager offers preconfigured packages that match your workload requirements, or you can create your own configuration. If you choose your own configuration, you should understand the options available to you.

[Learn more.](#)



## Set up your networking

- a. Ensure that your VPC and subnets will support connectivity between the Connector and Cloud Volumes ONTAP.
- b. Enable outbound internet access from the target VPC so the Connector and Cloud Volumes ONTAP can contact several endpoints.

This step is important because the Connector can't manage Cloud Volumes ONTAP without outbound internet access. If you need to limit outbound connectivity, refer to the list of endpoints for [the Connector and Cloud Volumes ONTAP](#).

- c. Set up a VPC endpoint to the S3 service.

A VPC endpoint is required if you want to tier cold data from Cloud Volumes ONTAP to low-cost object storage.

[Learn more about networking requirements.](#)



## Set up the AWS KMS

If you want to use Amazon encryption with Cloud Volumes ONTAP, then you need to ensure that an active Customer Master Key (CMK) exists. You also need to modify the key policy for each CMK by adding the IAM role that provides permissions to the Connector as a *key user*. [Learn more.](#)



## Launch Cloud Volumes ONTAP using Cloud Manager

Click **Add Working Environment**, select the type of system that you would like to deploy, and complete the steps in the wizard. [Read step-by-step instructions.](#)

### Related links

- [Evaluating](#)
- [Creating a Connector from Cloud Manager](#)
- [Launching a Connector from the AWS Marketplace](#)
- [Installing the Connector software on a Linux host](#)
- [What Cloud Manager does with AWS permissions](#)

## Planning your Cloud Volumes ONTAP configuration in AWS

When you deploy Cloud Volumes ONTAP in AWS, you can choose a preconfigured system that matches your workload requirements, or you can create your own configuration. If you choose your own configuration, you should understand the options available to you.

### Choosing a license type

Cloud Volumes ONTAP is available in two pricing options: pay-as-you-go and Bring Your Own License (BYOL). For pay-as-you-go, you can choose from three licenses: Explore, Standard, or Premium. Each license provides different capacity and compute options.

[Supported configurations for Cloud Volumes ONTAP 9.7 in AWS](#)

### Understanding storage limits

The raw capacity limit for a Cloud Volumes ONTAP system is tied to the license. Additional limits impact the size of aggregates and volumes. You should be aware of these limits as you plan your configuration.

[Storage limits for Cloud Volumes ONTAP 9.7 in AWS](#)

### Sizing your system in AWS

Sizing your Cloud Volumes ONTAP system can help you meet requirements for performance and capacity. You should be aware of a few key points when choosing an instance type, disk type, and disk size:

#### Instance type

- Match your workload requirements to the maximum throughput and IOPS for each EC2 instance type.
- If several users write to the system at the same time, choose an instance type that has enough CPUs to manage the requests.
- If you have an application that is mostly reads, then choose a system with enough RAM.
  - [AWS Documentation: Amazon EC2 Instance Types](#)
  - [AWS Documentation: Amazon EBS–Optimized Instances](#)

#### EBS disk type

General Purpose SSDs are the most common disk type for Cloud Volumes ONTAP. To view the use cases for EBS disks, refer to [AWS Documentation: EBS Volume Types](#).

#### EBS disk size

You need to choose an initial disk size when you launch a Cloud Volumes ONTAP system. After that, you can [let Cloud Manager manage a system's capacity for you](#), but if you want to [build aggregates yourself](#), be aware of the following:

- All disks in an aggregate must be the same size.
- The performance of EBS disks is tied to disk size. The size determines the baseline IOPS and maximum burst duration for SSD disks and the baseline and burst throughput for HDD disks.
- Ultimately, you should choose the disk size that gives you the *sustained performance* that you need.
- Even if you do choose larger disks (for example, six 4 TB disks), you might not get all of the IOPS because the EC2 instance can reach its bandwidth limit.



For more details about EBS disk performance, refer to [AWS Documentation: EBS Volume Types](#).

Watch the following video for more details about sizing your Cloud Volumes ONTAP system in AWS:



### Choosing a configuration that supports Flash Cache

Some Cloud Volumes ONTAP configurations in AWS include local NVMe storage, which Cloud Volumes ONTAP uses as *Flash Cache* for better performance. [Learn more about Flash Cache](#).

### AWS network information worksheet

When you launch Cloud Volumes ONTAP in AWS, you need to specify details about your VPC network. You can use a worksheet to collect the information from your administrator.

#### Network information for Cloud Volumes ONTAP

AWS information	Your value
Region	
VPC	
Subnet	
Security group (if using your own)	

#### Network information for an HA pair in multiple AZs



<b>AWS information</b>	<b>Your value</b>
Region	
VPC	
Security group (if using your own)	
Node 1 availability zone	
Node 1 subnet	
Node 2 availability zone	
Node 2 subnet	
Mediator availability zone	
Mediator subnet	
Key pair for the mediator	
Floating IP address for cluster management port	
Floating IP address for data on node 1	
Floating IP address for data on node 2	
Route tables for floating IP addresses	

## Choosing a write speed

Cloud Manager enables you to choose a write speed setting for single node Cloud Volumes ONTAP systems. Before you choose a write speed, you should understand the differences between the normal and high settings and risks and recommendations when using high write speed.

### Difference between normal write speed and high write speed

When you choose normal write speed, data is written directly to disk, thereby reducing the likelihood of data loss in the event of an unplanned system outage.

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

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

### When to use high write speed

High write speed is a good choice if fast write performance is required for your workload and you can withstand the risk of data loss in the event of an unplanned system outage.

## Recommendations when using high write speed

If you enable high write speed, you should ensure write protection at the application layer.

## Choosing a volume usage profile

ONTAP includes several storage efficiency features that can reduce the total amount of storage that you need. When you create a volume in Cloud Manager, you can choose a profile that enables these features or a profile that disables them. You should learn more about these features to help you decide which profile to use.

NetApp storage efficiency features provide the following benefits:

### Thin provisioning

Presents more logical storage to hosts or users than you actually have in your physical storage pool. Instead of preallocating storage space, storage space is allocated dynamically to each volume as data is written.

### Deduplication

Improves efficiency by locating identical blocks of data and replacing them with references to a single shared block. This technique reduces storage capacity requirements by eliminating redundant blocks of data that reside in the same volume.

### Compression

Reduces the physical capacity required to store data by compressing data within a volume on primary, secondary, and archive storage.

## Set up your networking

### Networking requirements for Cloud Volumes ONTAP in AWS

Set up your AWS networking so Cloud Volumes ONTAP systems can operate properly.

#### General requirements for Cloud Volumes ONTAP

The following requirements must be met in AWS.

#### Outbound internet access for Cloud Volumes ONTAP nodes

Cloud Volumes ONTAP nodes require outbound internet access to send messages to NetApp AutoSupport, which proactively monitors the health of your storage.

Routing and firewall policies must allow AWS HTTP/HTTPS traffic to the following endpoints so Cloud Volumes ONTAP can send AutoSupport messages:

- <https://support.netapp.com/aods/asupmessage>
- <https://support.netapp.com/asupprod/post/1.0/postAsup>

If you have a NAT instance, you must define an inbound security group rule that allows HTTPS traffic from the private subnet to the internet.

[Learn how to configure AutoSupport.](#)

## Outbound internet access for the HA mediator

The HA mediator instance must have an outbound connection to the AWS EC2 service so it can assist with storage failover. To provide the connection, you can add a public IP address, specify a proxy server, or use a manual option.

The manual option can be a NAT gateway or an interface VPC endpoint from the target subnet to the AWS EC2 service. For details about VPC endpoints, refer to [AWS Documentation: Interface VPC Endpoints \(AWS PrivateLink\)](#).

## Number of IP addresses

Cloud Manager allocates the following number of IP addresses to Cloud Volumes ONTAP in AWS:

- Single node: 6 IP addresses
- HA pairs in single AZs: 15 addresses
- HA pairs in multiple AZs: 15 or 16 IP addresses

Note that Cloud Manager creates an SVM management LIF on single node systems, but not on HA pairs in a single AZ. You can choose whether to create an SVM management LIF on HA pairs in multiple AZs.



A LIF is an IP address associated with a physical port. An SVM management LIF is required for management tools like SnapCenter.

## Security groups

You do not need to create security groups because Cloud Manager does that for you. If you need to use your own, refer to [Security group rules](#).

## Connection from Cloud Volumes ONTAP to AWS S3 for data tiering

If you want to use EBS as a performance tier and AWS S3 as a capacity tier, you must ensure that Cloud Volumes ONTAP has a connection to S3. The best way to provide that connection is by creating a VPC Endpoint to the S3 service. For instructions, see [AWS Documentation: Creating a Gateway Endpoint](#).

When you create the VPC Endpoint, be sure to select the region, VPC, and route table that corresponds to the Cloud Volumes ONTAP instance. You must also modify the security group to add an outbound HTTPS rule that enables traffic to the S3 endpoint. Otherwise, Cloud Volumes ONTAP cannot connect to the S3 service.

If you experience any issues, see [AWS Support Knowledge Center: Why can't I connect to an S3 bucket using a gateway VPC endpoint?](#)

## Connections to ONTAP systems in other networks

To replicate data between a Cloud Volumes ONTAP system in AWS and ONTAP systems in other networks, you must have a VPN connection between the AWS VPC and the other network—for example, an Azure VNet or your corporate network. For instructions, see [AWS Documentation: Setting Up an AWS VPN Connection](#).

## DNS and Active Directory for CIFS

If you want to provision CIFS storage, you must set up DNS and Active Directory in AWS or extend your on-premises setup to AWS.

The DNS server must provide name resolution services for the Active Directory environment. You can configure DHCP option sets to use the default EC2 DNS server, which must not be the DNS server used by

the Active Directory environment.

For instructions, refer to [AWS Documentation: Active Directory Domain Services on the AWS Cloud: Quick Start Reference Deployment](#).

### Requirements for HA pairs in multiple AZs

Additional AWS networking requirements apply to Cloud Volumes ONTAP HA configurations that use multiple Availability Zones (AZs). You should review these requirements before you launch an HA pair because you must enter the networking details in Cloud Manager.

To understand how HA pairs work, see [High-availability pairs](#).

### Availability Zones

This HA deployment model uses multiple AZs to ensure high availability of your data. You should use a dedicated AZ for each Cloud Volumes ONTAP instance and the mediator instance, which provides a communication channel between the HA pair.

### Floating IP addresses for NAS data and cluster/SVM management

HA configurations in multiple AZs use floating IP addresses that migrate between nodes if failures occur. They are not natively accessible from outside the VPC, unless you [set up an AWS transit gateway](#).

One floating IP address is for cluster management, one is for NFS/CIFS data on node 1, and one is for NFS/CIFS data on node 2. A fourth floating IP address for SVM management is optional.




A floating IP address is required for the SVM management LIF if you use SnapDrive for Windows or SnapCenter with the HA pair. If you don't specify the IP address when you deploy the system, you can create the LIF later. For details, see [Setting up Cloud Volumes ONTAP](#).

You need to enter the floating IP addresses in Cloud Manager when you create a Cloud Volumes ONTAP HA working environment. Cloud Manager allocates the IP addresses to the HA pair when it launches the system.

The floating IP addresses must be outside of the CIDR blocks for all VPCs in the AWS region in which you deploy the HA configuration. Think of the floating IP addresses as a logical subnet that's outside of the VPCs in your region.

The following example shows the relationship between floating IP addresses and the VPCs in an AWS region. While the floating IP addresses are outside the CIDR blocks for all VPCs, they're routable to subnets through route tables.

## AWS region



Cloud Manager automatically creates static IP addresses for iSCSI access and for NAS access from clients outside the VPC. You don't need to meet any requirements for these types of IP addresses.

### Transit gateway to enable floating IP access from outside the VPC

[Set up an AWS transit gateway](#) to enable access to an HA pair's floating IP addresses from outside the VPC where the HA pair resides.

### Route tables

After you specify the floating IP addresses in Cloud Manager, you need to select the route tables that should include routes to the floating IP addresses. This enables client access to the HA pair.

If you have just one route table for the subnets in your VPC (the main route table), then Cloud Manager automatically adds the floating IP addresses to that route table. If you have more than one route table, it's very important to select the correct route tables when launching the HA pair. Otherwise, some clients might not have access to Cloud Volumes ONTAP.

For example, you might have two subnets that are associated with different route tables. If you select route table A, but not route table B, then clients in the subnet associated with route table A can access the HA

pair, but clients in the subnet associated with route table B can't.

For more information about route tables, refer to [AWS Documentation: Route Tables](#).


### **Connection to NetApp management tools**

To use NetApp management tools with HA configurations that are in multiple AZs, you have two connection options:

1. Deploy the NetApp management tools in a different VPC and [set up an AWS transit gateway](#). The gateway enables access to the floating IP address for the cluster management interface from outside the VPC.
2. Deploy the NetApp management tools in the same VPC with a similar routing configuration as NAS clients.

### **Example HA configuration**

The following image shows an optimal HA configuration in AWS operating as an active-passive configuration:



### Requirements for the Connector

Set up your networking so that the Connector can manage resources and processes within your public cloud environment. The most important step is ensuring outbound internet access to various endpoints.



If your network uses a proxy server for all communication to the internet, you can specify the proxy server from the Settings page. Refer to [Configuring the Connector to use a proxy server](#).

### Connection to target networks

A Connector requires a network connection to the VPCs and VNets in which you want to deploy Cloud Volumes ONTAP.

For example, if you install a Connector in your corporate network, then you must set up a VPN connection to the VPC or VNet in which you launch Cloud Volumes ONTAP.

### Outbound internet access

The Connector requires outbound internet access to manage resources and processes within your public cloud environment. A Connector contacts the following endpoints when managing resources in AWS:

Endpoints	Purpose
<p>AWS services (amazonaws.com):</p> <ul style="list-style-type: none"> <li>• CloudFormation</li> <li>• Elastic Compute Cloud (EC2)</li> <li>• Key Management Service (KMS)</li> <li>• Security Token Service (STS)</li> <li>• Simple Storage Service (S3)</li> </ul> <p>The exact endpoint depends on the region in which you deploy Cloud Volumes ONTAP.  <a href="#">Refer to AWS documentation for details.</a></p>	Enables Cloud Manager to deploy and manage Cloud Volumes ONTAP in AWS.
https://api.services.cloud.netapp.com:443	API requests to NetApp Cloud Central.
https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com	Provides access to software images, manifests, and templates.
https://repo.cloud.support.netapp.com	Used to download Cloud Manager dependencies.
http://repo.mysql.com/	Used to download MySQL.
<p>https://cognito-idp.us-east-1.amazonaws.com</p> <p>https://cognito-identity.us-east-1.amazonaws.com</p> <p>https://sts.amazonaws.com</p> <p>https://cloud-support-netapp-com-accelerated.s3.amazonaws.com</p>	Enables Cloud Manager to access and download manifests, templates, and Cloud Volumes ONTAP upgrade images.
https://cloudmanagerinfraprod.azurecr.io	Access to software images of container components for an infrastructure that's running Docker and provides a solution for service integrations with Cloud Manager.
https://kinesis.us-east-1.amazonaws.com	Enables NetApp to stream data from audit records.
https://cloudmanager.cloud.netapp.com	Communication with the Cloud Manager service, which includes Cloud Central accounts.
https://netapp-cloud-account.auth0.com	Communication with NetApp Cloud Central for centralized user authentication.
https://w86yt021u5.execute-api.us-east-1.amazonaws.com/production/whitelist	Used to add your AWS account ID to the list of allowed users for Backup to S3.



Endpoints	Purpose
<a href="https://support.netapp.com/aods/asupmessage">https://support.netapp.com/aods/asupmessage</a> <a href="https://support.netapp.com/asupprod/post/1.0/postAsup">https://support.netapp.com/asupprod/post/1.0/postAsup</a>	Communication with NetApp AutoSupport.
<a href="https://support.netapp.com/svcgw">https://support.netapp.com/svcgw</a> <a href="https://support.netapp.com/ServiceGW/entitlement">https://support.netapp.com/ServiceGW/entitlement</a> <a href="https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com">https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com</a> <a href="https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com">https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com</a>	Communication with NetApp for system licensing and support registration.
<a href="https://ipa-signer.cloudmanager.netapp.com">https://ipa-signer.cloudmanager.netapp.com</a>	Enables Cloud Manager to generate licenses (for example, a FlexCache license for Cloud Volumes ONTAP)
<a href="https://packages.cloud.google.com/yum">https://packages.cloud.google.com/yum</a> <a href="https://github.com/NetApp/trident/releases/download/">https://github.com/NetApp/trident/releases/download/</a>	Required to connect Cloud Volumes ONTAP systems with a Kubernetes cluster. The endpoints enable installation of NetApp Trident.
<p>Various third-party locations, for example:</p> <ul style="list-style-type: none"> <li>• <a href="https://repo1.maven.org/maven2">https://repo1.maven.org/maven2</a></li> <li>• <a href="https://oss.sonatype.org/content/repositories">https://oss.sonatype.org/content/repositories</a></li> <li>• <a href="https://repo.typesafe.org">https://repo.typesafe.org</a></li> </ul> <p>Third-party locations are subject to change.</p>	During upgrades, Cloud Manager downloads the latest packages for third-party dependencies.

While you should perform almost all tasks from the SaaS user interface, a local user interface is still available on the Connector. The machine running the web browser must have connections to the following endpoints:

Endpoints	Purpose
The Connector host	<p>You must enter the host's IP address from a web browser to load the Cloud Manager console.</p> <p>Depending on your connectivity to your cloud provider, you can use the private IP or a public IP assigned to the host:</p> <ul style="list-style-type: none"> <li>• A private IP works if you have a VPN and direct connect access to your virtual network</li> <li>• A public IP works in any networking scenario</li> </ul> <p>In any case, you should secure network access by ensuring that security group rules allow access from only authorized IPs or subnets.</p>

Endpoints	Purpose
https://auth0.com https://cdn.auth0.com https://netapp-cloud-account.auth0.com https://services.cloud.netapp.com	Your web browser connects to these endpoints for centralized user authentication through NetApp Cloud Central.
https://widget.intercom.io	For in-product chat that enables you to talk to NetApp cloud experts.

### Setting up an AWS transit gateway for HA pairs in multiple AZs


Set up an AWS transit gateway to enable access to an HA pair's [floating IP addresses](#) from outside the VPC where the HA pair resides.

When a Cloud Volumes ONTAP HA configuration is spread across multiple AWS Availability Zones, floating IP addresses are required for NAS data access from within the VPC. These floating IP addresses can migrate between nodes when failures occur, but they are not natively accessible from outside the VPC. Separate private IP addresses provide data access from outside the VPC, but they don't provide automatic failover.

Floating IP addresses are also required for the cluster management interface and the optional SVM management LIF.

If you set up an AWS transit gateway, you enable access to the floating IP addresses from outside the VPC where the HA pair resides. That means NAS clients and NetApp management tools outside the VPC can access the floating IPs.

Here's an example that shows two VPCs connected by a transit gateway. An HA system resides in one VPC, while a client resides in the other. You could then mount a NAS volume on the client using the floating IP address.



The following steps illustrate how to set up a similar configuration.

### Steps

1. [Create a transit gateway and attach the VPCs to the gateway.](#)
2. Create routes in the transit gateway's route table by specifying the HA pair's floating IP addresses.

You can find the floating IP addresses on the Working Environment Information page in Cloud Manager. Here's an example:

## NFS & CIFS access from within the VPC using Floating IP

### Auto failover

Cluster Management : 172.23.0.1

Data (nfs,cifs) : Node 1: 172.23.0.2 | Node 2: 172.23.0.3

### Access

SVM Management : 172.23.0.4

The following sample image shows the route table for the transit gateway. It includes routes to the CIDR blocks of the two VPCs and four floating IP addresses used by Cloud Volumes ONTAP.

Transit Gateway Route Table: tgw-rtb-0ea8ee291c7aeddd3

Details Associations Propagations **Routes** Tags

The table below will return a maximum of 1000 routes. Narrow the filter or use export routes to view more routes.

Create route Replace route Delete route

Filter by attributes or search by keyword

<input type="checkbox"/>	CIDR	Attachment	Resource type	Route type	Route state
<input type="checkbox"/>	10.100.0.0/16	tgw-attach-05e77bd34e2ff91f8   vpc-0b2bc30e0dc8e0db1	VPC2	propagated	active
<input type="checkbox"/>	10.160.0.0/20	tgw-attach-00eba3eac3250d7db   vpc-673ae603	VPC1	propagated	active
<input type="checkbox"/>	172.23.0.1/32	tgw-attach-00eba3eac3250d7db   vpc-673ae603	VPC	static	active
<input type="checkbox"/>	172.23.0.2/32	tgw-attach-00eba3eac3250d7db   vpc-673ae603	VPC	static	active
<input type="checkbox"/>	172.23.0.3/32	tgw-attach-00eba3eac3250d7db   vpc-673ae603	VPC	static	active
<input type="checkbox"/>	172.23.0.4/32	tgw-attach-00eba3eac3250d7db   vpc-673ae603	VPC	static	active

**Floating IP Addresses**

3. Modify the route table of VPCs that need to access the floating IP addresses.
  - a. Add route entries to the floating IP addresses.
  - b. Add a route entry to the CIDR block of the VPC where the HA pair resides.

The following sample image shows the route table for VPC 2, which includes routes to VPC 1 and the floating IP addresses.

Route Table: rtb-0569a1bd740ed033f

Summary Routes Subnet Associations Route Propagation Tags

Edit routes

View All routes

Destination	Target	Status	Propagated
10.100.0.0/16	local	active	No
0.0.0.0/0	igw-07250bd01781e67df	active	No
10.160.0.0/20	tgw-015b7c249661ac279	active	No
172.23.0.1/32	tgw-015b7c249661ac279	active	No
172.23.0.2/32	tgw-015b7c249661ac279	active	No
172.23.0.3/32	tgw-015b7c249661ac279	active	No
172.23.0.4/32	tgw-015b7c249661ac279	active	No

VPC1  
Floating IP Addresses

- Modify the route table for the HA pair's VPC by adding a route to the VPC that needs access to the floating IP addresses.

This step is important because it completes the routing between the VPCs.

The following sample image shows the route table for VPC 1. It includes a route to the floating IP addresses and to VPC 2, which is where a client resides. Cloud Manager automatically added the floating IPs to the route table when it deployed the HA pair.

Summary Routes Subnet Associations Route Propagation Tags

Edit routes

View All routes

Destination	Target	Status
10.160.0.0/20	local	active
pl-68a54001 (com.amazonaws.us-west-2.s3, 54.231.160.0/19, 52.218.128.0/17, 52.92.32.0/22)	vpce-cb51a0a2	active
0.0.0.0/0	igw-b2182dd7	active
10.60.29.0/25	pcx-589c3331	active
10.100.0.0/16	tgw-015b7c249661ac279	active
10.129.0.0/20	pcx-ff7e1396	active
172.23.0.1/32	eni-0854d4715559c3cdb	active
172.23.0.2/32	eni-0854d4715559c3cdb	active
172.23.0.3/32	eni-0f76681216c3108ed	active
172.23.0.4/32	eni-0854d4715559c3cdb	active


VPC2  
Floating IP Addresses

- Mount volumes to clients using the floating IP address.

You can find the correct IP address in Cloud Manager by selecting a volume and clicking **Mount Command**.

# Volumes

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



## Related links

- [High-availability pairs in AWS](#)
- [Networking requirements for Cloud Volumes ONTAP in AWS](#)

## Security group rules for AWS

Cloud Manager creates AWS security groups that include the inbound and outbound rules that the Connector and Cloud Volumes ONTAP need to operate successfully. You might want to refer to the ports for testing purposes or if you prefer your to use own security groups.

### Rules for Cloud Volumes ONTAP

The security group for Cloud Volumes ONTAP requires both inbound and outbound rules.

### Inbound rules

The source for inbound rules in the predefined security group is 0.0.0.0/0.

Protocol	Port	Purpose
All ICMP	All	Pinging the instance
HTTP	80	HTTP access to the System Manager web console using the IP address of the cluster management LIF
HTTPS	443	HTTPS access to the System Manager web console using the IP address of the cluster management LIF
SSH	22	SSH access to the IP address of the cluster management LIF or a node management LIF
TCP	111	Remote procedure call for NFS
TCP	139	NetBIOS service session for CIFS

Protocol	Port	Purpose
TCP	161-162	Simple network management protocol
TCP	445	Microsoft SMB/CIFS over TCP with NetBIOS framing
TCP	635	NFS mount
TCP	749	Kerberos
TCP	2049	NFS server daemon
TCP	3260	iSCSI access through the iSCSI data LIF
TCP	4045	NFS lock daemon
TCP	4046	Network status monitor for NFS
TCP	10000	Backup using NDMP
TCP	11104	Management of intercluster communication sessions for SnapMirror
TCP	11105	SnapMirror data transfer using intercluster LIFs
UDP	111	Remote procedure call for NFS
UDP	161-162	Simple network management protocol
UDP	635	NFS mount
UDP	2049	NFS server daemon
UDP	4045	NFS lock daemon
UDP	4046	Network status monitor for NFS
UDP	4049	NFS rquotad protocol

## Outbound rules

The predefined security group for Cloud Volumes ONTAP opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

### Basic outbound rules

The predefined security group for Cloud Volumes ONTAP includes the following outbound rules.

Protocol	Port	Purpose
All ICMP	All	All outbound traffic
All TCP	All	All outbound traffic
All UDP	All	All outbound traffic

### Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by Cloud Volumes ONTAP.



The source is the interface (IP address) on the Cloud Volumes ONTAP system.

Service	Protocol	Port	Source	Destination	Purpose
Active Directory	TCP	88	Node management LIF	Active Directory forest	Kerberos V authentication
	UDP	137	Node management LIF	Active Directory forest	NetBIOS name service
	UDP	138	Node management LIF	Active Directory forest	NetBIOS datagram service
	TCP	139	Node management LIF	Active Directory forest	NetBIOS service session
	TCP & UDP	389	Node management LIF	Active Directory forest	LDAP
	TCP	445	Node management LIF	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	TCP	464	Node management LIF	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	UDP	464	Node management LIF	Active Directory forest	Kerberos key administration
	TCP	749	Node management LIF	Active Directory forest	Kerberos V change & set Password (RPCSEC_GSS)
	TCP	88	Data LIF (NFS, CIFS, iSCSI)	Active Directory forest	Kerberos V authentication
	UDP	137	Data LIF (NFS, CIFS)	Active Directory forest	NetBIOS name service
	UDP	138	Data LIF (NFS, CIFS)	Active Directory forest	NetBIOS datagram service
	TCP	139	Data LIF (NFS, CIFS)	Active Directory forest	NetBIOS service session
	TCP & UDP	389	Data LIF (NFS, CIFS)	Active Directory forest	LDAP
	TCP	445	Data LIF (NFS, CIFS)	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	TCP	464	Data LIF (NFS, CIFS)	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	UDP	464	Data LIF (NFS, CIFS)	Active Directory forest	Kerberos key administration
	TCP	749	Data LIF (NFS, CIFS)	Active Directory forest	Kerberos V change & set password (RPCSEC_GSS)
	Backup to S3	TCP	5010	Intercluster LIF	Backup endpoint or restore endpoint



Service	Protocol	Port	Source	Destination	Purpose
Cluster	All traffic	All traffic	All LIFs on one node	All LIFs on the other node	Intercluster communications (Cloud Volumes ONTAP HA only)
	TCP	3000	Node management LIF	HA mediator	ZAPI calls (Cloud Volumes ONTAP HA only)
	ICMP	1	Node management LIF	HA mediator	Keep alive (Cloud Volumes ONTAP HA only)
DHCP	UDP	68	Node management LIF	DHCP	DHCP client for first-time setup
DHCPS	UDP	67	Node management LIF	DHCP	DHCP server
DNS	UDP	53	Node management LIF and data LIF (NFS, CIFS)	DNS	DNS
NDMP	TCP	1860-18699	Node management LIF	Destination servers	NDMP copy
SMTP	TCP	25	Node management LIF	Mail server	SMTP alerts, can be used for AutoSupport
SNMP	TCP	161	Node management LIF	Monitor server	Monitoring by SNMP traps
	UDP	161	Node management LIF	Monitor server	Monitoring by SNMP traps
	TCP	162	Node management LIF	Monitor server	Monitoring by SNMP traps
	UDP	162	Node management LIF	Monitor server	Monitoring by SNMP traps
SnapMirror	TCP	11104	Intercluster LIF	ONTAP intercluster LIFs	Management of intercluster communication sessions for SnapMirror
	TCP	11105	Intercluster LIF	ONTAP intercluster LIFs	SnapMirror data transfer
Syslog	UDP	514	Node management LIF	Syslog server	Syslog forward messages

### Rules for the HA mediator external security group

The predefined external security group for the Cloud Volumes ONTAP HA mediator includes the following inbound and outbound rules.

#### Inbound rules

The source for inbound rules is 0.0.0.0/0.

Protocol	Port	Purpose
SSH	22	SSH connections to the HA mediator
TCP	3000	RESTful API access from the Connector

## Outbound rules

The predefined security group for the HA mediator opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

### Basic outbound rules

The predefined security group for the HA mediator includes the following outbound rules.

Protocol	Port	Purpose
All TCP	All	All outbound traffic
All UDP	All	All outbound traffic

### Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by the HA mediator.

Protocol	Port	Destination	Purpose
HTTP	80	Connector IP address	Download upgrades for the mediator
HTTPS	443	AWS API services	Assist with storage failover
UDP	53	AWS API services	Assist with storage failover



Rather than open ports 443 and 53, you can create an interface VPC endpoint from the target subnet to the AWS EC2 service.

### Rules for the HA mediator internal security group

The predefined internal security group for the Cloud Volumes ONTAP HA mediator includes the following rules. Cloud Manager always creates this security group. You do not have the option to use your own.

### Inbound rules

The predefined security group includes the following inbound rules.

Protocol	Port	Purpose
All traffic	All	Communication between the HA mediator and HA nodes

## Outbound rules

The predefined security group includes the following outbound rules.

Protocol	Port	Purpose
All traffic	All	Communication between the HA mediator and HA nodes

## Rules for the Connector

The security group for the Connector requires both inbound and outbound rules.

## Inbound rules

The source for inbound rules in the predefined security group is 0.0.0.0/0.

Protocol	Port	Purpose
SSH	22	Provides SSH access to the Connector host
HTTP	80	Provides HTTP access from client web browsers to the local user interface and connections from Cloud Compliance
HTTPS	443	Provides HTTPS access from client web browsers to the local user interface
TCP	3128	Provides the Cloud Compliance instance with internet access, if your AWS network doesn't use a NAT or proxy

## Outbound rules

The predefined security group for the Connector opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

## Basic outbound rules

The predefined security group for the Connector includes the following outbound rules.

Protocol	Port	Purpose
All TCP	All	All outbound traffic
All UDP	All	All outbound traffic

## Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by the Connector.



The source IP address is the Connector host.

Service	Protocol	Port	Destination	Purpose
Active Directory	TCP	88	Active Directory forest	Kerberos V authentication
	TCP	139	Active Directory forest	NetBIOS service session
	TCP	389	Active Directory forest	LDAP
	TCP	445	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	TCP	464	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	TCP	749	Active Directory forest	Active Directory Kerberos V change & set password (RPCSEC_GSS)
	UDP	137	Active Directory forest	NetBIOS name service
	UDP	138	Active Directory forest	NetBIOS datagram service
	UDP	464	Active Directory forest	Kerberos key administration
API calls and AutoSupport	HTTP	443	Outbound internet and ONTAP cluster management LIF	API calls to AWS and ONTAP, and sending AutoSupport messages to NetApp
API calls	TCP	3000	ONTAP cluster management LIF	API calls to ONTAP
	TCP	8088	Backup to S3	API calls to Backup to S3
DNS	UDP	53	DNS	Used for DNS resolve by Cloud Manager
Cloud Compliance	HTTP	80	Cloud Compliance instance	Cloud Compliance for Cloud Volumes ONTAP

## Setting up the AWS KMS

If you want to use Amazon encryption with Cloud Volumes ONTAP, then you need to set up the AWS Key Management Service (KMS).

### Steps

1. Ensure that an active Customer Master Key (CMK) exists.

The CMK can be an AWS-managed CMK or a customer-managed CMK. It can be in the same AWS account as Cloud Manager and Cloud Volumes ONTAP or in a different AWS account.

[AWS Documentation: Customer Master Keys \(CMKs\)](#)

2. Modify the key policy for each CMK by adding the IAM role that provides permissions to Cloud Manager as a *key user*.

Adding the IAM role as a key user gives Cloud Manager permissions to use the CMK with Cloud Volumes ONTAP.

[AWS Documentation: Editing Keys](#)


3. If the CMK is in a different AWS account, complete the following steps:

- a. Go to the KMS console from the account where the CMK resides.
- b. Select the key.
- c. In the **General configuration** pane, copy the ARN of the key.

You'll need to provide the ARN to Cloud Manager when you create the Cloud Volumes ONTAP system.

- d. In the **Other AWS accounts** pane, add the AWS account that provides Cloud Manager with permissions.

In most cases, this is the account where Cloud Manager resides. If Cloud Manager wasn't installed in AWS, it would be the account for which you provided AWS access keys to Cloud Manager.



- e. Now switch to the AWS account that provides Cloud Manager with permissions and open the IAM console.
- f. Create an IAM policy that includes the permissions listed below.
- g. Attach the policy to the IAM role or IAM user that provides permissions to Cloud Manager.

The following policy provides the permissions that Cloud Manager needs to use the CMK from the external AWS account. Be sure to modify the region and account ID in the "Resource" sections.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowUseOfTheKey",
      "Effect": "Allow",
      "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:DescribeKey"
      ],
      "Resource": [
        "arn:aws:kms:us-east-1:externalaccountid:key/externalkeyid"
      ]
    },
    {
      "Sid": "AllowAttachmentOfPersistentResources",
      "Effect": "Allow",
      "Action": [
        "kms:CreateGrant",
        "kms:ListGrants",
        "kms:RevokeGrant"
      ],
      "Resource": [
        "arn:aws:kms:us-east-1:externalaccountid:key/externalaccountid"
      ],
      "Condition": {
        "Bool": {
          "kms:GrantIsForAWSResource": true
        }
      }
    }
  ]
}

```

For additional details about this process, see [AWS Documentation: Allowing External AWS Accounts to Access a CMK](#).

## Launching Cloud Volumes ONTAP in AWS

You can launch Cloud Volumes ONTAP in a single-system configuration or as an HA pair

in AWS.

## Launching a single-node Cloud Volumes ONTAP system in AWS

If you want to launch Cloud Volumes ONTAP in AWS, you need to create a new working environment in Cloud Manager.

### Before you begin

- You should have a [Connector that is associated with your workspace](#).



You must be an Account Admin to create a Connector. When you create your first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to create a Connector if you don't have one yet.

- [You should be prepared to leave the Connector running at all times](#).
- You should have prepared by choosing a configuration and by obtaining AWS networking information from your administrator. For details, see [Planning your Cloud Volumes ONTAP configuration](#).
- If you want to launch a BYOL system, you must have the 20-digit serial number (license key).
- If you want to use CIFS, you must have set up DNS and Active Directory. For details, see [Networking requirements for Cloud Volumes ONTAP in AWS](#).

### About this task

Immediately after you create the working environment, Cloud Manager launches a test instance in the specified VPC to verify connectivity. If successful, Cloud Manager immediately terminates the instance and then starts deploying the Cloud Volumes ONTAP system. If Cloud Manager cannot verify connectivity, creation of the working environment fails. The test instance is either a t2.nano (for default VPC tenancy) or m3.medium (for dedicated VPC tenancy).

### Steps

1. On the Working Environments page, click **Add Working Environment** and follow the prompts.
2. **Choose a Location:** Select **Amazon Web Services** and **Cloud Volumes ONTAP Single Node**.
3. **Details and Credentials:** Optionally change the AWS credentials and subscription, enter a working environment name, add tags if needed, and then enter a password.

Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Working Environment Name	Cloud Manager uses the working environment name to name both the Cloud Volumes ONTAP system and the Amazon EC2 instance. It also uses the name as the prefix for the predefined security group, if you select that option.

Field	Description
Add tags	<p>AWS tags are metadata for your AWS resources. Cloud Manager adds the tags to the Cloud Volumes ONTAP instance and each AWS resource associated with the instance.</p> <p>You can add up to four tags from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four tags when creating a working environment.</p> <p>For information about tags, refer to <a href="#">AWS Documentation: Tagging your Amazon EC2 Resources</a>.</p>
User name and password	<p>These are the credentials for the Cloud Volumes ONTAP cluster admin account. You can use these credentials to connect to Cloud Volumes ONTAP through OnCommand System Manager or its CLI.</p>
Edit Credentials	<p>Choose the AWS credentials and marketplace subscription to use with this Cloud Volumes ONTAP system.</p> <p>Click <b>Add Subscription</b> to associate the selected credentials with a subscription.</p> <p>To create a pay-as-you-go Cloud Volumes ONTAP system, you need to select AWS credentials that are associated with a subscription to Cloud Volumes ONTAP from the AWS Marketplace. You'll be charged from this subscription for every Cloud Volumes ONTAP 9.6 and later PAYGO system that you create and each add-on feature that you enable.</p> <p><a href="#">Learn how to add additional AWS credentials to Cloud Manager.</a></p>

The following video shows how to associate a pay-as-you-go Marketplace subscription to your AWS credentials:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_aws.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_aws.mp4) (video)

If multiple IAM users work in the same AWS account, then each user needs to subscribe. After the first user subscribes, the AWS Marketplace informs subsequent users that they're already subscribed, as shown in the image below. While a subscription is in place for the AWS *account*, each IAM user needs to associate themselves with that subscription. If you see the message shown below, click the **click here** link to go to Cloud Central and complete the process.



### Cloud Manager (for Cloud Volumes ONTAP)

You are currently subscribed to this product and will be charged for your accumulated usage at the end of your next billing cycle, based on the costs listed in Pricing information on the right.

?

**Having issues signing up for your product?**  
If you were unable to complete the set-up process for this software, please [click here](#) to be taken to the product's registration area.

Subscribe

You are already subscribed to this product

---

**Pricing Details**

Software Fees

- Services:** Keep the services enabled or disable the individual services that you don't want to use with Cloud Volumes ONTAP.



- [Learn more about Cloud Compliance.](#)
- [Learn more about Backup to Cloud.](#)
- [Learn more about Monitoring.](#)

5. **Location & Connectivity:** Enter the network information that you recorded in the AWS worksheet.

The following image shows the page filled out:

Location	Connectivity
AWS Region <input type="text" value="US West   Oregon"/>	Security Group <input checked="" type="radio"/> Generated security group <input type="radio"/> Use existing security group
VPC <input type="text" value="vpc-3a01e05f - 172.31.0.0/16"/>	SSH Authentication Method <input checked="" type="radio"/> Password <input type="radio"/> Key Pair
Subnet <input type="text" value="172.31.5.0/24 (OCCM subnet)"/>	

6. **Data Encryption:** Choose no data encryption or AWS-managed encryption.

For AWS-managed encryption, you can choose a different Customer Master Key (CMK) from your account or another AWS account.



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

[Learn how to set up the AWS KMS for Cloud Volumes ONTAP.](#)

[Learn more about supported encryption technologies.](#)

7. **License and Support Site Account:** Specify whether you want to use pay-as-you-go or BYOL, and then specify a NetApp Support Site account.

To understand how licenses work, see [Licensing](#).

A NetApp Support Site Account is optional for pay-as-you-go, but required for BYOL systems. [Learn how to add NetApp Support Site accounts.](#)

8. **Preconfigured Packages:** Select one of the packages to quickly launch Cloud Volumes ONTAP, or click **Create my own configuration**.

If you choose one of the packages, you only need to specify a volume and then review and approve the configuration.

9. **IAM Role:** You should keep the default option to let Cloud Manager create the role for you.

If you prefer to use your own policy, it must meet [policy requirements for Cloud Volumes ONTAP nodes](#).

10. **Licensing:** Change the Cloud Volumes ONTAP version as needed, select a license, an instance type, and the instance tenancy.

If your needs change after you launch the instance, you can modify the license or instance type later.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then Cloud Manager updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.6 RC1 and 9.6 GA is available. The update does not occur from one release to another—for example, from 9.6 to 9.7.

11. **Underlying Storage Resources:** Choose settings for the initial aggregate: a disk type, a size for each disk, and whether data tiering should be enabled.

Note the following:

- The disk type is for the initial volume. You can choose a different disk type for subsequent volumes.
- The disk size is for all disks in the initial aggregate and for any additional aggregates that Cloud Manager creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.

For help choosing a disk type and size, see [Sizing your system in AWS](#).

- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates.

[Learn how data tiering works.](#)

12. **Write Speed & WORM:** Choose **Normal** or **High** write speed, and activate write once, read many (WORM) storage, if desired.

Choosing a write speed is supported with single node systems only.

[Learn more about write speed.](#)

WORM can't be enabled if data tiering was enabled.

[Learn more about WORM storage.](#)

13. **Create Volume:** Enter details for the new volume or click **Skip**.

Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, Cloud Manager enters a value that provides access to all instances in the subnet.

Field	Description
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, Cloud Manager automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

### Volume Details, Protection & Protocol

#### Details & Protection

Volume Name:  Size (GB):

Snapshot Policy:

Default Policy

#### Protocol

NFS   
 CIFS   
 iSCSI

Share name:  Permissions:

Users / Groups:

Valid users and groups separated by a semicolon

14. **CIFS Setup:** If you chose the CIFS protocol, set up a CIFS server.

Field	Description
DNS Primary and Secondary IP Address	The IP addresses of the DNS servers that provide name resolution for the CIFS server. The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers. If you configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=Computers,OU=corp</b> in this field.
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. See the <a href="#">Cloud Manager API Developer Guide</a> for details.

15. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and edit the volume tiering policy, if needed.

For more information, see [Understanding volume usage profiles](#) and [Data tiering overview](#).

16. **Review & Approve:** Review and confirm your selections.
- Review details about the configuration.
  - Click **More information** to review details about support and the AWS resources that Cloud Manager will purchase.
  - Select the **I understand...** check boxes.
  - Click **Go**.

## Result

Cloud Manager launches the Cloud Volumes ONTAP instance. You can track the progress in the timeline.

If you experience any issues launching the Cloud Volumes ONTAP instance, review the failure message. You can also select the working environment and click Re-create environment.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

## After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.
- If you want to apply quotas to volumes, use System Manager or the CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

## Launching a Cloud Volumes ONTAP HA pair in AWS

If you want to launch a Cloud Volumes ONTAP HA pair in AWS, you need to create an HA working environment in Cloud Manager.

### Before you begin

- You should have a [Connector that is associated with your workspace](#).



You must be an Account Admin to create a Connector. When you create your first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to create a Connector if you don't have one yet.

- [You should be prepared to leave the Connector running at all times](#).
- You should have prepared by choosing a configuration and by obtaining AWS networking information from your administrator. For details, see [Planning your Cloud Volumes ONTAP configuration](#).
- If you purchased BYOL licenses, you must have a 20-digit serial number (license key) for each node.
- If you want to use CIFS, you must have set up DNS and Active Directory. For details, see [Networking requirements for Cloud Volumes ONTAP in AWS](#).

### Limitation

At this time, HA pairs are not supported with AWS Outposts.

### About this task

Immediately after you create the working environment, Cloud Manager launches a test instance in the specified VPC to verify connectivity. If successful, Cloud Manager immediately terminates the instance and then starts deploying the Cloud Volumes ONTAP system. If Cloud Manager cannot verify connectivity, creation of the working environment fails. The test instance is either a t2.nano (for default VPC tenancy) or m3.medium (for dedicated VPC tenancy).

### Steps

1. On the Working Environments page, click **Add Working Environment** and follow the prompts.
2. **Choose a Location:** Select **Amazon Web Services** and **Cloud Volumes ONTAP Single Node**.
3. **Details and Credentials:** Optionally change the AWS credentials and subscription, enter a working environment name, add tags if needed, and then enter a password.

Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Working Environment Name	Cloud Manager uses the working environment name to name both the Cloud Volumes ONTAP system and the Amazon EC2 instance. It also uses the name as the prefix for the predefined security group, if you select that option.

Field	Description
Add tags	<p>AWS tags are metadata for your AWS resources. Cloud Manager adds the tags to the Cloud Volumes ONTAP instance and each AWS resource associated with the instance.</p> <p>You can add up to four tags from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four tags when creating a working environment.</p> <p>For information about tags, refer to <a href="#">AWS Documentation: Tagging your Amazon EC2 Resources</a>.</p>
User name and password	<p>These are the credentials for the Cloud Volumes ONTAP cluster admin account. You can use these credentials to connect to Cloud Volumes ONTAP through OnCommand System Manager or its CLI.</p>
Edit Credentials	<p>Choose the AWS credentials and marketplace subscription to use with this Cloud Volumes ONTAP system.</p> <p>Click <b>Add Subscription</b> to associate the selected credentials with a subscription.</p> <p>To create a pay-as-you-go Cloud Volumes ONTAP system, you need to select AWS credentials that are associated with a subscription to Cloud Volumes ONTAP from the AWS Marketplace. You'll be charged from this subscription for every Cloud Volumes ONTAP 9.6 and later PAYGO system that you create and each add-on feature that you enable.</p> <p><a href="#">Learn how to add additional AWS credentials to Cloud Manager.</a></p>

The following video shows how to associate a pay-as-you-go Marketplace subscription to your AWS credentials:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_aws.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_aws.mp4) (video)



If multiple IAM users work in the same AWS account, then each user needs to subscribe. After the first user subscribes, the AWS Marketplace informs subsequent users that they're already subscribed, as shown in the image below. While a subscription is in place for the *AWS account*, each IAM user needs to associate themselves with that subscription. If you see the message shown below, click the **click here** link to go to Cloud Central and complete the process.

4. **Services:** Keep the services enabled or disable the individual services that you don't want to use with this Cloud Volumes ONTAP system.
  - [Learn more about Cloud Compliance.](#)
  - [Learn more about Backup to Cloud.](#)
  - [Learn more about Monitoring.](#)
5. **HA Deployment Models:** Choose an HA configuration.

For an overview of the deployment models, see [Cloud Volumes ONTAP HA for AWS](#).

6. **Region & VPC:** Enter the network information that you recorded in the AWS worksheet.

The following image shows the page filled out for a multiple AZ configuration:

Region & VPC

AWS Region: US East | N. Virginia

VPC: vpc-a76d91c2 - 172.31.0.0/16

Security group: Use a generated security group

Node 1:  
Availability Zone: us-east-1a  
Subnet: 172.31.8.0/24

Node 2:  
Availability Zone: us-east-1b  
Subnet: 172.31.9.0/24

Mediator:  
Availability Zone: us-east-1c  
Subnet: 172.31.2.0/24

7. **Connectivity and SSH Authentication:** Choose connection methods for the HA pair and the mediator.

8. **Floating IPs:** If you chose multiple AZs, specify the floating IP addresses.

The IP addresses must be outside of the CIDR block for all VPCs in the region. For additional details, see [AWS networking requirements for Cloud Volumes ONTAP HA in multiple AZs](#).

9. **Route Tables:** If you chose multiple AZs, select the route tables that should include routes to the floating IP addresses.

If you have more than one route table, it is very important to select the correct route tables. Otherwise, some clients might not have access to the Cloud Volumes ONTAP HA pair. For more information about route tables, refer to [AWS Documentation: Route Tables](#).

10. **Data Encryption:** Choose no data encryption or AWS-managed encryption.

For AWS-managed encryption, you can choose a different Customer Master Key (CMK) from your account or another AWS account.



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

[Learn how to set up the AWS KMS for Cloud Volumes ONTAP.](#)

[Learn more about supported encryption technologies.](#)

11. **License and Support Site Account:** Specify whether you want to use pay-as-you-go or BYOL, and then specify a NetApp Support Site account.

To understand how licenses work, see [Licensing](#).

A NetApp Support Site Account is optional for pay-as-you-go, but required for BYOL systems. [Learn how to add NetApp Support Site accounts](#).

12. **Preconfigured Packages:** Select one of the packages to quickly launch a Cloud Volumes ONTAP system, or click **Create my own configuration**.

If you choose one of the packages, you only need to specify a volume and then review and approve the configuration.

13. **IAM Role:** You should keep the default option to let Cloud Manager create the roles for you.

If you prefer to use your own policy, it must meet [policy requirements for Cloud Volumes ONTAP nodes and the HA mediator](#).

14. **Licensing:** Change the Cloud Volumes ONTAP version as needed, select a license, an instance type, and the instance tenancy.

If your needs change after you launch the instances, you can modify the license or instance type later.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then Cloud Manager updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.6 RC1 and 9.6 GA is available. The update does not occur from one release to another—for example, from 9.6 to 9.7.

15. **Underlying Storage Resources:** Choose settings for the initial aggregate: a disk type, a size for each disk, and whether data tiering should be enabled.

Note the following:

- The disk type is for the initial volume. You can choose a different disk type for subsequent volumes.
- The disk size is for all disks in the initial aggregate and for any additional aggregates that Cloud Manager creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.

For help choosing a disk type and size, see [Sizing your system in AWS](#).

- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates.

[Learn how data tiering works](#).

16. **WORM:** Activate write once, read many (WORM) storage, if desired.

WORM can't be enabled if data tiering was enabled.

[Learn more about WORM storage](#).

17. **Create Volume:** Enter details for the new volume or click **Skip**.



Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, Cloud Manager enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, Cloud Manager automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

### Volume Details, Protection & Protocol

Details & Protection	Protocol
<p>Volume Name: <input style="width: 200px;" type="text" value="vol"/> Size (GB): <input style="width: 80px;" type="text" value="250"/></p> <p>Snapshot Policy: <input style="width: 300px;" type="text" value="default"/></p> <p><small>Default Policy</small></p>	<p style="text-align: center;"> <span style="margin-right: 20px;">NFS</span> <span style="border-bottom: 2px solid #0070C0; display: inline-block; width: 100px; margin: 0 10px;"></span> <span>iSCSI</span> </p> <p>Share name: <input style="width: 150px;" type="text" value="vol_share"/> Permissions: <input style="width: 150px;" type="text" value="Full Control"/></p> <p>Users / Groups: <input style="width: 300px;" type="text" value="engineering"/></p> <p style="font-size: small; color: #0070C0;">Valid users and groups separated by a semicolon</p>

18. **CIFS Setup:** If you selected the CIFS protocol, set up a CIFS server.

Field	Description
DNS Primary and Secondary IP Address	The IP addresses of the DNS servers that provide name resolution for the CIFS server. The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers. If you configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=Computers,OU=corp</b> in this field.
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. See the <a href="#">Cloud Manager API Developer Guide</a> for details.

19. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and edit the volume tiering policy, if needed.

For more information, see [Understanding volume usage profiles](#) and [Data tiering overview](#).

20. **Review & Approve:** Review and confirm your selections.

- a. Review details about the configuration.

- b. Click **More information** to review details about support and the AWS resources that Cloud Manager will purchase.
- c. Select the **I understand...** check boxes.
- d. Click **Go**.

## Result

Cloud Manager launches the Cloud Volumes ONTAP HA pair. You can track the progress in the timeline.

If you experience any issues launching the HA pair, review the failure message. You can also select the working environment and click Re-create environment.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

## After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.
- If you want to apply quotas to volumes, use System Manager or the CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

# Get started in Azure

## Getting started with Cloud Volumes ONTAP for Azure

Get started with Cloud Volumes ONTAP for Azure in a few steps.



### Create a Connector

If you don't have a [Connector](#) yet, an Account Admin needs to create one. [Learn how to create a Connector in Azure](#).

When you create your first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to deploy a Connector if you don't have one yet.



### Plan your configuration

Cloud Manager offers preconfigured packages that match your workload requirements, or you can create your own configuration. If you choose your own configuration, you should understand the options available to you. [Learn more](#).



### Set up your networking

- a. Ensure that your VNet and subnets will support connectivity between the Connector and Cloud Volumes ONTAP.
- b. Enable outbound internet access from the target VNet so the Connector and Cloud Volumes ONTAP can contact several endpoints.

This step is important because the Connector can't manage Cloud Volumes ONTAP without outbound internet access. If you need to limit outbound connectivity, refer to the list of endpoints for [the Connector and Cloud Volumes ONTAP](#).

[Learn more about networking requirements.](#)



## Launch Cloud Volumes ONTAP using Cloud Manager

Click **Add Working Environment**, select the type of system that you would like to deploy, and complete the steps in the wizard. [Read step-by-step instructions.](#)

### Related links

- [Evaluating](#)
- [Creating a Connector from Cloud Manager](#)
- [Creating a Connector from the Azure Marketplace](#)
- [Installing the Connector software on a Linux host](#)
- [What Cloud Manager does with Azure permissions](#)

## Planning your Cloud Volumes ONTAP configuration in Azure

When you deploy Cloud Volumes ONTAP in Azure, you can choose a preconfigured system that matches your workload requirements, or you can create your own configuration. If you choose your own configuration, you should understand the options available to you.

### Choosing a license type

Cloud Volumes ONTAP is available in two pricing options: pay-as-you-go and Bring Your Own License (BYOL). For pay-as-you-go, you can choose from three licenses: Explore, Standard, or Premium. Each license provides different capacity and compute options.

[Supported configurations for Cloud Volumes ONTAP 9.7 in Azure](#)

### Understanding storage limits

The raw capacity limit for a Cloud Volumes ONTAP system is tied to the license. Additional limits impact the size of aggregates and volumes. You should be aware of these limits as you plan your configuration.

[Storage limits for Cloud Volumes ONTAP 9.7 in Azure](#)

### Sizing your system in Azure

Sizing your Cloud Volumes ONTAP system can help you meet requirements for performance and capacity. You should be aware of a few key points when choosing a VM type, disk type, and disk size:

### Virtual machine type

Look at the supported virtual machine types in the [Cloud Volumes ONTAP Release Notes](#) and then review details about each supported VM type. Be aware that each VM type supports a specific number of data disks.

- [Azure documentation: General purpose virtual machine sizes](#)
- [Azure documentation: Memory optimized virtual machine sizes](#)

## Azure disk type

When you create volumes for Cloud Volumes ONTAP, you need to choose the underlying cloud storage that Cloud Volumes ONTAP uses as a disk.

HA systems use Premium page blobs. Meanwhile, single node systems can use two 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.

For additional details about the use cases for these disks, see [Microsoft Azure Documentation: What disk types are available in Azure?](#)

## Azure disk size

When you launch Cloud Volumes ONTAP instances, you must choose the default disk size for aggregates. Cloud Manager uses this disk size for the initial aggregate, and for any additional aggregates that it creates when you use the simple provisioning option. You can create aggregates that use a disk size different from the default by [using the advanced allocation option](#).



All disks in an aggregate must be the same size.

When choosing a disk size, you should take several factors into consideration. The disk size impacts how much you pay for storage, the size of volumes that you can create in an aggregate, the total capacity available to Cloud Volumes ONTAP, and storage performance.

The performance of Azure Premium Storage is tied to the disk size. Larger disks provide higher IOPS and throughput. For example, choosing 1 TB disks can provide better performance than 500 GB disks, at a higher cost.

There are no performance differences between disk sizes for Standard Storage. You should choose disk size based on the capacity that you need.

Refer to Azure for IOPS and throughput by disk size:

- [Microsoft Azure: Managed Disks pricing](#)
- [Microsoft Azure: Page Blobs pricing](#)

## Choosing a configuration that supports Flash Cache

A Cloud Volumes ONTAP configuration in Azure includes local NVMe storage, which Cloud Volumes ONTAP uses as *Flash Cache* for better performance. [Learn more about Flash Cache](#).

## Azure network information worksheet

When you deploy Cloud Volumes ONTAP in Azure, you need to specify details about your virtual network. You can use a worksheet to collect the information from your administrator.

Azure information	Your value
Region	
Virtual network (VNet)	
Subnet	
Network security group (if using your own)	

## Choosing a write speed

Cloud Manager enables you to choose a write speed setting for single node Cloud Volumes ONTAP systems. Before you choose a write speed, you should understand the differences between the normal and high settings and risks and recommendations when using high write speed.

### Difference between normal write speed and high write speed

When you choose normal write speed, data is written directly to disk, thereby reducing the likelihood of data loss in the event of an unplanned system outage.

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

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

### When to use high write speed

High write speed is a good choice if fast write performance is required for your workload and you can withstand the risk of data loss in the event of an unplanned system outage.

### Recommendations when using high write speed

If you enable high write speed, you should ensure write protection at the application layer.

## Choosing a volume usage profile

ONTAP includes several storage efficiency features that can reduce the total amount of storage that you need. When you create a volume in Cloud Manager, you can choose a profile that enables these features or a profile that disables them. You should learn more about these features to help you decide which profile to use.

NetApp storage efficiency features provide the following benefits:

### Thin provisioning

Presents more logical storage to hosts or users than you actually have in your physical storage pool. Instead of preallocating storage space, storage space is allocated dynamically to each volume as data is written.

## Deduplication

Improves efficiency by locating identical blocks of data and replacing them with references to a single shared block. This technique reduces storage capacity requirements by eliminating redundant blocks of data that reside in the same volume.

## Compression

Reduces the physical capacity required to store data by compressing data within a volume on primary, secondary, and archive storage.

## Networking requirements to deploy and manage Cloud Volumes ONTAP in Azure

Set up your Azure networking so Cloud Volumes ONTAP systems can operate properly. This includes networking for the Connector and Cloud Volumes ONTAP.

### Requirements for Cloud Volumes ONTAP

The following networking requirements must be met in Azure.

#### Outbound internet access for Cloud Volumes ONTAP

Cloud Volumes ONTAP requires outbound internet access to send messages to NetApp AutoSupport, which proactively monitors the health of your storage.

Routing and firewall policies must allow HTTP/HTTPS traffic to the following endpoints so Cloud Volumes ONTAP can send AutoSupport messages:

- <https://support.netapp.com/aods/asupmessage>
- <https://support.netapp.com/asupprod/post/1.0/postAsup>

[Learn how to configure AutoSupport.](#)

## Security groups

You do not need to create security groups because Cloud Manager does that for you. If you need to use your own, refer to the security group rules listed below.

### Number of IP addresses

Cloud Manager allocates the following number of IP addresses to Cloud Volumes ONTAP in Azure:

- Single node: 5 IP addresses
- HA pair: 16 IP addresses

Note that Cloud Manager creates an SVM management LIF on HA pairs, but not on single node systems in Azure.



A LIF is an IP address associated with a physical port. An SVM management LIF is required for management tools like SnapCenter.

### Connection from Cloud Volumes ONTAP to Azure Blob storage for data tiering

If you want to tier cold data to Azure Blob storage, you don't need to set up a connection between the performance tier and the capacity tier as long as Cloud Manager has the required permissions. Cloud Manager enables a VNet service endpoint for you if the Cloud Manager policy has these permissions:

```
"Microsoft.Network/virtualNetworks/subnets/write",  
"Microsoft.Network/routeTables/join/action",
```

These permissions are included in the latest [Cloud Manager policy](#).

For details about setting up data tiering, see [Tiering cold data to low-cost object storage](#).

### Connections to ONTAP systems in other networks

To replicate data between a Cloud Volumes ONTAP system in Azure and ONTAP systems in other networks, you must have a VPN connection between the Azure VNet and the other network—for example, an AWS VPC or your corporate network.

For instructions, refer to [Microsoft Azure Documentation: Create a Site-to-Site connection in the Azure portal](#).

### Requirements for the Connector

Set up your networking so that the Connector can manage resources and processes within your public cloud environment. The most important step is ensuring outbound internet access to various endpoints.



If your network uses a proxy server for all communication to the internet, you can specify the proxy server from the Settings page. Refer to [Configuring the Connector to use a proxy server](#).

### Connections to target networks

A Connector requires a network connection to the VPCs and VNets in which you want to deploy Cloud Volumes ONTAP.

For example, if you install a Connector in your corporate network, then you must set up a VPN connection to the VPC or VNet in which you launch Cloud Volumes ONTAP.

### Outbound internet access

The Connector requires outbound internet access to manage resources and processes within your public cloud environment. A Connector contacts the following endpoints when managing resources in Azure:

Endpoints	Purpose
<a href="https://management.azure.com">https://management.azure.com</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a>	Enables Cloud Manager to deploy and manage Cloud Volumes ONTAP in most Azure regions.
<a href="https://management.microsoftazure.de">https://management.microsoftazure.de</a> <a href="https://login.microsoftonline.de">https://login.microsoftonline.de</a>	Enables Cloud Manager to deploy and manage Cloud Volumes ONTAP in the Azure Germany regions.
<a href="https://management.usgovcloudapi.net">https://management.usgovcloudapi.net</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a>	Enables Cloud Manager to deploy and manage Cloud Volumes ONTAP in the Azure US Gov regions.
<a href="https://api.services.cloud.netapp.com:443">https://api.services.cloud.netapp.com:443</a>	API requests to NetApp Cloud Central.
<a href="https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com">https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com</a>	Provides access to software images, manifests, and templates.
<a href="https://repo.cloud.support.netapp.com">https://repo.cloud.support.netapp.com</a>	Used to download Cloud Manager dependencies.



Endpoints	Purpose
<a href="http://repo.mysql.com/">http://repo.mysql.com/</a>	Used to download MySQL.
<a href="https://cognito-idp.us-east-1.amazonaws.com">https://cognito-idp.us-east-1.amazonaws.com</a> <a href="https://cognito-identity.us-east-1.amazonaws.com">https://cognito-identity.us-east-1.amazonaws.com</a> <a href="https://sts.amazonaws.com">https://sts.amazonaws.com</a> <a href="https://cloud-support-netapp-com-accelerated.s3.amazonaws.com">https://cloud-support-netapp-com-accelerated.s3.amazonaws.com</a>	Enables Cloud Manager to access and download manifests, templates, and Cloud Volumes ONTAP upgrade images.
<a href="https://cloudmanagerinfraproduct.azurecr.io">https://cloudmanagerinfraproduct.azurecr.io</a>	Access to software images of container components for an infrastructure that's running Docker and provides a solution for service integrations with Cloud Manager.
<a href="https://kinesis.us-east-1.amazonaws.com">https://kinesis.us-east-1.amazonaws.com</a>	Enables NetApp to stream data from audit records.
<a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a>	Communication with the Cloud Manager service, which includes Cloud Central accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	Communication with NetApp Cloud Central for centralized user authentication.
<a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	Communication with NetApp AutoSupport.
<a href="https://support.netapp.com/svcgw">https://support.netapp.com/svcgw</a> <a href="https://support.netapp.com/ServiceGW/entitlement">https://support.netapp.com/ServiceGW/entitlement</a> <a href="https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com">https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com</a> <a href="https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com">https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com</a>	Communication with NetApp for system licensing and support registration.
<a href="https://ipa-signer.cloudmanager.netapp.com">https://ipa-signer.cloudmanager.netapp.com</a>	Enables Cloud Manager to generate licenses (for example, a FlexCache license for Cloud Volumes ONTAP)
<a href="https://packages.cloud.google.com/yum">https://packages.cloud.google.com/yum</a> <a href="https://github.com/NetApp/trident/releases/download/">https://github.com/NetApp/trident/releases/download/</a>	Required to connect Cloud Volumes ONTAP systems with a Kubernetes cluster. The endpoints enable installation of NetApp Trident.
*.blob.core.windows.net	Required for HA pairs when using a proxy.
Various third-party locations, for example: <ul style="list-style-type: none"> <li>• <a href="https://repo1.maven.org/maven2">https://repo1.maven.org/maven2</a></li> <li>• <a href="https://oss.sonatype.org/content/repositories">https://oss.sonatype.org/content/repositories</a></li> <li>• <a href="https://repo.typesafe.org">https://repo.typesafe.org</a></li> </ul> Third-party locations are subject to change.	During upgrades, Cloud Manager downloads the latest packages for third-party dependencies.

While you should perform almost all tasks from the SaaS user interface, a local user interface is still available on the Connector. The machine running the web browser must have connections to the following endpoints:

Endpoints	Purpose
The Connector host	<p>You must enter the host's IP address from a web browser to load the Cloud Manager console.</p> <p>Depending on your connectivity to your cloud provider, you can use the private IP or a public IP assigned to the host:</p> <ul style="list-style-type: none"> <li>• A private IP works if you have a VPN and direct connect access to your virtual network</li> <li>• A public IP works in any networking scenario</li> </ul> <p>In any case, you should secure network access by ensuring that security group rules allow access from only authorized IPs or subnets.</p>
<a href="https://auth0.com">https://auth0.com</a> <a href="https://cdn.auth0.com">https://cdn.auth0.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://services.cloud.netapp.com">https://services.cloud.netapp.com</a>	Your web browser connects to these endpoints for centralized user authentication through NetApp Cloud Central.
<a href="https://widget.intercom.io">https://widget.intercom.io</a>	For in-product chat that enables you to talk to NetApp cloud experts.

## Security group rules for Cloud Volumes ONTAP

Cloud Manager creates Azure security groups that include the inbound and outbound rules that Cloud Volumes ONTAP needs to operate successfully. You might want to refer to the ports for testing purposes or if you prefer your to use own security groups.

The security group for Cloud Volumes ONTAP requires both inbound and outbound rules.

### Inbound rules for single node systems

The rules listed below allow traffic, unless the description notes that it blocks specific inbound traffic.

Priority and name	Port and protocol	Source and destination	Description
1000 inbound_ssh	22 TCP	Any to Any	SSH access to the IP address of the cluster management LIF or a node management LIF
1001 inbound_http	80 TCP	Any to Any	HTTP access to the System Manager web console using the IP address of the cluster management LIF
1002 inbound_111_tcp	111 TCP	Any to Any	Remote procedure call for NFS
1003 inbound_111_udp	111 UDP	Any to Any	Remote procedure call for NFS
1004 inbound_139	139 TCP	Any to Any	NetBIOS service session for CIFS

Priority and name	Port and protocol	Source and destination	Description
1005 inbound_161-162_tcp	161-162 TCP	Any to Any	Simple network management protocol
1006 inbound_161-162_udp	161-162 UDP	Any to Any	Simple network management protocol
1007 inbound_443	443 TCP	Any to Any	HTTPS access to the System Manager web console using the IP address of the cluster management LIF
1008 inbound_445	445 TCP	Any to Any	Microsoft SMB/CIFS over TCP with NetBIOS framing
1009 inbound_635_tcp	635 TCP	Any to Any	NFS mount
1010 inbound_635_udp	635 UDP	Any to Any	NFS mount
1011 inbound_749	749 TCP	Any to Any	Kerberos
1012 inbound_2049_tcp	2049 TCP	Any to Any	NFS server daemon
1013 inbound_2049_udp	2049 UDP	Any to Any	NFS server daemon
1014 inbound_3260	3260 TCP	Any to Any	iSCSI access through the iSCSI data LIF
1015 inbound_4045-4046_tcp	4045-4046 TCP	Any to Any	NFS lock daemon and network status monitor
1016 inbound_4045-4046_udp	4045-4046 UDP	Any to Any	NFS lock daemon and network status monitor
1017 inbound_10000	10000 TCP	Any to Any	Backup using NDMP
1018 inbound_11104-11105	11104-11105 TCP	Any to Any	SnapMirror data transfer
3000 inbound_deny_all_tcp	Any port TCP	Any to Any	Block all other TCP inbound traffic
3001 inbound_deny_all_udp	Any port UDP	Any to Any	Block all other UDP inbound traffic

Priority and name	Port and protocol	Source and destination	Description
65000 AllowVnetInBound	Any port Any protocol	VirtualNetwork to VirtualNetwork	Inbound traffic from within the VNet
65001 AllowAzureLoadBalancerInBound	Any port Any protocol	AzureLoadBalancer to Any	Data traffic from the Azure Standard Load Balancer
65500 DenyAllInBound	Any port Any protocol	Any to Any	Block all other inbound traffic

### Inbound rules for HA systems

The rules listed below allow traffic, unless the description notes that it blocks specific inbound traffic.



HA systems have less inbound rules than single node systems because inbound data traffic goes through the Azure Standard Load Balancer. Because of this, traffic from the Load Balancer should be open, as shown in the "AllowAzureLoadBalancerInBound" rule.

Priority and name	Port and protocol	Source and destination	Description
100 inbound_443	443 Any protocol	Any to Any	HTTPS access to the System Manager web console using the IP address of the cluster management LIF
101 inbound_111_tcp	111 Any protocol	Any to Any	Remote procedure call for NFS
102 inbound_2049_tcp	2049 Any protocol	Any to Any	NFS server daemon
111 inbound_ssh	22 Any protocol	Any to Any	SSH access to the IP address of the cluster management LIF or a node management LIF
121 inbound_53	53 Any protocol	Any to Any	DNS and CIFS
65000 AllowVnetInBound	Any port Any protocol	VirtualNetwork to VirtualNetwork	Inbound traffic from within the VNet
65001 AllowAzureLoadBalancerInBound	Any port Any protocol	AzureLoadBalancer to Any	Data traffic from the Azure Standard Load Balancer

Priority and name	Port and protocol	Source and destination	Description
65500 DenyAllInBound	Any port Any protocol	Any to Any	Block all other inbound traffic

### Outbound rules

The predefined security group for Cloud Volumes ONTAP opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

### Basic outbound rules

The predefined security group for Cloud Volumes ONTAP includes the following outbound rules.

Port	Protocol	Purpose
All	All TCP	All outbound traffic
All	All UDP	All outbound traffic

### Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by Cloud Volumes ONTAP.



The source is the interface (IP address) on the Cloud Volumes ONTAP system.

Service	Port	Protocol	Source	Destination	Purpose
Active Directory	88	TCP	Node management LIF	Active Directory forest	Kerberos V authentication
	137	UDP	Node management LIF	Active Directory forest	NetBIOS name service
	138	UDP	Node management LIF	Active Directory forest	NetBIOS datagram service
	139	TCP	Node management LIF	Active Directory forest	NetBIOS service session
	389	TCP & UDP	Node management LIF	Active Directory forest	LDAP
	445	TCP	Node management LIF	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	464	TCP	Node management LIF	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	464	UDP	Node management LIF	Active Directory forest	Kerberos key administration
	749	TCP	Node management LIF	Active Directory forest	Kerberos V change & set Password (RPCSEC_GSS)
	88	TCP	Data LIF (NFS, CIFS, iSCSI)	Active Directory forest	Kerberos V authentication
	137	UDP	Data LIF (NFS, CIFS)	Active Directory forest	NetBIOS name service
	138	UDP	Data LIF (NFS, CIFS)	Active Directory forest	NetBIOS datagram service
	139	TCP	Data LIF (NFS, CIFS)	Active Directory forest	NetBIOS service session
	389	TCP & UDP	Data LIF (NFS, CIFS)	Active Directory forest	LDAP
	445	TCP	Data LIF (NFS, CIFS)	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	464	TCP	Data LIF (NFS, CIFS)	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	464	UDP	Data LIF (NFS, CIFS)	Active Directory forest	Kerberos key administration
	749	TCP	Data LIF (NFS, CIFS)	Active Directory forest	Kerberos V change & set password (RPCSEC_GSS)
	DHCP	68	UDP	Node management LIF	DHCP

Service	Port	Protocol	Source	Destination	Purpose
DHCPS	67	UDP	Node management LIF	DHCP	DHCP server
DNS	53	UDP	Node management LIF and data LIF (NFS, CIFS)	DNS	DNS
NDMP	18600–18699	TCP	Node management LIF	Destination servers	NDMP copy
SMTP	25	TCP	Node management LIF	Mail server	SMTP alerts, can be used for AutoSupport
SNMP	161	TCP	Node management LIF	Monitor server	Monitoring by SNMP traps
	161	UDP	Node management LIF	Monitor server	Monitoring by SNMP traps
	162	TCP	Node management LIF	Monitor server	Monitoring by SNMP traps
	162	UDP	Node management LIF	Monitor server	Monitoring by SNMP traps
SnapMirror	11104	TCP	Intercluster LIF	ONTAP intercluster LIFs	Management of intercluster communication sessions for SnapMirror
	11105	TCP	Intercluster LIF	ONTAP intercluster LIFs	SnapMirror data transfer
Syslog	514	UDP	Node management LIF	Syslog server	Syslog forward messages

### Security group rules for the Connector

The security group for the Connector requires both inbound and outbound rules.

#### Inbound rules

The source for inbound rules in the predefined security group is 0.0.0.0/0.

Port	Protocol	Purpose
22	SSH	Provides SSH access to the Connector host
80	HTTP	Provides HTTP access from client web browsers to the local user interface
443	HTTPS	Provides HTTPS access from client web browsers to the local user interface

#### Outbound rules

The predefined security group for the Connector opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

## Basic outbound rules

The predefined security group for the Connector includes the following outbound rules.

Port	Protocol	Purpose
All	All TCP	All outbound traffic
All	All UDP	All outbound traffic

## Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by the Connector.



The source IP address is the Connector host.

Service	Port	Protocol	Destination	Purpose
Active Directory	88	TCP	Active Directory forest	Kerberos V authentication
	139	TCP	Active Directory forest	NetBIOS service session
	389	TCP	Active Directory forest	LDAP
	445	TCP	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	464	TCP	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	749	TCP	Active Directory forest	Active Directory Kerberos V change & set password (RPCSEC_GSS)
	137	UDP	Active Directory forest	NetBIOS name service
	138	UDP	Active Directory forest	NetBIOS datagram service
	464	UDP	Active Directory forest	Kerberos key administration
API calls and AutoSupport	443	HTTPS	Outbound internet and ONTAP cluster management LIF	API calls to AWS and ONTAP, and sending AutoSupport messages to NetApp
API calls	3000	TCP	ONTAP cluster management LIF	API calls to ONTAP
DNS	53	UDP	DNS	Used for DNS resolve by Cloud Manager

## Launching Cloud Volumes ONTAP in Azure

You can launch a single node system or an HA pair in Azure by creating a Cloud Volumes ONTAP working environment in Cloud Manager.

### Before you begin

- You should have a [Connector that is associated with your workspace](#).





You must be an Account Admin to create a Connector. When you create your first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to create a Connector if you don't have one yet.

- You should be prepared to leave the Connector running at all times.
- You should have chose a configuration and obtained Azure networking information from your administrator. For details, see [Planning your Cloud Volumes ONTAP configuration](#).
- To deploy a BYOL system, you need the 20-digit serial number (license key) for each node.

### About this task

When Cloud Manager creates a Cloud Volumes ONTAP system in Azure, it creates several Azure objects, such as a resource group, network interfaces, and storage accounts. You can review a summary of the resources at the end of the wizard.



#### Potential for Data Loss

Deploying Cloud Volumes ONTAP in an existing, shared resource group is not recommended due to the risk of data loss. While rollback is currently disabled by default when using the API to deploy into an existing resource group, deleting Cloud Volumes ONTAP will potentially delete other resources from that shared group.

The best practice is to use a new, dedicated resource group for Cloud Volumes ONTAP. This is the default and only recommended option when deploying Cloud Volumes ONTAP in Azure from Cloud Manager.

### Steps

1. On the Working Environments page, click **Add Working Environment** and follow the prompts.
2. **Choose a Location:** Select **Microsoft Azure** and **Cloud Volumes ONTAP Single Node** or **Cloud Volumes ONTAP High Availability**.
3. **Details and Credentials:** Optionally change the Azure credentials and subscription, specify a cluster name and resource group name, add tags if needed, and then specify credentials.

The following table describes fields for which you might need guidance:

Field	Description
Working Environment Name	Cloud Manager uses the working environment name to name both the Cloud Volumes ONTAP system and the Azure virtual machine. It also uses the name as the prefix for the predefined security group, if you select that option.
Resource Group Name	Keep the default name for the new resource group or uncheck <b>Use Default</b> and enter your own name for the new resource group.  The best practice is to use a new, dedicated resource group for Cloud Volumes ONTAP. While it is possible to deploy Cloud Volumes ONTAP in an existing, shared resource group by using the API, it's not recommended due to the risk of data loss. See the warning above for more details.

Field	Description
Tags	<p>Tags are metadata for your Azure resources. When you enter tags in this field, Cloud Manager adds them to the resource group associated with the Cloud Volumes ONTAP system.</p> <p>You can add up to four tags from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four tags when creating a working environment.</p> <p>For information about tags, refer to <a href="#">Microsoft Azure Documentation: Using tags to organize your Azure resources</a>.</p>
User name and password	These are the credentials for the Cloud Volumes ONTAP cluster admin account. You can use these credentials to connect to Cloud Volumes ONTAP through OnCommand System Manager or its CLI.
Edit Credentials	You can choose different Azure credentials and a different Azure subscription to use with this Cloud Volumes ONTAP system. You need to associate an Azure Marketplace subscription with the selected Azure subscription in order to deploy a pay-as-you-go Cloud Volumes ONTAP system. <a href="#">Learn how to add credentials</a> .

The following video shows how to associate a Marketplace subscription to an Azure subscription:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_azure.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_azure.mp4) (video)

4. **Services:** Keep the services enabled or disable the individual services that you don't want to use with Cloud Volumes ONTAP.
  - [Learn more about Cloud Compliance](#).
  - [Learn more about Backup to Cloud](#).
5. **Location & Connectivity:** Select a location and security group and select the checkbox to confirm network connectivity between Cloud Manager and the target location.
6. **License and Support Site Account:** Specify whether you want to use pay-as-you-go or BYOL, and then specify a NetApp Support Site account.

To understand how licenses work, see [Licensing](#).

A NetApp Support Site Account is optional for pay-as-you-go, but required for BYOL systems. [Learn how to add NetApp Support Site accounts](#).

7. **Preconfigured Packages:** Select one of the packages to quickly deploy a Cloud Volumes ONTAP system, or click **Create my own configuration**.

If you choose one of the packages, you only need to specify a volume and then review and approve the configuration.

8. **Licensing:** Change the Cloud Volumes ONTAP version as needed, select a license, and select a virtual machine type.

If your needs change after you launch the system, you can modify the license or virtual machine type later.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then Cloud Manager updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.6 RC1 and 9.6 GA is available. The update does not occur from one release to another—for example, from 9.6 to 9.7.

9. **Subscribe from the Azure Marketplace:** Follow the steps if Cloud Manager could not enable programmatic deployments of Cloud Volumes ONTAP.
10. **Underlying Storage Resources:** Choose settings for the initial aggregate: a disk type, a size for each disk, and whether data tiering to Blob storage should be enabled.

Note the following:

- The disk type is for the initial volume. You can choose a different disk type for subsequent volumes.
- The disk size is for all disks in the initial aggregate and for any additional aggregates that Cloud Manager creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.

For help choosing a disk type and size, see [Sizing your system in Azure](#).

- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates.

[Learn more about data tiering](#).

11. **Write Speed & WORM** (single node systems only): Choose **Normal** or **High** write speed, and activate write once, read many (WORM) storage, if desired.

Choosing a write speed is supported with single node systems only.

[Learn more about write speed](#).

WORM can't be enabled if data tiering was enabled.

[Learn more about WORM storage](#).

12. **Secure Communication to Storage & WORM** (HA only): Choose whether to enable an HTTPS connection to Azure storage accounts, and activate write once, read many (WORM) storage, if desired.

The HTTPS connection is from a Cloud Volumes ONTAP 9.7 HA pair to Azure storage accounts. Note that enabling this option can impact write performance. You can't change the setting after you create the working environment.

[Learn more about WORM storage](#).

13. **Create Volume:** Enter details for the new volume or click **Skip**.

Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, Cloud Manager enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, Cloud Manager automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

**Volume Details, Protection & Protocol**

Details & Protection	Protocol
<p>Volume Name: <input style="width: 200px;" type="text" value="vol"/> Size (GB): <input style="width: 80px;" type="text" value="250"/></p> <p>Snapshot Policy: <input style="width: 300px;" type="text" value="default"/></p> <p><small>Default Policy</small></p>	<p style="text-align: center;"> <input type="radio"/> NFS    <input checked="" type="radio"/> CIFS    <input type="radio"/> iSCSI         </p> <hr style="border: 0; border-top: 1px solid #ccc; margin: 5px 0;"/> <p>Share name: <input style="width: 150px;" type="text" value="vol_share"/> Permissions: <input style="width: 150px;" type="text" value="Full Control"/></p> <p>Users / Groups: <input style="width: 300px;" type="text" value="engineering"/></p> <p><small>Valid users and groups separated by a semicolon</small></p>

14. **CIFS Setup:** If you chose the CIFS protocol, set up a CIFS server.

Field	Description
DNS Primary and Secondary IP Address	The IP addresses of the DNS servers that provide name resolution for the CIFS server. The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.  To configure Azure AD Domain Services as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=AADDCC Computers</b> or <b>OU=AADDCC Users</b> in this field. <a href="#">Azure Documentation: Create an Organizational Unit (OU) in an Azure AD Domain Services managed domain</a>
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. See the <a href="#">Cloud Manager API Developer Guide</a> for details.

15. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and change the volume tiering policy, if needed.

For more information, see [Understanding volume usage profiles](#) and [Data tiering overview](#).

16. **Review & Approve:** Review and confirm your selections.

- a. Review details about the configuration.
- b. Click **More information** to review details about support and the Azure resources that Cloud Manager will purchase.
- c. Select the **I understand...** check boxes.
- d. Click **Go**.

### Result

Cloud Manager deploys the Cloud Volumes ONTAP system. You can track the progress in the timeline.

If you experience any issues deploying the Cloud Volumes ONTAP system, review the failure message. You can also select the working environment and click **Re-create environment**.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

### After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.
- If you want to apply quotas to volumes, use System Manager or the CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

## Get started in GCP

### Getting started with Cloud Volumes ONTAP for Google Cloud

Get started with Cloud Volumes ONTAP for GCP in a few steps.



#### Create a Connector

If you don't have a [Connector](#) yet, an Account Admin needs to create one. [Learn how to create a Connector in GCP](#).

When you create your first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to deploy a Connector if you don't have one yet.



#### Plan your configuration

Cloud Manager offers preconfigured packages that match your workload requirements, or you can create your own configuration. If you choose your own configuration, you should understand the options available to you. [Learn more](#).



#### Set up your networking

- a. Ensure that your VPC and subnets will support connectivity between the Connector and Cloud Volumes

ONTAP.

- b. Enable outbound internet access from the target VPC so the Connector and Cloud Volumes ONTAP can contact several endpoints.

This step is important because the Connector can't manage Cloud Volumes ONTAP without outbound internet access. If you need to limit outbound connectivity, refer to the list of endpoints for [the Connector and Cloud Volumes ONTAP](#).

[Learn more about networking requirements.](#)



## Set up GCP for data tiering

Two requirements must be met to tier cold data from Cloud Volumes ONTAP to low-cost object storage (a Google Cloud Storage bucket):

- a. [Configure the Cloud Volumes ONTAP subnet for Private Google Access.](#)
- b. [Set up a service account for data tiering:](#)
  - Assign the predefined *Storage Admin* role to the tiering service account.
  - Add the Connector service account as a *Service Account User* to the tiering service account.

You can provide the user role [in step 3 of the wizard when you create the tiering service account](#), or [grant the role after the service account was created](#).

You'll need to select the tiering service account later when you create a Cloud Volumes ONTAP working environment.

If you don't enable data tiering and select a service account when you create the Cloud Volumes ONTAP system, then you'll need to turn off the system and add the service account to Cloud Volumes ONTAP from the GCP console.



## Enable Google Cloud APIs

[Enable the following Google Cloud APIs in your project.](#) These APIs are required to deploy the Connector and Cloud Volumes ONTAP.

- Cloud Deployment Manager V2 API
- Cloud Logging API
- Cloud Resource Manager API
- Compute Engine API
- Identity and Access Management (IAM) API



## Launch Cloud Volumes ONTAP using Cloud Manager

Click **Add Working Environment**, select the type of system that you would like to deploy, and complete the steps in the wizard. [Read step-by-step instructions.](#)

## Related links

- [Evaluating](#)
- [Creating a Connector from Cloud Manager](#)
- [Installing the Connector software on a Linux host](#)
- [What Cloud Manager does with GCP permissions](#)

## Planning your Cloud Volumes ONTAP configuration in Google Cloud

When you deploy Cloud Volumes ONTAP in Google Cloud, you can choose a preconfigured system that matches your workload requirements, or you can create your own configuration. If you choose your own configuration, you should understand the options available to you.

### Choosing a license type

Cloud Volumes ONTAP is available in two pricing options: pay-as-you-go and Bring Your Own License (BYOL). For pay-as-you-go, you can choose from three licenses: Explore, Standard, or Premium. Each license provides different capacity and compute options.

[Supported configurations for Cloud Volumes ONTAP 9.7 in GCP](#)

### Understanding storage limits

The raw capacity limit for a Cloud Volumes ONTAP system is tied to the license. Additional limits impact the size of aggregates and volumes. You should be aware of these limits as you plan your configuration.

[Storage limits for Cloud Volumes ONTAP 9.7 in GCP](#)

### Sizing your system in GCP

Sizing your Cloud Volumes ONTAP system can help you meet requirements for performance and capacity. You should be aware of a few key points when choosing a machine type, disk type, and disk size:

#### Machine type

Look at the supported machine types in the [Cloud Volumes ONTAP Release Notes](#) and then review details from Google about each supported machine type. Match your workload requirements to the number of vCPUs and memory for the machine type. Note that each CPU core increases networking performance.

Refer to the following for more details:

- [Google Cloud documentation: N1 standard machine types](#)
- [Google Cloud documentation: Performance](#)

#### GCP disk type

When you create volumes for Cloud Volumes ONTAP, you need to choose the underlying cloud storage that Cloud Volumes ONTAP uses for a disk. The disk type can be either *Zonal SSD persistent disks* or *Zonal standard persistent disks*.

SSD persistent disks are best for workloads that require high rates of random IOPS, while Standard persistent disks are economical and can handle sequential read/write operations. For more details, see [Google Cloud documentation: Zonal Persistent disks \(Standard and SSD\)](#).



## GCP disk size

You need to choose an initial disk size when you deploy a Cloud Volumes ONTAP system. After that you can let Cloud Manager manage a system's capacity for you, but if you want to build aggregates yourself, be aware of the following:

- All disks in an aggregate must be the same size.
- Determine the space that you need, while taking performance into consideration.
- The performance of persistent disks scales automatically with disk size and the number of vCPUs available to the system.

Refer to the following for more details:

- [Google Cloud documentation: Zonal Persistent disks \(Standard and SSD\)](#)
- [Google Cloud documentation: Optimizing Persistent Disk and Local SSD Performance](#)

## GCP network information worksheet

When you deploy Cloud Volumes ONTAP in GCP, you need to specify details about your virtual network. You can use a worksheet to collect the information from your administrator.

GCP information	Your value
Region	
Zone	
VPC network	
Subnet	
Firewall policy (if using your own)	

## Choosing a write speed

Cloud Manager enables you to choose a write speed setting for single node Cloud Volumes ONTAP systems. Before you choose a write speed, you should understand the differences between the normal and high settings and risks and recommendations when using high write speed.

### Difference between normal write speed and high write speed

When you choose normal write speed, data is written directly to disk, thereby reducing the likelihood of data loss in the event of an unplanned system outage.

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

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

## When to use high write speed

High write speed is a good choice if fast write performance is required for your workload and you can withstand the risk of data loss in the event of an unplanned system outage.

## Recommendations when using high write speed

If you enable high write speed, you should ensure write protection at the application layer.

## Choosing a volume usage profile

ONTAP includes several storage efficiency features that can reduce the total amount of storage that you need. When you create a volume in Cloud Manager, you can choose a profile that enables these features or a profile that disables them. You should learn more about these features to help you decide which profile to use.

NetApp storage efficiency features provide the following benefits:

### Thin provisioning

Presents more logical storage to hosts or users than you actually have in your physical storage pool. Instead of preallocating storage space, storage space is allocated dynamically to each volume as data is written.

### Deduplication

Improves efficiency by locating identical blocks of data and replacing them with references to a single shared block. This technique reduces storage capacity requirements by eliminating redundant blocks of data that reside in the same volume.

### Compression

Reduces the physical capacity required to store data by compressing data within a volume on primary, secondary, and archive storage.

## Networking requirements to deploy and manage Cloud Volumes ONTAP in GCP

Set up your Google Cloud Platform networking so Cloud Volumes ONTAP systems can operate properly. This includes networking for the Connector and Cloud Volumes ONTAP.

### Requirements for Cloud Volumes ONTAP

The following requirements must be met in GCP.

#### Virtual Private Cloud

Cloud Volumes ONTAP and the Connector are supported in a Google Cloud shared VPC and also in non-shared VPCs.

A shared VPC enables you to configure and centrally manage virtual networks across multiple projects. You can set up shared VPC networks in the *host project* and deploy the Connector and Cloud Volumes ONTAP virtual machine instances in a *service project*. [Google Cloud documentation: Shared VPC overview](#).

The only requirement when using a shared VPC is to provide the [Compute Network User role](#) to the Connector service account. Cloud Manager needs these permissions to query the firewalls, VPC, and subnets in the host project.

## Outbound internet access for Cloud Volumes ONTAP

Cloud Volumes ONTAP requires outbound internet access to send messages to NetApp AutoSupport, which proactively monitors the health of your storage.

Routing and firewall policies must allow HTTP/HTTPS traffic to the following endpoints so Cloud Volumes ONTAP can send AutoSupport messages:

- <https://support.netapp.com/aods/asupmessage>
- <https://support.netapp.com/asupprod/post/1.0/postAsup>

[Learn how to configure AutoSupport.](#)

## Number of IP addresses

Cloud Manager allocates 5 IP addresses to Cloud Volumes ONTAP in GCP.

Note that Cloud Manager doesn't create an SVM management LIF for Cloud Volumes ONTAP in GCP.



A LIF is an IP address associated with a physical port. An SVM management LIF is required for management tools like SnapCenter.

## Firewall rules

You don't need to create firewall rules because Cloud Manager does that for you. If you need to use your own, refer to the firewall rules listed below.

## Connection from Cloud Volumes ONTAP to Google Cloud Storage for data tiering

If you want to tier cold data to a Google Cloud Storage bucket, the subnet in which Cloud Volumes ONTAP resides must be configured for Private Google Access. For instructions, refer to [Google Cloud documentation: Configuring Private Google Access](#).

For additional steps required to set up data tiering in Cloud Manager, see [Tiering cold data to low-cost object storage](#).

## Connections to ONTAP systems in other networks

To replicate data between a Cloud Volumes ONTAP system in GCP and ONTAP systems in other networks, you must have a VPN connection between the VPC and the other network—for example, your corporate network.

For instructions, refer to [Google Cloud documentation: Cloud VPN overview](#).

## Requirements for the Connector

Set up your networking so that the Connector can manage resources and processes within your public cloud environment. The most important step is ensuring outbound internet access to various endpoints.



If your network uses a proxy server for all communication to the internet, you can specify the proxy server from the Settings page. Refer to [Configuring the Connector to use a proxy server](#).

## Connection to target networks

A Connector requires a network connection to the VPCs and VNets in which you want to deploy Cloud Volumes ONTAP.

For example, if you install a Connector in your corporate network, then you must set up a VPN connection to the VPC or VNet in which you launch Cloud Volumes ONTAP.

### Outbound internet access

The Connector requires outbound internet access to manage resources and processes within your public cloud environment. A Connector contacts the following endpoints when managing resources in GCP:

Endpoints	Purpose
<a href="https://www.googleapis.com">https://www.googleapis.com</a>	Enables the Connector to contact Google APIs for deploying and managing Cloud Volumes ONTAP in GCP.
<a href="https://api.services.cloud.netapp.com:443">https://api.services.cloud.netapp.com:443</a>	API requests to NetApp Cloud Central.
<a href="https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com">https://cloud.support.netapp.com.s3.us-west-1.amazonaws.com</a>	Provides access to software images, manifests, and templates.
<a href="https://repo.cloud.support.netapp.com">https://repo.cloud.support.netapp.com</a>	Used to download Cloud Manager dependencies.
<a href="http://repo.mysql.com/">http://repo.mysql.com/</a>	Used to download MySQL.
<a href="https://cognito-idp.us-east-1.amazonaws.com">https://cognito-idp.us-east-1.amazonaws.com</a> <a href="https://cognito-identity.us-east-1.amazonaws.com">https://cognito-identity.us-east-1.amazonaws.com</a> <a href="https://sts.amazonaws.com">https://sts.amazonaws.com</a> <a href="https://cloud-support-netapp-com-accelerated.s3.amazonaws.com">https://cloud-support-netapp-com-accelerated.s3.amazonaws.com</a>	Enables the Connector to access and download manifests, templates, and Cloud Volumes ONTAP upgrade images.
<a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a>	Access to software images of container components for an infrastructure that's running Docker and provides a solution for service integrations with Cloud Manager.
<a href="https://kinesis.us-east-1.amazonaws.com">https://kinesis.us-east-1.amazonaws.com</a>	Enables NetApp to stream data from audit records.
<a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a>	Communication with the Cloud Manager service, which includes Cloud Central accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	Communication with NetApp Cloud Central for centralized user authentication.
<a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	Communication with NetApp AutoSupport.
<a href="https://support.netapp.com/svcgw">https://support.netapp.com/svcgw</a> <a href="https://support.netapp.com/ServiceGW/entitlement">https://support.netapp.com/ServiceGW/entitlement</a> <a href="https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com">https://eval.lic.netapp.com.s3.us-west-1.amazonaws.com</a> <a href="https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com">https://cloud-support-netapp-com.s3.us-west-1.amazonaws.com</a>	Communication with NetApp for system licensing and support registration.
<a href="https://ipa-signer.cloudmanager.netapp.com">https://ipa-signer.cloudmanager.netapp.com</a>	Enables Cloud Manager to generate licenses (for example, a FlexCache license for Cloud Volumes ONTAP)
<a href="https://packages.cloud.google.com/yum">https://packages.cloud.google.com/yum</a> <a href="https://github.com/NetApp/trident/releases/download/">https://github.com/NetApp/trident/releases/download/</a>	Required to connect Cloud Volumes ONTAP systems with a Kubernetes cluster. The endpoints enable installation of NetApp Trident.

Endpoints	Purpose
<p>Various third-party locations, for example:</p> <ul style="list-style-type: none"> <li>• <a href="https://repo1.maven.org/maven2">https://repo1.maven.org/maven2</a></li> <li>• <a href="https://oss.sonatype.org/content/repositories">https://oss.sonatype.org/content/repositories</a></li> <li>• <a href="https://repo.typesafe.org">https://repo.typesafe.org</a></li> </ul> <p>Third-party locations are subject to change.</p>	<p>During upgrades, Cloud Manager downloads the latest packages for third-party dependencies.</p>

While you should perform almost all tasks from the SaaS user interface, a local user interface is still available on the Connector. The machine running the web browser must have connections to the following endpoints:

Endpoints	Purpose
The Connector host	<p>You must enter the host's IP address from a web browser to load the Cloud Manager console.</p> <p>Depending on your connectivity to your cloud provider, you can use the private IP or a public IP assigned to the host:</p> <ul style="list-style-type: none"> <li>• A private IP works if you have a VPN and direct connect access to your virtual network</li> <li>• A public IP works in any networking scenario</li> </ul> <p>In any case, you should secure network access by ensuring that security group rules allow access from only authorized IPs or subnets.</p>
<a href="https://auth0.com">https://auth0.com</a> <a href="https://cdn.auth0.com">https://cdn.auth0.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://services.cloud.netapp.com">https://services.cloud.netapp.com</a>	Your web browser connects to these endpoints for centralized user authentication through NetApp Cloud Central.
<a href="https://widget.intercom.io">https://widget.intercom.io</a>	For in-product chat that enables you to talk to NetApp cloud experts.

## Firewall rules for Cloud Volumes ONTAP

Cloud Manager creates GCP firewall rules that include the inbound and outbound rules that Cloud Manager and Cloud Volumes ONTAP need to operate successfully. You might want to refer to the ports for testing purposes or if you prefer your to use own security groups.

The firewall rules for Cloud Volumes ONTAP requires both inbound and outbound rules.

### Inbound rules

The source for inbound rules in the predefined security group is 0.0.0.0/0.

Protocol	Port	Purpose
All ICMP	All	Pinging the instance
HTTP	80	HTTP access to the System Manager web console using the IP address of the cluster management LIF
HTTPS	443	HTTPS access to the System Manager web console using the IP address of the cluster management LIF
SSH	22	SSH access to the IP address of the cluster management LIF or a node management LIF
TCP	111	Remote procedure call for NFS
TCP	139	NetBIOS service session for CIFS
TCP	161-162	Simple network management protocol
TCP	445	Microsoft SMB/CIFS over TCP with NetBIOS framing
TCP	635	NFS mount
TCP	749	Kerberos
TCP	2049	NFS server daemon
TCP	3260	iSCSI access through the iSCSI data LIF
TCP	4045	NFS lock daemon
TCP	4046	Network status monitor for NFS
TCP	10000	Backup using NDMP
TCP	11104	Management of intercluster communication sessions for SnapMirror
TCP	11105	SnapMirror data transfer using intercluster LIFs
UDP	111	Remote procedure call for NFS
UDP	161-162	Simple network management protocol
UDP	635	NFS mount
UDP	2049	NFS server daemon
UDP	4045	NFS lock daemon
UDP	4046	Network status monitor for NFS
UDP	4049	NFS rquotad protocol

### Outbound rules

The predefined security group for Cloud Volumes ONTAP opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

### Basic outbound rules

The predefined security group for Cloud Volumes ONTAP includes the following outbound rules.

Protocol	Port	Purpose
All ICMP	All	All outbound traffic
All TCP	All	All outbound traffic
All UDP	All	All outbound traffic

### Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by Cloud Volumes ONTAP.



The source is the interface (IP address) on the Cloud Volumes ONTAP system.

Service	Protocol	Port	Source	Destination	Purpose
Active Directory	TCP	88	Node management LIF	Active Directory forest	Kerberos V authentication
	UDP	137	Node management LIF	Active Directory forest	NetBIOS name service
	UDP	138	Node management LIF	Active Directory forest	NetBIOS datagram service
	TCP	139	Node management LIF	Active Directory forest	NetBIOS service session
	TCP & UDP	389	Node management LIF	Active Directory forest	LDAP
	TCP	445	Node management LIF	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	TCP	464	Node management LIF	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	UDP	464	Node management LIF	Active Directory forest	Kerberos key administration
	TCP	749	Node management LIF	Active Directory forest	Kerberos V change & set Password (RPCSEC_GSS)
	TCP	88	Data LIF (NFS, CIFS, iSCSI)	Active Directory forest	Kerberos V authentication
	UDP	137	Data LIF (NFS, CIFS)	Active Directory forest	NetBIOS name service
	UDP	138	Data LIF (NFS, CIFS)	Active Directory forest	NetBIOS datagram service
	TCP	139	Data LIF (NFS, CIFS)	Active Directory forest	NetBIOS service session
	TCP & UDP	389	Data LIF (NFS, CIFS)	Active Directory forest	LDAP
	TCP	445	Data LIF (NFS, CIFS)	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	TCP	464	Data LIF (NFS, CIFS)	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	UDP	464	Data LIF (NFS, CIFS)	Active Directory forest	Kerberos key administration
	TCP	749	Data LIF (NFS, CIFS)	Active Directory forest	Kerberos V change & set password (RPCSEC_GSS)



Service	Protocol	Port	Source	Destination	Purpose
Cluster	All traffic	All traffic	All LIFs on one node	All LIFs on the other node	Intercluster communications (Cloud Volumes ONTAP HA only)
	TCP	3000	Node management LIF	HA mediator	ZAPI calls (Cloud Volumes ONTAP HA only)
	ICMP	1	Node management LIF	HA mediator	Keep alive (Cloud Volumes ONTAP HA only)
DHCP	UDP	68	Node management LIF	DHCP	DHCP client for first-time setup
DHCPS	UDP	67	Node management LIF	DHCP	DHCP server
DNS	UDP	53	Node management LIF and data LIF (NFS, CIFS)	DNS	DNS
NDMP	TCP	1860-18699	Node management LIF	Destination servers	NDMP copy
SMTP	TCP	25	Node management LIF	Mail server	SMTP alerts, can be used for AutoSupport
SNMP	TCP	161	Node management LIF	Monitor server	Monitoring by SNMP traps
	UDP	161	Node management LIF	Monitor server	Monitoring by SNMP traps
	TCP	162	Node management LIF	Monitor server	Monitoring by SNMP traps
	UDP	162	Node management LIF	Monitor server	Monitoring by SNMP traps
SnapMirror	TCP	11104	Intercluster LIF	ONTAP intercluster LIFs	Management of intercluster communication sessions for SnapMirror
	TCP	11105	Intercluster LIF	ONTAP intercluster LIFs	SnapMirror data transfer
Syslog	UDP	514	Node management LIF	Syslog server	Syslog forward messages

### Firewall rules for the Connector

The firewall rules for the Connector requires both inbound and outbound rules.

#### Inbound rules

The source for inbound rules in the predefined firewall rules is 0.0.0.0/0.

Protocol	Port	Purpose
SSH	22	Provides SSH access to the Connector host
HTTP	80	Provides HTTP access from client web browsers to the local user interface
HTTPS	443	Provides HTTPS access from client web browsers to the local user interface

### Outbound rules

The predefined firewall rules for the Connector opens all outbound traffic. If that is acceptable, follow the basic outbound rules. If you need more rigid rules, use the advanced outbound rules.

### Basic outbound rules

The predefined firewall rules for the Connector includes the following outbound rules.

Protocol	Port	Purpose
All TCP	All	All outbound traffic
All UDP	All	All outbound traffic

### Advanced outbound rules

If you need rigid rules for outbound traffic, you can use the following information to open only those ports that are required for outbound communication by the Connector.



The source IP address is the Connector host.

Service	Protocol	Port	Destination	Purpose
Active Directory	TCP	88	Active Directory forest	Kerberos V authentication
	TCP	139	Active Directory forest	NetBIOS service session
	TCP	389	Active Directory forest	LDAP
	TCP	445	Active Directory forest	Microsoft SMB/CIFS over TCP with NetBIOS framing
	TCP	464	Active Directory forest	Kerberos V change & set password (SET_CHANGE)
	TCP	749	Active Directory forest	Active Directory Kerberos V change & set password (RPCSEC_GSS)
	UDP	137	Active Directory forest	NetBIOS name service
	UDP	138	Active Directory forest	NetBIOS datagram service
	UDP	464	Active Directory forest	Kerberos key administration
API calls and AutoSupport	HTTPS	443	Outbound internet and ONTAP cluster management LIF	API calls to GCP and ONTAP, and sending AutoSupport messages to NetApp
API calls	TCP	3000	ONTAP cluster management LIF	API calls to ONTAP


Service	Protocol	Port	Destination	Purpose
DNS	UDP	53	DNS	Used for DNS resolve by Cloud Manager

## Using customer-managed encryption keys with Cloud Volumes ONTAP

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.

### Steps

1. Give the Connector service account permission to use the encryption key.



2. Obtain the "id" of the key by invoking the get command for the /gcp/vsa/metadata/gcp-encryption-keys API.
3. Use the "GcpEncryption" parameter with your API request when creating a working environment.

### Example

```
"gcpEncryptionParameters": {
  "key": "projects/tlv-support/locations/us-east4/keyRings/Nikiskeys/cryptoKeys/generatedkey1"
}
```

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

## Launching Cloud Volumes ONTAP in GCP

You can launch a single node Cloud Volumes ONTAP system in GCP by creating a working environment.

### What you'll need

- You should have a [Connector that is associated with your workspace](#).



You must be an Account Admin to create a Connector. When you create your first Cloud Volumes ONTAP working environment, Cloud Manager prompts you to create a Connector if you don't have one yet.


- You should be prepared to leave the Connector running at all times.
- You should have chose a configuration and obtained GCP networking information from your administrator. For details, see [Planning your Cloud Volumes ONTAP configuration](#).
- To deploy a BYOL system, you need the 20-digit serial number (license key) for each node.
- The following Google Cloud APIs should be [enabled in your project](#):
  - Cloud Deployment Manager V2 API
  - Cloud Logging API
  - Cloud Resource Manager API
  - Compute Engine API
  - Identity and Access Management (IAM) API

### Steps

- On the Working Environments page, click **Add Working Environment** and follow the prompts.
- Choose a Location:** Select **Google Cloud** and **Cloud Volumes ONTAP**.
- Details & Credentials:** Select a project, specify a cluster name, optionally add labels, and then specify credentials.

The following table describes fields for which you might need guidance:

Field	Description
Working Environment Name	Cloud Manager uses the working environment name to name both the Cloud Volumes ONTAP system and the GCP VM instance. It also uses the name as the prefix for the predefined security group, if you select that option.

Field	Description
Add Labels	<p>Labels are metadata for your GCP resources. Cloud Manager adds the labels to the Cloud Volumes ONTAP system and GCP resources associated with the system.</p> <p>You can add up to four labels from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four labels when creating a working environment.</p> <p>For information about labels, refer to <a href="#">Google Cloud Documentation: Labeling Resources</a>.</p>
User name and password	<p>These are the credentials for the Cloud Volumes ONTAP cluster admin account. You can use these credentials to connect to Cloud Volumes ONTAP through System Manager or its CLI.</p>
Edit Project	<p>Select the project where you want Cloud Volumes ONTAP to reside. The default project is the project where Cloud Manager resides.</p> <p>If you don't see any additional projects in the drop-down list, then you haven't yet associated the Cloud Manager service account with other projects. Go to the Google Cloud console, open the IAM service, and select the project. Add the service account with the Cloud Manager role to that project. You'll need to repeat this step for each project.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  This is the service account that you set up for Cloud Manager, <a href="#">as described in step 2b on this page</a>. </div> <p>Click <b>Add Subscription</b> to associate the selected credentials with a subscription.</p> <p>To create a pay-as-you-go Cloud Volumes ONTAP system, you need to select a GCP project that's associated with a subscription to Cloud Volumes ONTAP from the GCP Marketplace.</p>

The following video shows how to associate a pay-as-you-go Marketplace subscription to your GCP project:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_gcp.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_gcp.mp4) (video)

- Location & Connectivity:** Select a location, choose a firewall policy, and select the checkbox to confirm network connectivity to Google Cloud storage for data tiering.

If you want to tier cold data to a Google Cloud Storage bucket, the subnet in which Cloud Volumes ONTAP resides must be configured for Private Google Access. For instructions, refer to [Google Cloud Documentation: Configuring Private Google Access](#).

- License & Support Site Account:** Specify whether you want to use pay-as-you-go or BYOL, and then specify a NetApp Support Site account.

To understand how licenses work, see [Licensing](#).

A NetApp Support Site Account is optional for pay-as-you-go, but required for BYOL systems. [Learn how to](#)

[add NetApp Support Site accounts.](#)

6. **Preconfigured Packages:** Select one of the packages to quickly deploy a Cloud Volumes ONTAP system, or click **Create my own configuration**.

If you choose one of the packages, you only need to specify a volume and then review and approve the configuration.

7. **Licensing:** Change the Cloud Volumes ONTAP version as needed, select a license, and select a virtual machine type.

If your needs change after you launch the system, you can modify the license or virtual machine type later.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then Cloud Manager updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.6 RC1 and 9.6 GA is available. The update does not occur from one release to another—for example, from 9.6 to 9.7.

8. **Underlying Storage Resources:** Choose settings for the initial aggregate: a disk type and the size for each disk.

The disk type is for the initial volume. You can choose a different disk type for subsequent volumes.

The disk size is for all disks in the initial aggregate and for any additional aggregates that Cloud Manager creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.

For help choosing a disk type and size, see [Sizing your system in GCP](#).

9. **Write Speed & WORM:** Choose **Normal** or **High** write speed, and activate write once, read many (WORM) storage, if desired.

Choosing a write speed is supported with single node systems only.

[Learn more about write speed.](#)

WORM can't be enabled if data tiering was enabled.

[Learn more about WORM storage.](#)

10. **Data Tiering in Google Cloud Platform:** Choose whether to enable data tiering on the initial aggregate, choose a storage class for the tiered data, and then either select a service account that has the predefined Storage Admin role (required for Cloud Volumes ONTAP 9.7), or select a GCP account (required for Cloud Volumes ONTAP 9.6).

Note the following:

- Cloud Manager sets the service account on the Cloud Volumes ONTAP instance. This service account provides permissions for data tiering to a Google Cloud Storage bucket. Be sure to add the Cloud Manager service account as a user of the tiering service account, otherwise, you can't select it from Cloud Manager.

- For help with adding a GCP account, see [Setting up and adding GCP accounts for data tiering with 9.6](#).
- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates, but you'll need to turn off the system and add a service account from the GCP console.

[Learn more about data tiering.](#)

11. **Create Volume:** Enter details for the new volume or click **Skip**.

Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, Cloud Manager enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, Cloud Manager automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

**Volume Details, Protection & Protocol**

Details & Protection	Protocol
<p>Volume Name: <input style="width: 200px;" type="text" value="vol"/> Size (GB): <input style="width: 80px;" type="text" value="250"/></p> <p>Snapshot Policy: <input style="width: 300px;" type="text" value="default"/></p> <p><small>Default Policy</small></p>	<p style="text-align: center;"> <input type="radio"/> NFS    <input checked="" type="radio"/> CIFS    <input type="radio"/> iSCSI         </p> <hr style="border: 0; border-top: 1px solid #ccc; margin: 5px 0;"/> <p>Share name: <input style="width: 150px;" type="text" value="vol_share"/> Permissions: <input style="width: 150px;" type="text" value="Full Control"/></p> <p>Users / Groups: <input style="width: 300px;" type="text" value="engineering"/></p> <p><small>Valid users and groups separated by a semicolon</small></p>

12. **CIFS Setup:** If you chose the CIFS protocol, set up a CIFS server.

Field	Description
DNS Primary and Secondary IP Address	The IP addresses of the DNS servers that provide name resolution for the CIFS server. The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. See the <a href="#">Cloud Manager API Developer Guide</a> for details.

13. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and change the volume tiering policy, if needed.

For more information, see [Understanding volume usage profiles](#) and [Data tiering overview](#).

14. **Review & Approve:** Review and confirm your selections.

- a. Review details about the configuration.
- b. Click **More information** to review details about support and the GCP resources that Cloud Manager will purchase.



- c. Select the **I understand...** check boxes.
- d. Click **Go**.

## Result

Cloud Manager deploys the Cloud Volumes ONTAP system. You can track the progress in the timeline.

If you experience any issues deploying the Cloud Volumes ONTAP system, review the failure message. You can also select the working environment and click **Re-create environment**.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

## After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.
- If you want to apply quotas to volumes, use System Manager or the CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

# Provision and manage storage

## Provisioning storage

You can provision additional storage for your Cloud Volumes ONTAP systems from Cloud Manager by managing volumes and aggregates.



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.

## Creating FlexVol volumes

If you need more storage after you launch a Cloud Volumes ONTAP system, you can create new FlexVol volumes for NFS, CIFS, or iSCSI from Cloud Manager.

### About this task

When you create an iSCSI volume, Cloud Manager automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, [use the IQN to connect to the LUN from your hosts](#).



You can create additional LUNs from System Manager or the CLI.

## Before you begin

If you want to use CIFS in AWS, you must have set up DNS and Active Directory. For details, see [Networking requirements for Cloud Volumes ONTAP for AWS](#).

## Steps

1. On the Working Environments page, double-click the name of the Cloud Volumes ONTAP system on which you want to provision FlexVol volumes.

2. Create a new volume on any aggregate or on a specific aggregate:

Action	Steps
Create a new volume and let Cloud Manager choose the containing aggregate	Click <b>Add New Volume</b> .
Create a new volume on a specific aggregate	<ol style="list-style-type: none"> <li>a. Click the menu icon, and then click <b>Advanced &gt; Advanced allocation</b>.</li> <li>b. Click the menu for an aggregate.</li> <li>c. Click <b>Create volume</b>.</li> </ol>

3. Enter details for the new volume, and then click **Continue**.

Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, Cloud Manager enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.

Field	Description
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, Cloud Manager automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

4. If you chose the CIFS protocol and the CIFS server has not been set up, specify details for the server in the Create a CIFS Server dialog box, and then click **Save and continue**:

Field	Description
DNS Primary and Secondary IP Address	<p>The IP addresses of the DNS servers that provide name resolution for the CIFS server.</p> <p>The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.</p>
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	<p>The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.</p> <ul style="list-style-type: none"> <li>• To configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=Computers,OU=corp</b> in this field.</li> <li>• To configure Azure AD Domain Services as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=AADDCC Computers</b> or <b>OU=AADDCC Users</b> in this field. <a href="#">Azure Documentation: Create an Organizational Unit (OU) in an Azure AD Domain Services managed domain</a></li> </ul>
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.

Field	Description
NTP Server	Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. See the <a href="#">Cloud Manager API Developer Guide</a> for details.

- On the Usage Profile, Disk Type, and Tiering Policy page, choose whether you want to enable storage efficiency features, choose a disk type, and edit the tiering policy, if needed.

For help, refer to the following:

- [Understanding volume usage profiles](#)
- [Sizing your system in AWS](#)
- [Sizing your system in Azure](#)
- [Data tiering overview](#)

- Click **Go**.

### Result

Cloud Volumes ONTAP provisions the volume.

### After you finish

If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.

If you want to apply quotas to volumes, you must use System Manager or the CLI. Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

### Creating FlexVol volumes on the second node in an HA configuration

By default, Cloud Manager creates volumes on the first node in an HA configuration. If you need an active-active configuration, in which both nodes serve data to clients, you must create aggregates and volumes on the second node.

### Steps

- On the Working Environments page, double-click the name of the Cloud Volumes ONTAP working environment on which you want to manage aggregates.
- Click the menu icon and then click **Advanced > Advanced allocation**.
- Click **Add Aggregate** and then create the aggregate.
- For Home Node, choose the second node in the HA pair.
- After Cloud Manager creates the aggregate, select it and then click **Create volume**.
- Enter details for the new volume, and then click **Create**.

### After you finish

You can create additional volumes on this aggregate if required.



For HA pairs deployed in multiple AWS Availability Zones, you must mount the volume to clients by using the floating IP address of the node on which the volume resides.

## Creating aggregates

You can create aggregates yourself or let Cloud Manager do it for you when it creates volumes. The benefit of creating aggregates yourself is that you can choose the underlying disk size, which enables you to size your aggregate for the capacity or the performance that you need.

### Steps

1. On the Working Environments page, double-click the name of the Cloud Volumes ONTAP instance on which you want to manage aggregates.
2. Click the menu icon, and then click **Advanced > Advanced allocation**.
3. Click **Add Aggregate** and then specify details for the aggregate.

For help with disk type and disk size, see [Planning your configuration](#).

4. Click **Go**, and then click **Approve and Purchase**.

## Connecting a LUN to a host

When you create an iSCSI volume, Cloud Manager automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, use the IQN to connect to the LUN from your hosts.

Note the following:

1. Cloud Manager's automatic capacity management doesn't apply to LUNs. When Cloud Manager creates a LUN, it disables the autogrow feature.
2. You can create additional LUNs from System Manager or the CLI.

### Steps

1. On the Working Environments page, double-click the Cloud Volumes ONTAP working environment on which you want to manage volumes.
2. Select a volume, and then click **Target iQN**.
3. Click **Copy** to copy the iQN name.
4. Set up an iSCSI connection from the host to the LUN.
  - [ONTAP 9 iSCSI express configuration for Red Hat Enterprise Linux: Starting the iSCSI sessions with the target](#)
  - [ONTAP 9 iSCSI express configuration for Windows: Starting iSCSI sessions with the target](#)

## Using FlexCache volumes to accelerate data access

A FlexCache volume is a storage volume that caches NFS read data from an origin (or source) volume. Subsequent reads to the cached data result in faster access to that data.

You can use FlexCache volumes to speed up access to data or to offload traffic from heavily accessed volumes. FlexCache volumes help improve performance, especially when clients need to access the same data repeatedly, because the data can be served directly without having to access the origin volume. FlexCache volumes work well for system workloads that are read-intensive.

Cloud Manager does not provide management of FlexCache volumes at this time, but you can use the ONTAP CLI or ONTAP System Manager to create and manage FlexCache volumes:

- [FlexCache Volumes for Faster Data Access Power Guide](#)
- [Creating FlexCache volumes in System Manager](#)

Starting with the 3.7.2 release, Cloud Manager generates a FlexCache license for all new Cloud Volumes ONTAP systems. The license includes a 500 GB usage limit.



To generate the license, Cloud Manager needs to access <https://ipa-signer.cloudmanager.netapp.com>. Make sure that this URL is accessible from your firewall.



## Managing existing storage

Cloud Manager enables you to manage volumes, aggregates, and CIFS servers. It also prompts you to move volumes to avoid capacity issues.



### Managing existing volumes


You can manage existing volumes as your storage needs change. You can view, edit, clone, restore, and delete volumes.

#### Steps

1. On the Working Environments page, double-click the Cloud Volumes ONTAP working environment on which you want to manage volumes.
2. Manage your volumes:

Task	Action
View information about a volume	Select a volume, and then click <b>Info</b> .

Task	Action
Edit a volume (read-write volumes only)	<p>a. Select a volume, and then click <b>Edit</b>.</p> <p>b. Modify the volume's Snapshot policy, NFS protocol version, NFS access control list, or share permissions, and then click <b>Update</b>.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  If you need custom Snapshot policies, you can create them by using System Manager. </div>
Clone a volume	<p>a. Select a volume, and then click <b>Clone</b>.</p> <p>b. Modify the clone name as needed, and then click <b>Clone</b>.</p> <p>This process creates a FlexClone volume. A FlexClone volume is a writable, point-in-time copy that is space-efficient because it uses a small amount of space for metadata, and then only consumes additional space as data is changed or added.</p> <p>To learn more about FlexClone volumes, see the <a href="#">ONTAP 9 Logical Storage Management Guide</a>.</p>
Restore data from a Snapshot copy to a new volume	<p>a. Select a volume, and then click <b>Restore from Snapshot copy</b>.</p> <p>b. Select a Snapshot copy, enter a name for the new volume, and then click <b>Restore</b>.</p>
Create a Snapshot copy on demand	<p>a. Select a volume, and then click <b>Create a Snapshot copy</b>.</p> <p>b. Change the name, if needed, and then click <b>Create</b>.</p>
Get the NFS mount command	<p>a. Select a volume, and then click <b>Mount Command</b>.</p> <p>b. Click <b>Copy</b>.</p>
View the target iQN for an iSCSI volume	<p>a. Select a volume, and then click <b>Target iQN</b>.</p> <p>b. Click <b>Copy</b>.</p> <p>c. <a href="#">Use the IQN to connect to the LUN from your hosts</a>.</p>
Change the underlying disk type	<p>a. Select a volume, and then click <b>Change Disk Type &amp; Tiering Policy</b>.</p> <p>b. Select the disk type, and then click <b>Change</b>.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  Cloud Manager moves the volume to an existing aggregate that uses the selected disk type or it creates a new aggregate for the volume. </div>

Task	Action
Change the tiering policy	<ol style="list-style-type: none"> <li>Select a volume, and then click <b>Change Disk Type &amp; Tiering Policy</b>.</li> <li>Click <b>Edit Policy</b>.</li> <li>Select a different policy and click <b>Change</b>.</li> </ol> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>Cloud Manager moves the volume to an existing aggregate that uses the selected disk type with tiering, or it creates a new aggregate for the volume.</p> </div>
Delete a volume	<ol style="list-style-type: none"> <li>Select a volume, and then click <b>Delete</b>.</li> <li>Click <b>Delete</b> again to confirm.</li> </ol>

## Managing existing aggregates

Manage aggregates yourself by adding disks, viewing information about the aggregates, and by deleting them.

### Before you begin


If you want to delete an aggregate, you must have first deleted the volumes in the aggregate.

### About this task

If an aggregate is running out of space, you can move volumes to another aggregate by using OnCommand System Manager.

### Steps

- On the Working Environments page, double-click the Cloud Volumes ONTAP working environment on which you want to manage aggregates.
- Click the menu icon and then click **Advanced > Advanced allocation**.
- Manage your aggregates:

Task	Action
View information about an aggregate	Select an aggregate and click <b>Info</b> .
Create a volume on a specific aggregate	Select an aggregate and click <b>Create volume</b> .
Add disks to an aggregate	<ol style="list-style-type: none"> <li>Select an aggregate and click <b>Add AWS disks</b> or <b>Add Azure disks</b>.</li> <li>Select the number of disks that you want to add and click <b>Add</b>.</li> </ol> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>All disks in an aggregate must be the same size.</p> </div>
Delete an aggregate	<ol style="list-style-type: none"> <li>Select an aggregate that does not contain any volumes and click <b>Delete</b>.</li> <li>Click <b>Delete</b> again to confirm.</li> </ol>



## Modifying the CIFS server

If you change your DNS servers or Active Directory domain, you need to modify the CIFS server in Cloud Volumes ONTAP so that it can continue to serve storage to clients.

### Steps

1. From the working environment, click the menu icon and then click **Advanced > CIFS setup**.
2. Specify settings for the CIFS server:

Task	Action
DNS Primary and Secondary IP Address	The IP addresses of the DNS servers that provide name resolution for the CIFS server.  The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers. If you configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=Computers,OU=corp</b> in this field.
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. See the <a href="#">Cloud Manager API Developer Guide</a> for details.

3. Click **Save**.

### Result

Cloud Volumes ONTAP updates the CIFS server with the changes.

## Moving a volume

Move volumes for capacity utilization, improved performance, and to satisfy service-level agreements.

You can move a volume in System Manager by selecting a volume and the destination aggregate, starting the volume move operation, and optionally monitoring the volume move job. When using System Manager, a volume move operation finishes automatically.

### Steps

1. Use System Manager or the CLI to move the volumes to the aggregate.

In most situations, you can use System Manager to move volumes.

For instructions, see the [ONTAP 9 Volume Move Express Guide](#).

## Moving a volume when Cloud Manager displays an Action Required message

Cloud Manager might display an Action Required message that says moving a volume is necessary to avoid capacity issues, but that it cannot provide recommendations to correct the issue. If this happens, you need to identify how to correct the issue and then move one or more volumes.

### Steps

1. [Identify how to correct the issue](#).
2. Based on your analysis, move volumes to avoid capacity issues:
  - [Move volumes to another system](#).
  - [Move volumes to another aggregate on the same system](#).


### Identifying how to correct capacity issues

If Cloud Manager cannot provide recommendations for moving a volume to avoid capacity issues, you must identify the volumes that you need to move and whether you should move them to another aggregate on the same system or to another system.

### Steps

1. View the advanced information in the Action Required message to identify the aggregate that has reached its capacity limit.

For example, the advanced information should say something similar to the following: Aggregate aggr1 has reached its capacity limit.
2. Identify one or more volumes to move out of the aggregate:
  - a. In the working environment, click the menu icon, and then click **Advanced > Advanced allocation**.
  - b. Select the aggregate, and then click **Info**.
  - c. Expand the list of volumes.



- d. Review the size of each volume and choose one or more volumes to move out of the aggregate.

You should choose volumes that are large enough to free space in the aggregate so that you avoid additional capacity issues in the future.

3. If the system has not reached the disk limit, you should move the volumes to an existing aggregate or a new aggregate on the same system.

For details, see [Moving volumes to another aggregate to avoid capacity issues](#).

4. If the system has reached the disk limit, do any of the following:
  - a. Delete any unused volumes.
  - b. Rearrange volumes to free space on an aggregate.

For details, see [Moving volumes to another aggregate to avoid capacity issues](#).

- c. Move two or more volumes to another system that has space.

For details, see [Moving volumes to another system to avoid capacity issues](#).

### **Moving volumes to another system to avoid capacity issues**

You can move one or more volumes to another Cloud Volumes ONTAP system to avoid capacity issues. You might need to do this if the system reached its disk limit.

#### **About this task**

You can follow the steps in this task to correct the following Action Required message:

```
Moving a volume is necessary to avoid capacity issues; however, Cloud Manager cannot perform this action for you because the system has reached the disk limit.
```

#### **Steps**

1. Identify a Cloud Volumes ONTAP system that has available capacity, or deploy a new system.
2. Drag and drop the source working environment on the target working environment to perform a one-time data replication of the volume.

For details, see [Replicating data between systems](#).

3. Go to the Replication Status page, and then break the SnapMirror relationship to convert the replicated volume from a data protection volume to a read/write volume.

For details, see [Managing data replication schedules and relationships](#).

4. Configure the volume for data access.

For information about configuring a destination volume for data access, see the [ONTAP 9 Volume Disaster Recovery Express Guide](#).

5. Delete the original volume.

For details, see [Managing existing volumes](#).

## Moving volumes to another aggregate to avoid capacity issues

You can move one or more volumes to another aggregate to avoid capacity issues.

### About this task

You can follow the steps in this task to correct the following Action Required message:

```
Moving two or more volumes is necessary to avoid capacity issues; however,
Cloud Manager cannot perform this action for you.
```

### Steps

1. Verify whether an existing aggregate has available capacity for the volumes that you need to move:
  - a. In the working environment, click the menu icon, and then click **Advanced > Advanced allocation**.
  - b. Select each aggregate, click **Info**, and then view the available capacity (aggregate capacity minus used aggregate capacity).

**aggr1**

Aggregate Capacity: 442.94 GB

---

Used Aggregate Capacity: 105.66 GB

---

2. If needed, add disks to an existing aggregate:
  - a. Select the aggregate, and then click **Add disks**.
  - b. Select the number of disks to add, and then click **Add**.
3. If no aggregates have available capacity, create a new aggregate.  
For details, see [Creating aggregates](#).
4. Use System Manager or the CLI to move the volumes to the aggregate.
5. In most situations, you can use System Manager to move volumes.

For instructions, see the [ONTAP 9 Volume Move Express Guide](#).

### Reasons why a volume move might perform slowly

Moving a volume might take longer than you expect if any of the following conditions are true for Cloud Volumes ONTAP:

- The volume is a clone.
- The volume is a parent of a clone.
- The source or destination aggregate has a single Throughput Optimized HDD (st1) disk.
- The Cloud Volumes ONTAP system is in AWS and one aggregate uses an older naming scheme for objects. Both aggregates have to use the same name format.

An older naming scheme is used if data tiering was enabled on an aggregate in the 9.4 release or earlier.

- The encryption settings don't match on the source and destination aggregates, or a rekey is in progress.
- The `-tiering-policy` option was specified on the volume move to change the tiering policy.
- The `-generate-destination-key` option was specified on the volume move.

## Tiering inactive data to low-cost object storage

You can reduce storage costs for Cloud Volumes ONTAP by combining an SSD or HDD performance tier for hot data with an object storage capacity tier for inactive data. For a high-level overview, see [Data tiering overview](#).

To set up data tiering, you simply need to do the following:



### Choose a supported configuration

Most configurations are supported. If you have a Cloud Volumes ONTAP Standard, Premium, or BYOL system running the most recent version, then you should be good to go. [Learn more](#).



### Ensure connectivity between Cloud Volumes ONTAP and object storage

- For AWS, you'll need a VPC Endpoint to S3. [Learn more](#).
- For Azure, you won't need to do anything as long as Cloud Manager has the required permissions. [Learn more](#).
- For GCP, you need to configure the subnet for Private Google Access and set up a service account. [Learn more](#).



### Choose a tiering policy when creating, modifying, or replicating a volume

Cloud Manager prompts you to choose a tiering policy when you create, modify, or replicate a volume.

- [Tiering data on read-write volumes](#)
- [Tiering data on data protection volumes](#)



#### What's not required for data tiering

- You don't need to install a feature license to enable data tiering.
- You don't need to create the capacity tier (an S3 bucket, Azure Blob container, or GCP bucket). Cloud Manager does that for you.

## Configurations that support data tiering

You can enable data tiering when using specific configurations and features:

- Data tiering is supported with Cloud Volumes ONTAP Standard, Premium, and BYOL, starting with the following versions:

- Version 9.2 in AWS
- Version 9.4 in Azure with single node systems
- Version 9.6 in Azure with HA pairs
- Version 9.6 in GCP



Data tiering is not supported in Azure with the DS3\_v2 virtual machine type.

- In AWS, the performance tier can be General Purpose SSDs, Provisioned IOPS SSDs, or Throughput Optimized HDDs.
- In Azure, the performance tier can be Premium SSD managed disks, Standard SSD managed disks, or Standard HDD managed disks.
- In GCP, the performance tier can be either SSDs or HDDs (standard disks).
- Data tiering is supported with encryption technologies.
- Thin provisioning must be enabled on volumes.

### Requirements to tier cold data to AWS S3

Ensure that Cloud Volumes ONTAP has a connection to S3. The best way to provide that connection is by creating a VPC Endpoint to the S3 service. For instructions, see [AWS Documentation: Creating a Gateway Endpoint](#).

When you create the VPC Endpoint, be sure to select the region, VPC, and route table that corresponds to the Cloud Volumes ONTAP instance. You must also modify the security group to add an outbound HTTPS rule that enables traffic to the S3 endpoint. Otherwise, Cloud Volumes ONTAP cannot connect to the S3 service.

If you experience any issues, see [AWS Support Knowledge Center: Why can't I connect to an S3 bucket using a gateway VPC endpoint?](#)

### Requirements to tier cold data to Azure Blob storage

You don't need to set up a connection between the performance tier and the capacity tier as long as Cloud Manager has the required permissions. Cloud Manager enables a VNet service endpoint for you if the Cloud Manager policy has these permissions:

```
"Microsoft.Network/virtualNetworks/subnets/write",
"Microsoft.Network/routeTables/join/action",
```

The permissions are included in the latest [Cloud Manager policy](#).

### Requirements to tier cold data to a Google Cloud Storage bucket

- The subnet in which Cloud Volumes ONTAP resides must be configured for Private Google Access. For instructions, refer to [Google Cloud Documentation: Configuring Private Google Access](#).
- You need a service account that has the predefined Storage Admin role. You'll need to select this service account when you create a Cloud Volumes ONTAP working environment.

[Set up this tiering service account as follows:](#)

- a. Assign the predefined *Storage Admin* role to the tiering service account.
- b. Add the Connector service account as a *Service Account User* to the tiering service account.

You can provide the user role [in step 3 of the wizard when you create the tiering service account](#), or [grant the role after the service account was created](#).

You'll need to select the tiering service account later when you create a Cloud Volumes ONTAP working environment.

If you don't enable data tiering and select a service account when you create the Cloud Volumes ONTAP system, then you'll need to turn off the system and add the service account to Cloud Volumes ONTAP from the GCP console.

## Tiering data from read-write volumes

Cloud Volumes ONTAP can tier inactive data on read-write volumes to cost-effective object storage, freeing up the performance tier for hot data.

### Steps


1. In the working environment, create a new volume or change the tier of an existing volume:

Task	Action
Create a new volume	Click <b>Add New Volume</b> .
Modify an existing volume	Select the volume and click <b>Change Disk Type &amp; Tiering Policy</b> .

2. Select a tiering policy.

For a description of these policies, see [Data tiering overview](#).

### Example



#### Tiering data to object storage

---

**Volume Tiering Policy**

- All** - Immediately tiers all data (not including metadata) to object storage.
- Auto** - Tiers cold Snapshot copies and cold user data from the active file system to object storage.
- Snapshot Only** - Tiers cold Snapshot copies to object storage
- None** - Data tiering is disabled.

---

**Working Environment S3 Storage classes:** Standard

Cloud Manager creates a new aggregate for the volume if a data tiering-enabled aggregate does not already exist.



If you prefer to create aggregates yourself, you can enable data tiering on aggregates when you create them.


## Tiering data from data protection volumes

Cloud Volumes ONTAP can tier data from a data protection volume to a capacity tier. If you activate the destination volume, the data gradually moves to the performance tier as it is read.

### Steps

1. On the Working Environments page, select the working environment that contains the source volume, and then drag it to the working environment to which you want to replicate the volume.
2. Follow the prompts until you reach the tiering page and enable data tiering to object storage.

### Example



For help with replicating data, see [Replicating data to and from the cloud](#).

## Changing the storage class for tiered data

After you deploy Cloud Volumes ONTAP, you can reduce your storage costs by changing the storage class for inactive data that hasn't been accessed for 30 days. The access costs are higher if you do access the data, so you must take that into consideration before you change the storage class.

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

For information about supported storage classes, see [Data tiering overview](#).

### Steps

1. From the working environment, click the menu icon and then click **Storage Classes** or **Blob Storage Tiering**.
2. Choose a storage class and then click **Save**.

## Can I enable data tiering on an existing aggregate?

No, you can't enable data tiering on an existing aggregate. You can only enable data tiering on new aggregates.

You can enable data tiering on a new aggregate either [by creating an aggregate yourself](#) or [by creating a new volume with data tiering enabled](#). Cloud Manager would then create a new aggregate for the volume if a data tiering-enabled aggregate doesn't already exist.



## Managing storage VMs

A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*. Cloud Volumes ONTAP is configured with one storage VM by default, but some configurations support additional storage VMs.

### Supported number of storage VMs

Cloud Volumes ONTAP 9.7 supports multiple storage VMs in AWS with certain configurations and an add-on license. [View the number of supported storage VMs in AWS](#). Contact your account team to obtain an SVM add-on license.

All other Cloud Volumes ONTAP configurations support one data-serving storage VM and one destination storage VM used for disaster recovery. You can activate the destination storage VM for data access if there's an outage on the source storage VM.

A storage VM spans the entire Cloud Volumes ONTAP system (HA pair or single node).

### Creating additional storage VMs

If supported by your configuration, you can create additional storage VMs using [System Manager or the CLI](#).

- [Creating an SVM for SMB access](#)
- [Creating an SVM for NFS access](#)
- [Creating an SVM for iSCSI access](#)
- [Creating a destination SVM for disaster recovery](#)

### Working with multiple storage VMs in Cloud Manager

Cloud Manager supports any additional storage VMs that you create from System Manager or the CLI.

For example, the following image shows how you can choose a storage VM when you create a volume.

### Details & Protection

Storage VM Name ?

svm\_name1 v

Volume Name Size (GiB) ?

Snapshot Policy

default v

? Default Policy

And the following image shows how you can choose a storage VM when replicating a volume to another system.

Destination Volume Name

volume\_copy

Destination Storage VM Name

svm\_name1 v

Destination Aggregate

Automatically select the best aggregate v

### Managing storage VM disaster recovery

Cloud Manager doesn't provide any setup or orchestration support for storage VM disaster recovery. You must use System Manager or the CLI.

- [SVM Disaster Recovery Preparation Express Guide](#)
- [SVM Disaster Recovery Express Guide](#)


## Modifying the storage VM name

Cloud Manager automatically names the single storage VM that it creates for Cloud Volumes ONTAP. You can modify the name of the storage VM if you have strict naming standards. For example, you might want the name to match how you name the storage VMs for your ONTAP clusters.

If you created any additional storage VMs for Cloud Volumes ONTAP, then you can't rename the storage VMs from Cloud Manager. You'll need to do so directly from Cloud Volumes ONTAP by using System Manager or the CLI.

### Steps

1. From the working environment, click the menu icon, and then click **Information**.
2. Click the edit icon to the right of the storage VM name.



3. In the Modify SVM Name dialog box, change the name, and then click **Save**.

## Using Cloud Volumes ONTAP as persistent storage for Kubernetes

Cloud Manager can automate the deployment of NetApp Trident on Kubernetes clusters so you can use Cloud Volumes ONTAP as persistent storage for containers.

Trident is a fully-supported open source project maintained by NetApp. Trident integrates natively with Kubernetes and its Persistent Volume framework to seamlessly provision and manage volumes from systems running any combination of NetApp's storage platforms. [Learn more about Trident.](#)



The Kubernetes feature isn't supported with on-prem ONTAP clusters. It's supported with Cloud Volumes ONTAP only.

## Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



### Review prerequisites

Ensure that your environment can meet the prerequisites, which includes connectivity between Kubernetes clusters and Cloud Volumes ONTAP, connectivity between Kubernetes clusters and a Connector, a minimum Kubernetes version of 1.14, at least one worker node in a cluster, and more. [See the complete list.](#)



### Add your Kubernetes clusters to Cloud Manager

In Cloud Manager, click **Kubernetes** and discover clusters directly from your cloud provider's managed service or import a cluster by providing a kubeconfig file.



### Connect your clusters to Cloud Volumes ONTAP

After you add a Kubernetes cluster, click **Connect to Working Environment** to connect the cluster to one or more Cloud Volumes ONTAP systems.



### Start provisioning Persistent Volumes

Request and manage Persistent Volumes using native Kubernetes interfaces and constructs. Cloud Manager creates NFS and iSCSI storage classes that you can use when provisioning Persistent Volumes.

[Learn more about provisioning your first volume with Trident for Kubernetes.](#)

## Reviewing prerequisites

Before you get started, ensure that your Kubernetes clusters and Connector meet specific requirements.

### Kubernetes cluster requirements

- Network connectivity is required between a Kubernetes cluster and the Connector and between a Kubernetes cluster and Cloud Volumes ONTAP.

Both the Connector and Cloud Volumes ONTAP need a connection to the Kubernetes API endpoint:

- For managed clusters, set up a route between a cluster's VPC and the VPC where the Connector and Cloud Volumes ONTAP reside.
- For other clusters, the IP address of the master node or load balancer (as listed in the kubeconfig file) must be reachable by the Connector and Cloud Volumes ONTAP, and it must present a valid TLS certificate.
- A Kubernetes cluster can be in any location that has the network connectivity listed above.
- A Kubernetes cluster must be running version 1.14 at a minimum.

The maximum supported version is defined by Trident. [Click here to see the maximum supported](#)

## Kubernetes version.

- A Kubernetes cluster must have at least one worker node.
- For clusters running in Amazon Elastic Kubernetes Service (Amazon EKS), each cluster needs an IAM role added in order to resolve a permissions error. After you add the cluster, Cloud Manager will prompt you with the exact `eksctl` command that resolves the error.

## Learn about IAM permissions boundaries.

- For clusters running in Azure Kubernetes Service (AKS), those clusters must be assigned the *Azure Kubernetes Service RBAC Cluster Admin* role. This is required so Cloud Manager can install Trident and configure storage classes on the cluster.
- For clusters running in Google Kubernetes Engine (GKE), those clusters must not use the default Container Optimized OS. You should switch them to use Ubuntu.

GKE defaults to using the Google [container-optimized image](#), which doesn't have the utilities that Trident needs to mount volumes.

## Connector requirements

Ensure that the following networking and permissions are in place for the Connector.

## Networking

- The Connector needs an outbound internet connection to access the following endpoints when installing Trident:

<https://packages.cloud.google.com/yum>  
<https://github.com/NetApp/trident/releases/download/>

Cloud Manager installs Trident on a Kubernetes cluster when you connect a working environment to the cluster.

## Required permissions to discover and manage EKS clusters

The Connector needs Admin permissions to discover and manage Kubernetes clusters running in Amazon Elastic Kubernetes Service (EKS):

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": "eks:*",
      "Resource": "*"
    }
  ]
}
```

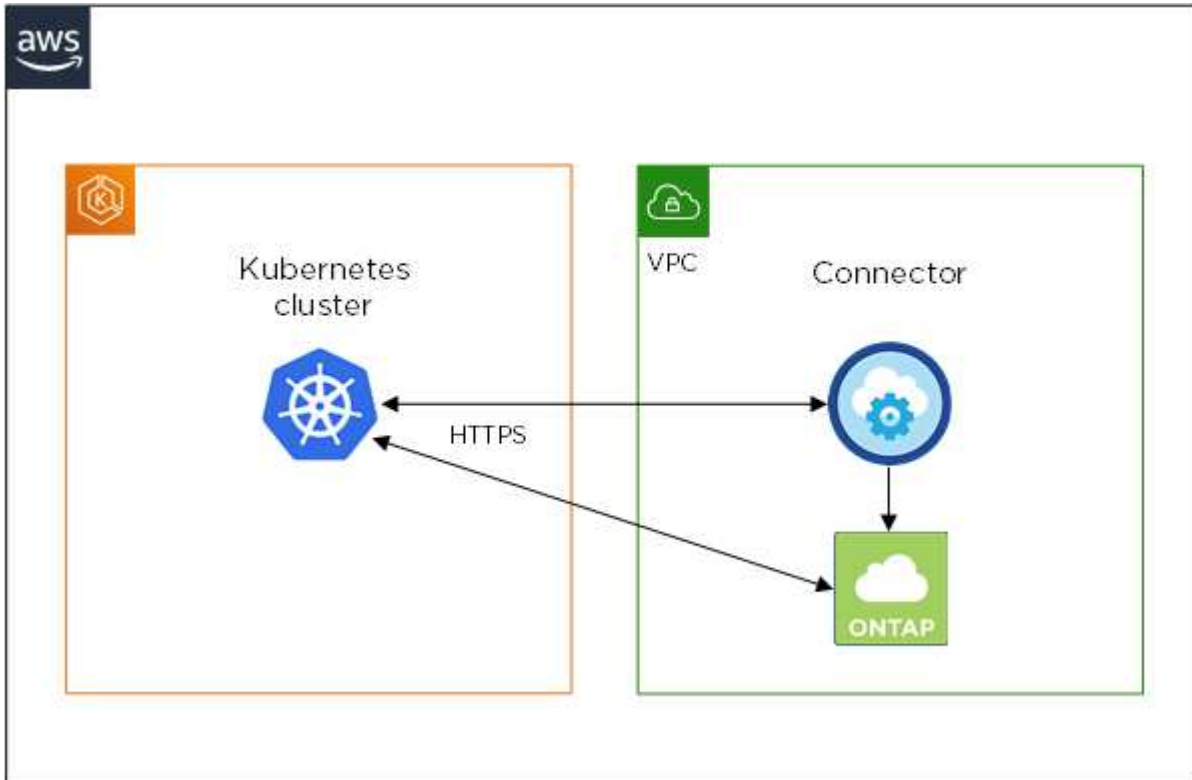
## Required permissions to discover and manage GKE clusters

The Connector needs the following permissions to discover and manage Kubernetes clusters running in Google Kubernetes Engine (GKE):

```
container.*
```

### Example setup

The following image shows an example of a Kubernetes cluster running in Amazon Elastic Kubernetes Service (Amazon EKS) and its connections to the Connector and Cloud Volumes ONTAP.



### Adding Kubernetes clusters

Add Kubernetes clusters to Cloud Manager by discovering the clusters running in your cloud provider's managed Kubernetes service or by importing a cluster's kubeconfig file.

#### Steps

1. At the top of Cloud Manager, click **Kubernetes**.
2. Click **Add Cluster**.
3. Choose one of the available options:
  - Click **Discover Clusters** to discover the managed clusters that Cloud Manager has access to based on permissions that you provided to the Connector.

For example, if your Connector is running in Google Cloud, Cloud Manager uses the permissions from the Connector's service account to discover clusters running in Google Kubernetes Engine (GKE).

- Click **Import Cluster** to import a cluster using a kubeconfig file.

After you upload the file, Cloud Manager verifies connectivity to the cluster and saves an encrypted copy of the kubeconfig file.

## Result

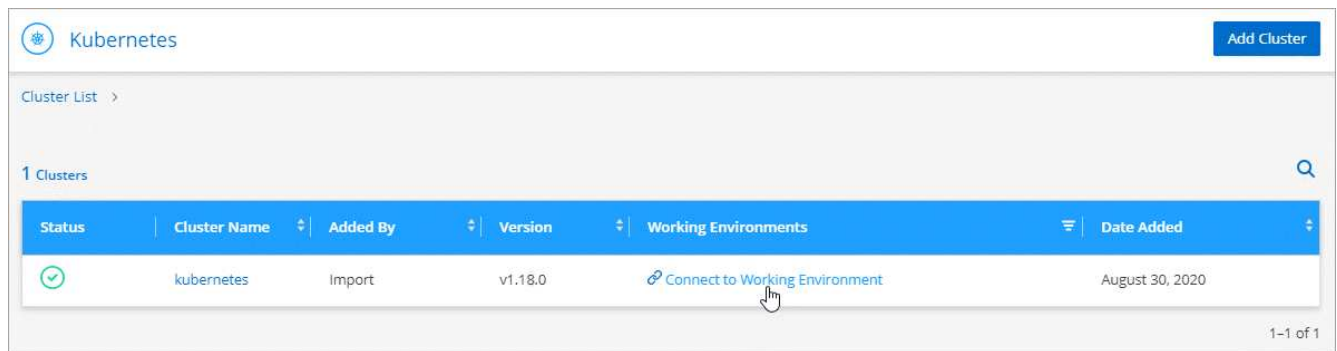
Cloud Manager adds the Kubernetes cluster. You can now connect the cluster to Cloud Volumes ONTAP.

## Connecting a cluster to Cloud Volumes ONTAP

Connect a Kubernetes cluster to Cloud Volumes ONTAP so you can use Cloud Volumes ONTAP as persistent storage for containers.

## Steps


1. At the top of Cloud Manager, click **Kubernetes**.
2. Click **Connect to Working Environment** for the cluster that you just added.



3. Select a working environment and click **Continue**.
4. Choose the NetApp storage class to use as the default storage class for the Kubernetes cluster and click **Continue**.

When a user creates a persistent volume, the Kubernetes cluster can use this storage class as the backend storage by default.

5. Choose whether to use default auto export policies or whether to add a custom CIDR block.



6. Click **Add Working Environment**.

### Result

Cloud Manager connects the working environment to the cluster, which can take up to 15 minutes.

### Managing your clusters


Cloud Manager enables you to manage your Kubernetes clusters by changing the default storage class, upgrading Trident, and more.

#### Changing the default storage class

Make sure that you've set a Cloud Volumes ONTAP storage class as the default storage class so clusters use Cloud Volumes ONTAP as the backend storage.

### Steps

1. At the top of Cloud Manager, click **Kubernetes**.
2. Click the name of the Kubernetes cluster.
3. In the **Storage Classes** table, click the actions menu on the far right for the storage class that you'd like to set as the default.



4. Click **Set as Default**.



## Upgrading Trident

You can upgrade Trident from Cloud Manager when a new version of Trident is available.

### Steps

1. At the top of Cloud Manager, click **Kubernetes**.
2. Click the name of the Kubernetes cluster.
3. If a new version is available, click **Upgrade** next to the Trident version.



## Updating the kubeconfig file

If you added your cluster to Cloud Manager by importing the kubeconfig file, you can upload the latest kubeconfig file to Cloud Manager at any time. You might do this if you've updated the credentials, if you've changed users or roles, or if something changed that affects the cluster, user, namespaces, or authentication.

### Steps

1. At the top of Cloud Manager, click **Kubernetes**.
2. Click the name of the Kubernetes cluster.
3. Click **Update Kubeconfig**.
4. When prompted through your web browser, select the updated kubeconfig file and click **Open**.

### Result


Cloud Manager updates information about the Kubernetes cluster based on the latest kubeconfig file.

## Disconnecting a cluster

When you disconnect a cluster from Cloud Volumes ONTAP, you can no longer use that Cloud Volumes ONTAP system as persistent storage for containers. Existing Persistent Volumes are not deleted.

### Steps

1. At the top of Cloud Manager, click **Kubernetes**.
2. Click the name of the Kubernetes cluster.
3. In the **Working Environments** table, click the actions menu on the far right for the working environment that you want to disconnect.



4. Click **Disconnect**.

## Result


Cloud Manager disconnects the cluster from the Cloud Volumes ONTAP system.

## Removing a cluster

Remove decommissioned clusters from Cloud Manager after you disconnect all working environments from the cluster.

## Steps

1. At the top of Cloud Manager, click **Kubernetes**.
2. Click the name of the Kubernetes cluster.
3. Click **Remove Cluster**.



## Encrypting volumes with NetApp encryption solutions

Cloud Volumes ONTAP supports both NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE) with an external key manager. NVE and NAE are software-based solutions that enable (FIPS) 140-2-compliant data-at-rest encryption of volumes. [Learn more about these encryption solutions.](#)

Starting with Cloud Volumes ONTAP 9.7, new aggregates will have NAE enabled by default after you set up an

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

Cloud Volumes ONTAP doesn't support onboard key management.

### What you'll need

Your Cloud Volumes ONTAP system should be registered with NetApp support. Starting with Cloud Manager 3.7.1, a NetApp Volume Encryption license is automatically installed on each Cloud Volumes ONTAP system that is registered with NetApp Support.

- [Adding NetApp Support Site accounts to Cloud Manager](#)
- [Registering pay-as-you-go systems](#)



Cloud Manager doesn't install the NVE license on systems that reside in the China region.

### Steps

1. Review the list of supported key managers in the [NetApp Interoperability Matrix Tool](#).



Search for the **Key Managers** solution.

2. [Connect to the Cloud Volumes ONTAP CLI](#).
3. Install SSL certificates and connect to the external key management servers.

[ONTAP 9 NetApp Encryption Power Guide: Configuring external key management](#)

## Replicating data between systems

You can replicate data between working environments by choosing a one-time data replication for data transfer, or a recurring schedule for disaster recovery or long-term retention. For example, you can set up data replication from an on-prem ONTAP system to Cloud Volumes ONTAP for disaster recovery.

Cloud Manager simplifies data replication between volumes on separate systems using SnapMirror and SnapVault technologies. You simply need to identify the source volume and the destination volume, and then choose a replication policy and schedule. Cloud Manager purchases the required disks, configures relationships, applies the replication policy, and then initiates the baseline transfer between volumes.



The baseline transfer includes a full copy of the source data. Subsequent transfers contain differential copies of the source data.

Cloud Manager enables data replication between the following types of working environments:

- From a Cloud Volumes ONTAP system to another Cloud Volumes ONTAP system
- Between a Cloud Volumes ONTAP system and an on-prem ONTAP cluster
- From an on-prem ONTAP cluster to another on-prem ONTAP cluster

## Data replication requirements

Before you can replicate data, you should confirm that specific requirements are met for both Cloud Volumes ONTAP systems and ONTAP clusters.

### Version requirements

You should verify that the source and destination volumes are running compatible ONTAP versions before replicating data. For details, see the [Data Protection Power Guide](#).

### Requirements specific to Cloud Volumes ONTAP

- The instance's security group must include the required inbound and outbound rules: specifically, rules for ICMP and ports 11104 and 11105.

These rules are included in the predefined security group.

- To replicate data between two Cloud Volumes ONTAP systems in different subnets, the subnets must be routed together (this is the default setting).
- To replicate data between a Cloud Volumes ONTAP system in AWS and a system in Azure, you must have a VPN connection between the AWS VPC and the Azure VNet.

### Requirements specific to ONTAP clusters

- An active SnapMirror license must be installed.
- If the cluster is on your premises, you should have a connection from your corporate network to AWS or Azure, which is typically a VPN connection.
- ONTAP clusters must meet additional subnet, port, firewall, and cluster requirements.

For details, see the Cluster and SVM Peering Express Guide for your version of ONTAP.

## Setting up data replication between systems

You can replicate data between Cloud Volumes ONTAP systems and ONTAP clusters by choosing a one-time data replication, which can help you move data to and from the cloud, or a recurring schedule, which can help with disaster recovery or long-term retention.

### About this task


Cloud Manager supports simple, fanout, and cascade data protection configurations:

- In a simple configuration, replication occurs from volume A to volume B.
- In a fanout configuration, replication occurs from volume A to multiple destinations.
- In a cascade configuration, replication occurs from volume A to volume B and from volume B to volume C.

You can configure fanout and cascade configurations in Cloud Manager by setting up multiple data replications between systems. For example, by replicating a volume from system A to system B and then by replicating the same volume from system B to system C.

### Steps

1. On the Working Environments page, select the working environment that contains the source volume, and then drag it to the working environment to which you want to replicate the volume:



2. If the Source and Destination Peering Setup pages appear, select all of the intercluster LIFs for the cluster peer relationship.

The intercluster network should be configured so that cluster peers have *pair-wise full-mesh connectivity*, which means that each pair of clusters in a cluster peer relationship has connectivity among all of their intercluster LIFs.

These pages appear if an ONTAP cluster that has multiple LIFs is the source or destination.

3. On the Source Volume Selection page, select the volume that you want to replicate.
4. On the Destination Volume Name and Tiering page, specify the destination volume name, choose an underlying disk type, change any of the advanced options, and then click **Continue**.

If the destination is an ONTAP cluster, you must also specify the destination SVM and aggregate.

5. On the Max Transfer Rate page, specify the maximum rate (in megabytes per second) at which data can be transferred.
6. On the Replication Policy page, choose one of the default policies or click **Additional Policies**, and then select one of the advanced policies.

For help, see [Choosing a replication policy](#).

If you choose a custom backup (SnapVault) policy, the labels associated with the policy must match the labels of the Snapshot copies on the source volume. For more information, see [How backup policies work](#).

7. On the Schedule page, choose a one-time copy or a recurring schedule.

Several default schedules are available. If you want a different schedule, you must create a new schedule on the *destination* cluster using System Manager.

8. On the Review page, review your selections, and then click **Go**.

## Result

Cloud Manager starts the data replication process. You can view details about the replication in the Replication Status page.

## Managing data replication schedules and relationships

After you set up data replication between two systems, you can manage the data replication schedule and relationship from Cloud Manager.

## Steps

1. On the Working Environments page, view the replication status for all working environments in the

workspace or for a specific working environment:


Option	Action
All working environments in the workspace	At the top of Cloud Manager, click <b>Replication</b> .
A specific working environment	Open the working environment and click <b>Replications</b> .

2. Review the status of the data replication relationships to verify that they are healthy.




If the Status of a relationship is idle and the Mirror State is uninitialized, you must initialize the relationship from the destination system for the data replication to occur according to the defined schedule. You can initialize the relationship by using System Manager or the command-line interface (CLI). These states can appear when the destination system fails and then comes back online.

3. Select the menu icon next to the source volume, and then choose one of the available actions.



The following table describes the available actions:

Action	Description
Break	<p>Breaks the relationship between the source and destination volumes, and activates the destination volume for data access.</p> <p>This option is typically used when the source volume cannot serve data due to events such as data corruption, accidental deletion, or an offline state.</p> <p>For information about configuring a destination volume for data access and reactivating a source volume, see the <a href="#">ONTAP 9 Volume Disaster Recovery Express Guide</a>.</p>
Resync	<p>Reestablishes a broken relationship between volumes and resumes data replication according to the defined schedule.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-left: 20px;">  <p>When you resynchronize the volumes, the contents on the destination volume are overwritten by the contents on the source volume.</p> </div> <p>To perform a reverse resync, which resynchronizes the data from the destination volume to the source volume, see the <a href="#">ONTAP 9 Volume Disaster Recovery Express Guide</a>.</p>
Reverse Resync	<p>Reverses the roles of the source and destination volumes. Contents from the original source volume are overwritten by contents of the destination volume. This is helpful when you want to reactivate a source volume that went offline.</p> <p>Any data written to the original source volume between the last data replication and the time that the source volume was disabled is not preserved.</p>
Edit Schedule	Enables you to choose a different schedule for data replication.
Policy Info	Shows you the protection policy assigned to the data replication relationship.
Edit Max Transfer Rate	Enables you to edit the maximum rate (in kilobytes per second) at which data can be transferred.
Update	Starts an incremental transfer to update the destination volume.
Delete	Deletes the data protection relationship between the source and destination volumes, which means that data replication no longer occurs between the volumes. This action does not activate the destination volume for data access. This action also deletes the cluster peer relationship and the storage virtual machine (SVM) peer relationship, if there are no other data protection relationships between the systems.

## Result

After you select an action, Cloud Manager updates the relationship or schedule.

## Choosing a replication policy

You might need help choosing a replication policy when you set up data replication in Cloud Manager. A replication policy defines how the storage system replicates data from a source volume to a destination volume.

### What replication policies do

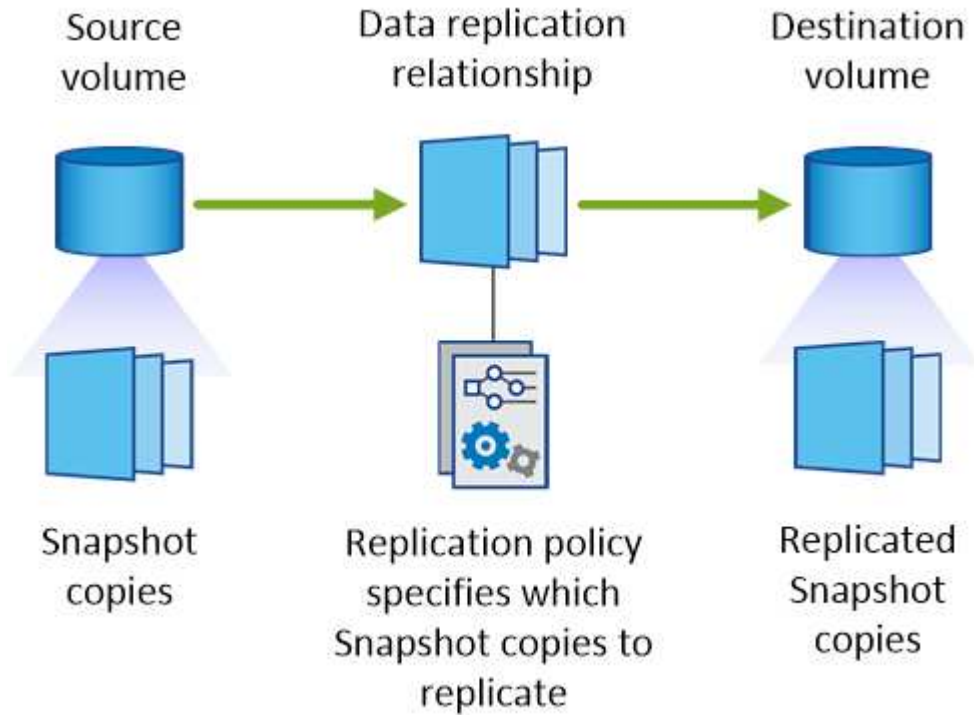
The ONTAP operating system automatically creates backups called Snapshot copies. A Snapshot copy is a read-only image of a volume that captures the state of the file system at a point in time.

When you replicate data between systems, you replicate Snapshot copies from a source volume to a destination volume. A replication policy specifies which Snapshot copies to replicate from the source volume to the destination volume.



Replication policies are also referred to as *protection* policies because they are powered by SnapMirror and SnapVault technologies, which provide disaster recovery protection and disk-to-disk backup and recovery.

The following image shows the relationship between Snapshot copies and replication policies:



### Types of replication policies

There are three types of replication policies:

- A *Mirror* policy replicates newly created Snapshot copies to a destination volume.

You can use these Snapshot copies to protect the source volume in preparation for disaster recovery or for one-time data replication. You can activate the destination volume for data access at any time.

- A *Backup* policy replicates specific Snapshot copies to a destination volume and typically retains them for a longer period of time than you would on the source volume.


You can restore data from these Snapshot copies when data is corrupted or lost, and retain them for standards compliance and other governance-related purposes.

- A *Mirror and Backup* policy provides both disaster recovery and long-term retention.


Each system includes a default Mirror and Backup policy, which works well for many situations. If you find that you need custom policies, you can create your own using System Manager.

The following images show the difference between the Mirror and Backup policies. A Mirror policy mirrors the Snapshot copies available on the source volume.





A Backup policy typically retains Snapshot copies longer than they are retained on the source volume:



### How Backup policies work

Unlike Mirror policies, Backup (SnapVault) policies replicate specific Snapshot copies to a destination volume. It is important to understand how Backup policies work if you want to use your own policies instead of the default policies.

#### Understanding the relationship between Snapshot copy labels and Backup policies

A Snapshot policy defines how the system creates Snapshot copies of volumes. The policy specifies when to create the Snapshot copies, how many copies to retain, and how to label them. For example, a system might create one Snapshot copy every day at 12:10 a.m., retain the two most recent copies, and label them "daily".

A Backup policy includes rules that specify which labeled Snapshot copies to replicate to a destination volume and how many copies to retain. The labels defined in a Backup policy must match one or more labels defined in a Snapshot policy. Otherwise, the system cannot replicate any Snapshot copies.

For example, a Backup policy that includes the labels "daily" and "weekly" results in replication of Snapshot

copies that include only those labels. No other Snapshot copies are replicated, as shown in the following image:

### **Default policies and custom policies**

The default Snapshot policy creates hourly, daily, and weekly Snapshot copies, retaining six hourly, two daily, and two weekly Snapshot copies.

You can easily use a default Backup policy with the default Snapshot policy. The default Backup policies replicate daily and weekly Snapshot copies, retaining seven daily and 52 weekly Snapshot copies.

If you create custom policies, the labels defined by those policies must match. You can create custom policies using System Manager.

## **Data replication from NetApp HCI to Cloud Volumes ONTAP**

If you're trying to replicate data from NetApp HCI to Cloud Volumes ONTAP, you can do so on a NetApp HCI system running NetApp Element software using SnapMirror. Alternatively, you can replicate data on volumes created on an ONTAP Select system running as a virtual guest in a NetApp HCI solution to Cloud Volumes ONTAP.

Refer to the following technical reports for details:

- [Technical Report 4641: NetApp HCI Data Protection](#)
- [Technical Report 4651: NetApp SolidFire SnapMirror Architecture and Configuration](#)

## **Monitor performance**

### **Learn about the Monitoring service**

By leveraging the [NetApp Cloud Insights service](#), Cloud Manager gives you insights into the health and performance of your Cloud Volumes ONTAP instances and helps you troubleshoot and optimize the performance of your cloud storage environment.

### **Features**

- Automatically monitor all volumes
- View volume performance data in terms of IOPS, throughput, and latency
- Identify performance issues to minimize impact on your users and apps

### **Supported cloud providers**

The Monitoring service is supported with Cloud Volumes ONTAP for AWS.

### **Cost**

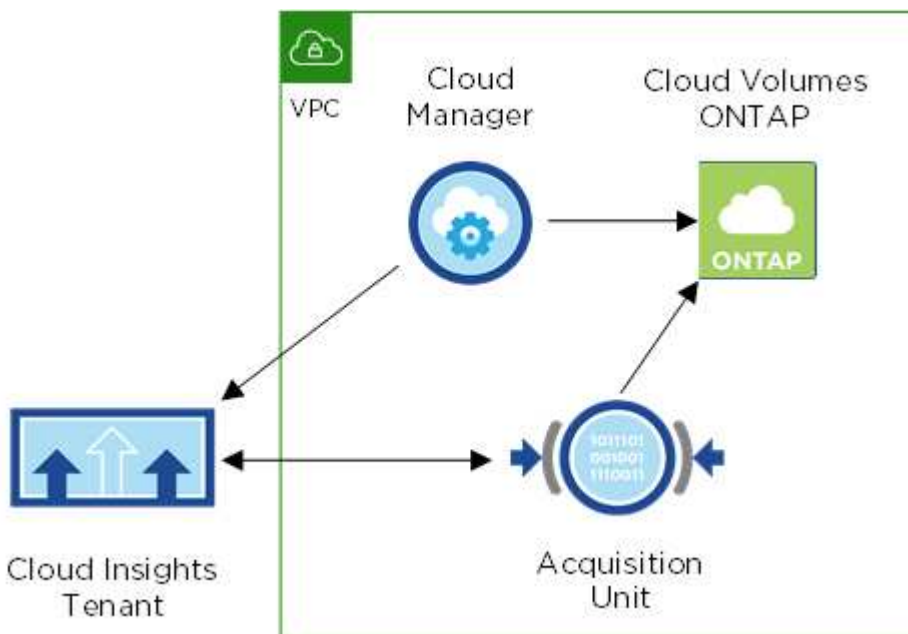
Monitoring is currently available as a Preview. Activation is free, but Cloud Manager launches a virtual machine in your VPC to facilitate monitoring. This VM results in charges from your cloud provider.

## How Cloud Insights works with Cloud Manager

At a high-level, Cloud Insights integration with Cloud Manager works like this:

1. You enable the Monitoring service on Cloud Volumes ONTAP.
2. Cloud Manager configures your environment. It does the following:
  1. Creates a Cloud Insights tenant (also called *environment*) and associates all users in your Cloud Central account to the tenant.
  2. Enables a 30-day free trial of Cloud Insights.
  3. Deploys a virtual machine in your VPC called an Acquisition Unit, which facilitates monitoring of volumes (this is the VM mentioned in the Cost section above).
  4. Connects the Acquisition Unit to Cloud Volumes ONTAP and to the Cloud Insights tenant.
3. In Cloud Manager, you click Monitoring and use the performance data to troubleshoot and optimize performance.

The following image shows the relationship between these components:



### The Acquisition Unit

When you enable Monitoring, Cloud Manager deploys an Acquisition Unit in the same subnet as the Connector.

An *Acquisition Unit* collects performance data from Cloud Volumes ONTAP and sends it to the Cloud Insights tenant. Cloud Manager then queries that data and presents it to you.

Note the following about the Acquisition Unit instance:

- The Acquisition Unit runs on a t3.xlarge instance with a 100 GB GP2 volume.
- The instance is named *AcquisitionUnit* with a generated hash (UUID) concatenated to it. For example: *AcquisitionUnit-FAN7FqeH*
- Only one Acquisition Unit is deployed per Connector.

- The instance must be running to access performance information in the Monitoring tab.

### Cloud Insights tenant

Cloud Manager sets up a *tenant* for you when you enable Monitoring. A Cloud Insights tenant enables you to access the performance data that the Acquisition Unit collects. The tenant is a secure data partition within the NetApp Cloud Insights service.

### Cloud Insights web interface

The Monitoring tab in Cloud Manager provides basic performance data for your volumes. You can go to the Cloud Insights web interface from your browser to perform more in-depth monitoring and to configure alerts for your Cloud Volumes ONTAP systems.

### Free trial and subscription

Cloud Manager enables a 30-day free trial of Cloud Insights to provide performance data within Cloud Manager and for you to explore the features that Cloud Insights Standard Edition has to offer.

You need to subscribe by the end of the free trial or your Cloud Insights tenant will eventually be deleted. You can subscribe to either the Basic, Standard, or Premium edition to continue using the Monitoring feature within Cloud Manager.

[Learn how to subscribe to Cloud Insights.](#)

## Monitoring Cloud Volumes ONTAP in AWS

Complete a few steps to get started monitoring Cloud Volumes ONTAP performance.

### Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



#### Verify support for your configuration

You need a new installation of Cloud Manager 3.8.4 or later in AWS, Cloud Volumes ONTAP in AWS, and you must be a new Cloud Insights customer.



#### Enable Monitoring on your new or existing system

- New working environments: Be sure to keep Monitoring enabled when you create the working environment (it's enabled by default).
- Existing working environments: Select a working environment and click **Start Monitoring**.



#### View performance data

Click **Monitoring** and view performance data for your volumes.



## Subscribe to Cloud Insights

Subscribe before your 30-day free trial ends to continue seeing performance data within Cloud Manager and Cloud Insights. [Learn how to subscribe.](#)

### Requirements

Read the following requirements to make sure that you have a supported configuration.

#### Supported Cloud Manager versions

You need a new installation of Cloud Manager 3.8.4 or later. A new installation is needed because a new infrastructure is required to enable the Monitoring service. This infrastructure is available starting with new installations of Cloud Manager 3.8.4.

#### Supported Cloud Volumes ONTAP versions

Any version of Cloud Volumes ONTAP in AWS.

#### Cloud Insights requirement

You must be a new Cloud Insights customer. Monitoring isn't supported if you already have a Cloud Insights tenant.

#### Email address for Cloud Central

The email address for your Cloud Central user account should be your business email address. Free email domains like gmail and hotmail aren't supported when creating a Cloud Insights tenant.

#### Networking for the Acquisition Unit

The Acquisition Unit uses 2-way/mutual authentication to connect to the Cloud Insights server. The client certificate must be passed to the Cloud Insights server to be authenticated. To accomplish this, the proxy must be set up to forward the http request to the Cloud Insights server without decrypting the data.

The Acquisition Unit uses the following two endpoints to communicate with Cloud Insights. If you have a firewall between the Acquisition Unit server and Cloud Insights, you need these endpoints when configuring firewall rules:

```
https://aLOGIN.<Cloud Insights Domain>  
https://<your-tenant-ID>.<Cloud Insights Domain>
```

For example:

```
https://aLOGIN.c01.cloudinsights.netapp.com  
https://cg0c586a-ee05-45rb-a5ac-  
333b5ae7718d7.c01.cloudinsights.netapp.com
```

Contact us through the in-product chat if you need help identifying your Cloud Insights domain and tenant ID.

## Networking for the Connector

Similar to the Acquisition Unit, the Connector must have outbound connectivity to the Cloud Insights tenant. But the endpoint that the Connector contacts is slightly different. It contacts the tenant host URL using the shortened tenant ID:

```
https://<your-short-tenant-ID>.<Cloud Insights Domain>
```

For example:

```
https://abcd12345.c01.cloudinsights.netapp.com
```

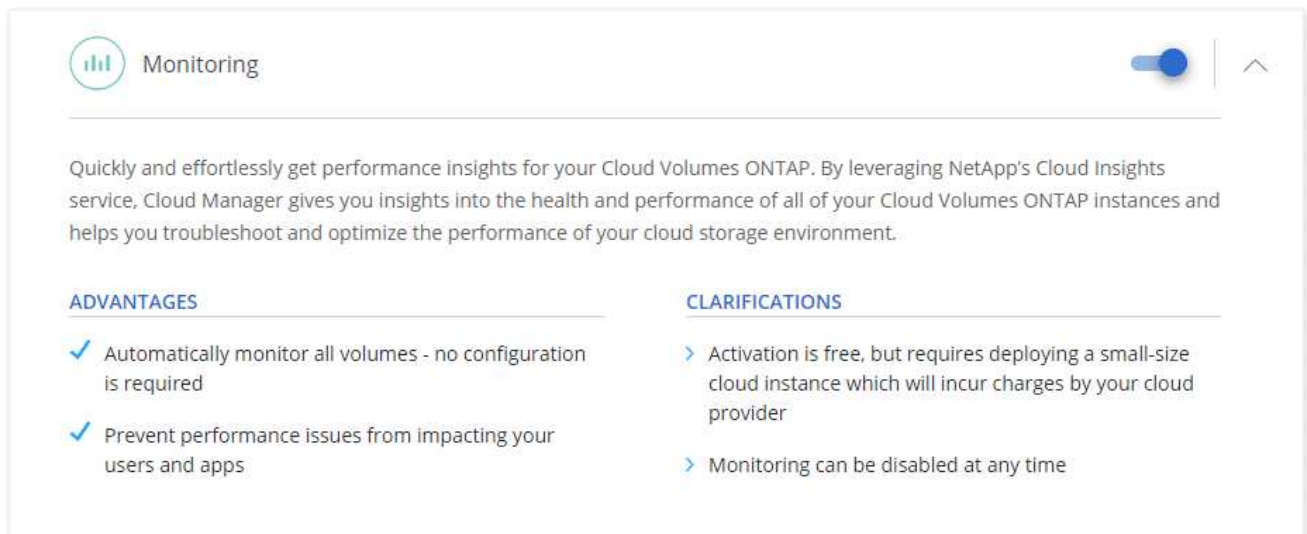
Again, you can contact us through the in-product chat if you need help identifying the tenant host URL.

## Enabling monitoring on a new system

The Monitoring service is enabled by default in the working environment wizard. Be sure to keep the option enabled.

### Steps

1. Click **Create Cloud Volumes ONTAP**.
2. Select Amazon Web Services as the cloud provider and then choose a single node or HA system.
3. Fill out the Details & Credentials page.
4. On the Services page, leave the service enabled and click **Continue**.



**Monitoring**

Quickly and effortlessly get performance insights for your Cloud Volumes ONTAP. By leveraging NetApp's Cloud Insights service, Cloud Manager gives you insights into the health and performance of all of your Cloud Volumes ONTAP instances and helps you troubleshoot and optimize the performance of your cloud storage environment.

ADVANTAGES	CLARIFICATIONS
<ul style="list-style-type: none"><li>✓ Automatically monitor all volumes - no configuration is required</li><li>✓ Prevent performance issues from impacting your users and apps</li></ul>	<ul style="list-style-type: none"><li>&gt; Activation is free, but requires deploying a small-size cloud instance which will incur charges by your cloud provider</li><li>&gt; Monitoring can be disabled at any time</li></ul>


## Enabling monitoring on an existing system

Enable monitoring at any time from the working environment.

### Steps

1. At the top of Cloud Manager, click **Working Environments**.

2. Select a working environment.
3. In the pane on the right, click **Start Monitoring**.




## Monitoring your volumes

Monitor performance by viewing IOPS, throughput, and latency for each of your volumes.


### Steps

1. At the top of Cloud Manager, click **Monitoring**.
2. Filter the contents of the dashboard to get the information that you need.
  - Select a specific working environment.
  - Select a different timeframe.
  - Select a specific SVM.
  - Search for a specific volume.

The following image highlights each of these options:



3. Click a volume in the table to expand the row and view a timeline for IOPS, throughput, and latency.



4. Use the data to identify performance issues to minimize impact on your users and apps.


### Getting more information from Cloud Insights

The Monitoring tab in Cloud Manager provides basic performance data for your volumes. You can go to the Cloud Insights web interface from your browser to perform more in-depth monitoring and to configure alerts for your Cloud Volumes ONTAP systems.

#### Steps

1. At the top of Cloud Manager, click **Monitoring**.
2. Click the **Cloud Insights** link.





## Result

Cloud Insights open in a new browser tab. If you need help, refer to the [Cloud Insights documentation](#).

## Disabling monitoring

If you no longer want to monitor Cloud Volumes ONTAP, you can disable the service at any time.



If you disable monitoring from each of your working environments, you'll need to delete the EC2 instance yourself. The instance is named *AcquisitionUnit* with a generated hash (UUID) concatenated to it. For example: *AcquisitionUnit-FAN7FqeH*

## Steps


1. At the top of Cloud Manager, click **Working Environments**.
2. Select a working environment.
3. In the pane on the right, click the icon and select **Deactivate Scan**.

## Improving protection against ransomware

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.

## Steps

1. From the working environment, click the **Ransomware** icon.



## 2. Implement the NetApp solution for ransomware:

- a. Click **Activate Snapshot Policy**, if you have volumes that do not have a Snapshot policy enabled.

NetApp Snapshot technology provides the industry's best solution for ransomware remediation. The key to a successful recovery is restoring from uninfected backups. 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.

- b. Click **Activate FPolicy** to enable ONTAP's FPolicy solution, which can block file operations based on a file's extension.

This preventative solution improves protection from ransomware attacks by blocking common ransomware file types.

### Ransomware Protection

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

#### 1 Enable Snapshot Copy Protection

50 %  
Protection

1 Volumes without a Snapshot Policy

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

Activate Snapshot Policy

#### 2 Block Ransomware File Extensions

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

View Denied File Names

Activate FPolicy

## Administer

### Registering pay-as-you-go systems

Support from NetApp is included with Cloud Volumes ONTAP Explore, Standard, and Premium systems, but you must first activate support by registering the systems with

## NetApp.

### Steps

1. If you have not yet added your NetApp Support Site account to Cloud Manager, go to **Account Settings** and add it now.

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

2. On the Working Environments page, double-click the name of the system that you want to register.
3. Click the menu icon and then click **Support registration**:



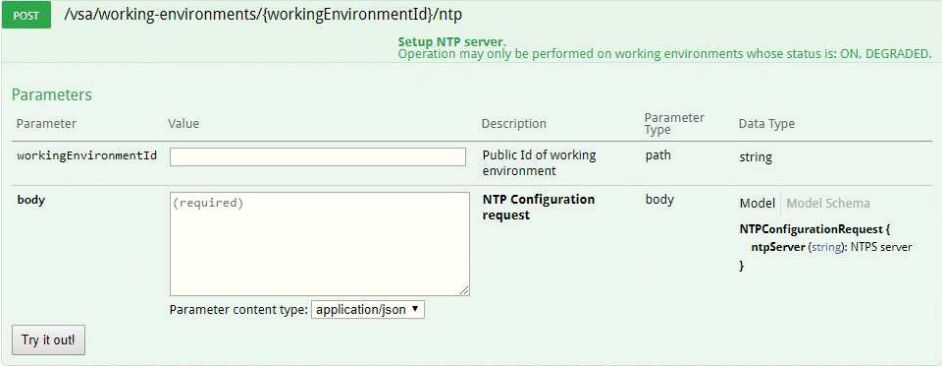
4. Select a NetApp Support Site account and click **Register**.

### Result

Cloud Manager registers the system with NetApp.

## Setting up Cloud Volumes ONTAP

After you deploy Cloud Volumes ONTAP, you can set it up by synchronizing the system time using NTP and by performing a few optional tasks from either System Manager or the CLI.

Task	Description
<p>Synchronize the system time using NTP</p>	<p>Specifying an NTP server synchronizes the time between the systems in your network, which can help prevent issues due to time differences.</p> <p>Specify an NTP server using the Cloud Manager API or from the user interface when you set up a CIFS server.</p> <ul style="list-style-type: none"> <li>• <a href="#">Modifying the CIFS server</a></li> <li>• <a href="#">Cloud Manager API Developer Guide</a></li> </ul> <p>For example, here's the API for a single-node system in AWS:</p> 
<p>Optional: Configure AutoSupport</p>	<p>AutoSupport proactively monitors the health of your system and automatically sends messages to NetApp technical support by default.</p> <p>If the Account Admin added a proxy server to Cloud Manager before you launched your instance, Cloud Volumes ONTAP is configured to use that proxy server for AutoSupport messages.</p> <p>You should test AutoSupport to ensure that it can send messages. For instructions, see the System Manager Help or the <a href="#">ONTAP 9 System Administration Reference</a>.</p>
<p>Optional: Configure Cloud Manager as the AutoSupport proxy</p>	<p>If your environment requires a proxy server to send AutoSupport messages, you can configure Cloud Manager to act as the proxy. No configuration for Cloud Manager is required, other than internet access. You simply need to go to the CLI for Cloud Volumes ONTAP and run the following command:</p> <pre data-bbox="548 1495 1484 1633">system node autosupport modify -proxy-url &lt;cloud-manager-ip-address&gt;</pre>
<p>Optional: Configure EMS</p>	<p>The Event Management System (EMS) collects and displays information about events that occur on Cloud Volumes ONTAP systems. To receive event notifications, you can set event destinations (email addresses, SNMP trap hosts, or syslog servers) and event routes for a particular event severity.</p> <p>You can configure EMS using the CLI. For instructions, see the <a href="#">ONTAP 9 EMS Configuration Express Guide</a>.</p>

Task	Description
<p>Optional: Create an SVM management network interface (LIF) for HA systems in multiple AWS Availability Zones</p>	<p>A storage virtual machine (SVM) management network interface (LIF) is required if you want to use SnapCenter or SnapDrive for Windows with an HA pair. The SVM management LIF must use a <i>floating</i> IP address when using an HA pair across multiple AWS Availability Zones.</p> <p>Cloud Manager prompts you to specify the floating IP address when you launch the HA pair. If you did not specify the IP address, you can create the SVM Management LIF yourself from System Manager or the CLI. The following example shows how to create the LIF from the CLI:</p> <pre data-bbox="548 495 1485 751">network interface create -vserver svm_cloud -lif svm_mgmt -role data -data-protocol none -home-node cloud-01 -home-port e0a -address 10.0.2.126 -netmask 255.255.255.0 -status-admin up -firewall -policy mgmt</pre>
<p>Optional: Change the backup location of configuration files</p>	<p>Cloud Volumes ONTAP automatically creates configuration backup files that contain information about the configurable options that it needs to operate properly.</p> <p>By default, Cloud Volumes ONTAP backs up the files to the Connector host every eight hours. If you want to send the backups to an alternate location, you can change the location to an FTP or HTTP server in your data center or in AWS. For example, you might already have a backup location for your FAS storage systems.</p> <p>You can change the backup location using the CLI. See the <a href="#">ONTAP 9 System Administration Reference</a>.</p>

## Managing BYOL licenses for Cloud Volumes ONTAP

Add a Cloud Volumes ONTAP BYOL system license to add additional capacity, update an existing system license, and manage BYOL licenses for Backup to Cloud.

### Managing system licenses

You can purchase multiple licenses for a Cloud Volumes ONTAP BYOL system to allocate more than 368 TB of capacity. For example, you might purchase two licenses to allocate up to 736 TB of capacity to Cloud Volumes ONTAP. Or you could purchase four licenses to get up to 1.4 PB.

The number of licenses that you can purchase for a single node system or HA pair is unlimited.

### Obtaining a system license file

In most cases, Cloud Manager can automatically obtain your license file using your NetApp Support Site account. But if it can't, then you'll need to manually upload the license file. If you don't have the license file, you can obtain it from netapp.com.

### Steps

1. Go to the [NetApp License File Generator](#) and log in using your NetApp Support Site credentials.
2. Enter your password, choose your product, enter the serial number, confirm that you have read and accepted the privacy policy, and then click **Submit**.

### Example


3. Choose whether you want to receive the serialnumber.NLF JSON file through email or direct download.

### Adding a new system license

Add a new BYOL system license at any time to allocate an additional 368 TB of capacity to your Cloud Volumes ONTAP BYOL system.

### Steps

1. In Cloud Manager, open the Cloud Volumes ONTAP BYOL working environment.
2. Click the menu icon and then click **License**.



3. Click **Add CVO System License**.



4. Choose to enter the serial number or to upload the license file.

5. Click **Add License**.

### Result

Cloud Manager installs the new license file on the Cloud Volumes ONTAP system.

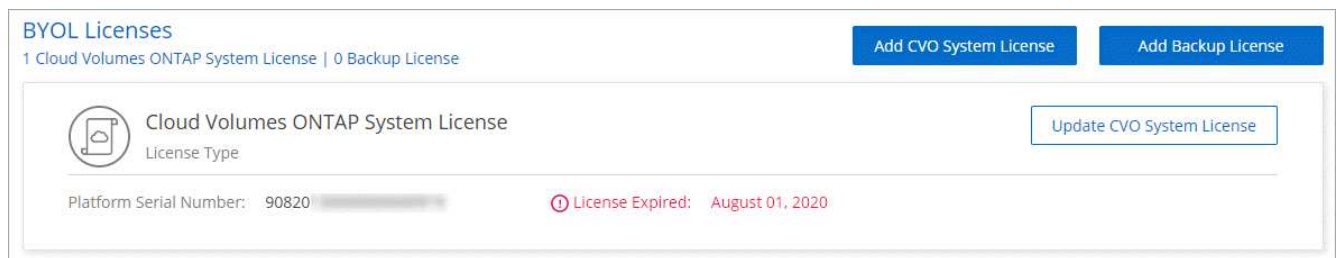
### Updating a system license

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 can't access the license file over the secure internet connection, you can obtain the file yourself and then manually upload the file to Cloud Manager.

### Steps

1. In Cloud Manager, open the Cloud Volumes ONTAP BYOL working environment.
2. Click the menu icon and then click **License**.
3. Click **Update CVO System License**.



4. Click **Upload File** and select the license file.
5. Click **Update License**.

### Result


Cloud Manager updates the license on the Cloud Volumes ONTAP system.

### Adding and updating your Backup BYOL license

You use the BYOL Licenses page to add or update your Backup BYOL license.

### Steps

1. In Cloud Manager, open the Cloud Volumes ONTAP BYOL working environment.
2. Click the menu icon and then click **License**.



3. Click **Add Backup License** or **Update Backup License** depending on whether you are adding a new license or updating an existing license.


**Total License Information**

Instance Type :	m5.2xlarge	Total Attached EBS Capacity :	200 TB	Total Used Tiering Capacity :	60 TB
Total License Limit :	368 TB	Total Used EBS Capacity :	180 TB	Total Allocated ONTAP Capacity :	100 TB
Total Backup Capacity Limit :	368 TB	Total Used Backup Capacity :	200 TB		

**BYOL Licenses**

1 Cloud Volumes ONTAP System License | 1 Backup License

[Add CVO System License](#) [Add Backup License](#)


 **Cloud Volumes ONTAP System License**  
License Type [Update CVO System License](#)

---

Platform Serial Number Node 1 : 9012013000000000020      License Expiry: April 10, 2021

---

Platform Serial Number Node 2 : 9012013000000000021      License Expiry: April 10, 2021

 **Backup License**  
License Type [Update Backup License](#)

---

Platform Serial Number : 9012013000000000022      License Expiry: April 10, 2021      License Capacity Limit : 368 TB (Used Capacity 200 TB)

4. Enter the license information and click **Add License**:
  - If you have the serial number, select the **Enter Backup BYOL Serial Number** option and enter the serial number.



- If you have the backup license file, select the **Upload Backup BYOL License** option and follow the prompts to attach the file.

**Add Backup License**

A Backup license enables Backup to Cloud for a certain period of time and for a maximum amount backup space.

Enter Backup BYOL Serial Number     Upload Backup BYOL License

Enter Backup BYOL Serial Number

### Result

Cloud Manager adds or updates the license so that your Backup to Cloud service is active.

## Updating Cloud Volumes ONTAP software

Cloud Manager includes several options that you can use to upgrade to the current Cloud Volumes ONTAP release or to downgrade Cloud Volumes ONTAP to an earlier release. You should prepare Cloud Volumes ONTAP systems before you upgrade or downgrade the software.

### Software updates must be completed by Cloud Manager

Upgrades of Cloud Volumes ONTAP must be completed from Cloud Manager. You should not upgrade Cloud Volumes ONTAP by using System Manager or the CLI. Doing so can impact system stability.

### Ways to update Cloud Volumes ONTAP

Cloud Manager displays a notification in Cloud Volumes ONTAP working environments when a new version of Cloud Volumes ONTAP is available:

The screenshot shows the Cloud Manager interface for a system named 'cloudvolumesontap1'. At the top, there is a 'Visual View' dropdown menu. Below it, the system name 'cloudvolumesontap1' is displayed with a status of 'On | AWS'. A red box highlights a notification titled 'NOTIFICATIONS' with a star icon and the text 'New version available'. Below the notification, there are two service cards: 'Cloud Compliance' with a status of 'On' and 'No Personal Files Found', and 'Backup to S3' with a status of 'On' and '3 Volumes Backed Up'.

You can start the upgrade process from this notification, which automates the process by obtaining the software image from an S3 bucket, installing the image, and then restarting the system. For details, see [Upgrading Cloud Volumes ONTAP from Cloud Manager notifications](#).



For HA systems in AWS, Cloud Manager might upgrade the HA mediator as part of the upgrade process.

#### Advanced options for software updates

Cloud Manager also provides the following advanced options for updating Cloud Volumes ONTAP software:

- Software updates using an image on an external URL

This option is helpful if Cloud Manager cannot access the S3 bucket to upgrade the software, if you were provided with a patch, or if you want to downgrade the software to a specific version.

For details, see [Upgrading or downgrading Cloud Volumes ONTAP by using an HTTP or FTP server](#).

- Software updates using the alternate image on the system

You can use this option to downgrade to the previous version by making the alternate software image the default image. This option is not available for HA pairs.

For details, see [Downgrading Cloud Volumes ONTAP by using a local image](#).

## Preparing to update Cloud Volumes ONTAP software

Before performing an upgrade or downgrade, you must verify that your systems are ready and make any required configuration changes.

- [Planning for downtime](#)
- [Reviewing version requirements](#)
- [Verifying that automatic giveback is still enabled](#)
- [Suspending SnapMirror transfers](#)
- [Verifying that aggregates are online](#)

### Planning for downtime

When you upgrade a single-node system, the upgrade process takes the system offline for up to 25 minutes, during which I/O is interrupted.

Upgrading an HA pair is nondisruptive and I/O is uninterrupted. During this nondisruptive upgrade process, each node is upgraded in tandem to continue serving I/O to clients.

### Reviewing version requirements

The version of ONTAP that you can upgrade or downgrade to varies based on the version of ONTAP currently running on your system.

To understand version requirements, refer to [ONTAP 9 Documentation: Cluster update requirements](#).

### Verifying that automatic giveback is still enabled

Automatic giveback must be enabled on a Cloud Volumes ONTAP HA pair (this is the default setting). If it isn't, then the operation will fail.

[ONTAP 9 Documentation: Commands for configuring automatic giveback](#)

### Suspending SnapMirror transfers

If a Cloud Volumes ONTAP system has active SnapMirror relationships, it is best to suspend transfers before you update the Cloud Volumes ONTAP software. Suspending the transfers prevents SnapMirror failures. You must suspend the transfers from the destination system.

### About this task

These steps describe how to use System Manager for version 9.3 and later.

### Steps

1. [Log in to System Manager](#) from the destination system.
2. Click **Protection > Relationships**.
3. Select the relationship and click **Operations > Quiesce**.

### Verifying that aggregates are online

Aggregates for Cloud Volumes ONTAP must be online before you update the software. Aggregates should be online in most configurations, but if they are not, then you should bring them online.

## About this task

These steps describe how to use System Manager for version 9.3 and later.

### Steps

1. In the working environment, click the menu icon, and then click **Advanced > Advanced allocation**.
2. Select an aggregate, click **Info**, and then verify that the state is online.

<b>aggr1</b>		
Aggregate Capacity:	88.57 GB	
-----		
Used Aggregate Capacity:	1.07 GB	
-----		
Volumes:	2	▼
-----		
AWS Disks:	1	▼
-----		
State:	online	

3. If the aggregate is offline, use System Manager to bring the aggregate online:
  - a. [Log in to System Manager](#).
  - b. Click **Storage > Aggregates & Disks > Aggregates**.
  - c. Select the aggregate, and then click **More Actions > Status > Online**.

## Upgrading Cloud Volumes ONTAP from Cloud Manager notifications

Cloud Manager notifies you when a new version of Cloud Volumes ONTAP is available. Click the notification to start the upgrade process.

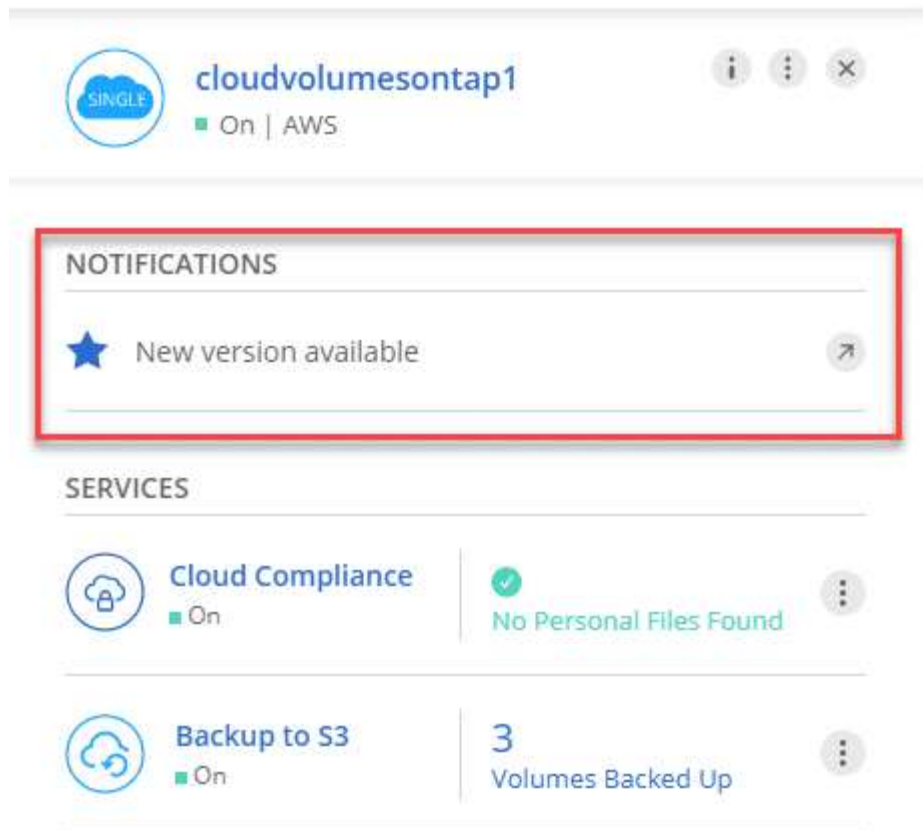
### Before you begin




Cloud Manager operations such as volume or aggregate creation must not be in progress for the Cloud Volumes ONTAP system.

### Steps

1. Click **Working Environments**.
2. Select a working environment.


A notification appears in the right pane if a new version is available:






**cloudvolumesontap1**   



■ On | AWS

**NOTIFICATIONS**

★ New version available 

**SERVICES**

 **Cloud Compliance** ■ On  No Personal Files Found 

 **Backup to S3** ■ On **3** Volumes Backed Up 

3. If a new version is available, click **Upgrade**.
4. In the Release Information page, click the link to read the Release Notes for the specified version, and then select the **I have read...** check box.
5. In the End User License Agreement (EULA) page, read the EULA, and then select **I read and approve the EULA**.
6. In the Review and Approve page, read the important notes, select **I understand...**, and then click **Go**.

### Result

Cloud Manager starts the software upgrade. You can perform actions on the working environment once the software update is complete.

### After you finish

If you suspended SnapMirror transfers, use System Manager to resume the transfers.

### Upgrading or downgrading Cloud Volumes ONTAP by using an HTTP or FTP server

You can place the Cloud Volumes ONTAP software image on an HTTP or FTP server and then initiate the software update from Cloud Manager. You might use this option if Cloud Manager cannot access the S3 bucket to upgrade the software or if you want to downgrade the software.

### Steps

1. Set up an HTTP server or FTP server that can host the Cloud Volumes ONTAP software image.

2. If you have a VPN connection to the virtual network, you can place the Cloud Volumes ONTAP software image on an HTTP server or FTP server in your own network. Otherwise, you must place the file on an HTTP server or FTP server in the cloud.
3. If you use your own security group for Cloud Volumes ONTAP, ensure that the outbound rules allow HTTP or FTP connections so Cloud Volumes ONTAP can access the software image.



The predefined Cloud Volumes ONTAP security group allows outbound HTTP and FTP connections by default.

4. Obtain the software image from [the NetApp Support Site](#).
5. Copy the software image to the directory on the HTTP or FTP server from which the file will be served.
6. From the working environment in Cloud Manager, click the menu icon, and then click **Advanced > Update Cloud Volumes ONTAP**.
7. On the update software page, choose **Select an image available from a URL**, enter the URL, and then click **Change Image**.
8. Click **Proceed** to confirm.

### Result

Cloud Manager starts the software update. You can perform actions on the working environment once the software update is complete.

### After you finish

If you suspended SnapMirror transfers, use System Manager to resume the transfers.

### Downgrading Cloud Volumes ONTAP by using a local image

Transitioning Cloud Volumes ONTAP to an earlier release in the same release family (for example, 9.5 to 9.4) is referred to as a downgrade. You can downgrade without assistance when downgrading new or test clusters, but you should contact technical support if you want to downgrade a production cluster.

Each Cloud Volumes ONTAP system can hold two software images: the current image that is running, and an alternate image that you can boot. Cloud Manager can change the alternate image to be the default image. You can use this option to downgrade to the previous version of Cloud Volumes ONTAP, if you are experiencing issues with the current image.

### About this task

This downgrade process is available for single Cloud Volumes ONTAP systems only. It is not available for HA pairs.

### Steps

1. From the working environment, click the menu icon, and then click **Advanced > Update Cloud Volumes ONTAP**.
2. On the update software page, select the alternate image, and then click **Change Image**.
3. Click **Proceed** to confirm.

### Result

Cloud Manager starts the software update. You can perform actions on the working environment once the software update is complete.

### After you finish

If you suspended SnapMirror transfers, use System Manager to resume the transfers.

## Modifying Cloud Volumes ONTAP systems

You might need to change the configuration of Cloud Volumes ONTAP systems as your storage needs change. For example, you can change between pay-as-you-go configurations, change the instance or VM type, and more.

### Changing the instance or machine type for Cloud Volumes ONTAP

You can choose from several instance or machine types when you launch Cloud Volumes ONTAP in AWS, Azure, or GCP. You can change the instance or machine type at any time if you determine that it is undersized or oversized for your needs.

#### About this task

- Automatic giveback must be enabled on a Cloud Volumes ONTAP HA pair (this is the default setting). If it isn't, then the operation will fail.

[ONTAP 9 Documentation: Commands for configuring automatic giveback](#)

- Changing the instance or machine type affects cloud provider service charges.
- The operation restarts Cloud Volumes ONTAP.

For single node systems, I/O is interrupted.

For HA pairs, the change is nondisruptive. HA pairs continue to serve data.



Cloud Manager gracefully changes one node at a time by initiating takeover and waiting for give back. NetApp's QA team tested both writing and reading files during this process and didn't see any issues on the client side. As connections changed, we did see retries on the I/O level, but the application layer overcame these short "re-wire" of NFS/CIFS connections.

#### Steps

1. From the working environment, click the menu icon, and then click **Change license or instance** for AWS, **Change license or VM** for Azure, or **Change license or machine** for GCP.
2. If you are using a pay-as-you-go configuration, you can optionally choose a different license.
3. Select an instance or machine type, select the check box to confirm that you understand the implications of the change, and then click **OK**.

#### Result

Cloud Volumes ONTAP reboots with the new configuration.

### Changing between pay-as-you-go configurations

After you launch pay-as-you-go Cloud Volumes ONTAP systems, you can change between the Explore, Standard, and Premium configurations at any time by modifying the license. Changing the license increases or decreases the raw capacity limit and enables you to choose from different AWS instance types or Azure virtual machine types.



In GCP, a single machine type is available for each pay-as-you-go configuration. You can't choose between different machine types.

### About this task

Note the following about changing between pay-as-you-go licenses:

- The operation restarts Cloud Volumes ONTAP.  
For single node systems, I/O is interrupted.  
For HA pairs, the change is nondisruptive. HA pairs continue to serve data.
- Changing the instance or machine type affects cloud provider service charges.

### Steps

1. From the working environment, click the menu icon, and then click **Change license or instance** for AWS, **Change license or VM** for Azure, or **Change license or machine** for GCP.
2. Select a license type and an instance type or machine type, select the check box to confirm that you understand the implications of the change, and then click **OK**.

### Result

Cloud Volumes ONTAP reboots with the new license, instance type or machine type, or both.

### Moving to an alternate Cloud Volumes ONTAP configuration

If you want to switch between a pay-as-you-go subscription and a BYOL subscription or between a single Cloud Volumes ONTAP system and an HA pair, then you need to deploy a new system and then replicate data from the existing system to the new system.

### Steps

1. Create a new Cloud Volumes ONTAP working environment.

[Launching Cloud Volumes ONTAP in AWS](#)

[Launching Cloud Volumes ONTAP in Azure](#)

[Launching Cloud Volumes ONTAP in GCP](#)

2. [Set up one-time data replication](#) between the systems for each volume that you must replicate.
3. Terminate the Cloud Volumes ONTAP system that you no longer need by [deleting the original working environment](#).

### Changing write speed to normal or high

Cloud Manager enables you to choose a write speed setting for single node Cloud Volumes ONTAP systems. The default write speed is normal. You can change to high write speed if fast write performance is required for your workload. Before you change the write speed, you should [understand the differences between the normal and high settings](#).

### About this task

- Ensure that operations such as volume or aggregate creation are not in progress.
- Be aware that this change restarts Cloud Volumes ONTAP, which means I/O is interrupted.

### Steps



1. From the working environment, click the menu icon, and then click **Advanced > Writing Speed**.
2. Select **Normal** or **High**.

If you choose High, then you'll need to read the "I understand..." statement and confirm by checking the box.

3. Click **Save**, review the confirmation message, and then click **Proceed**.

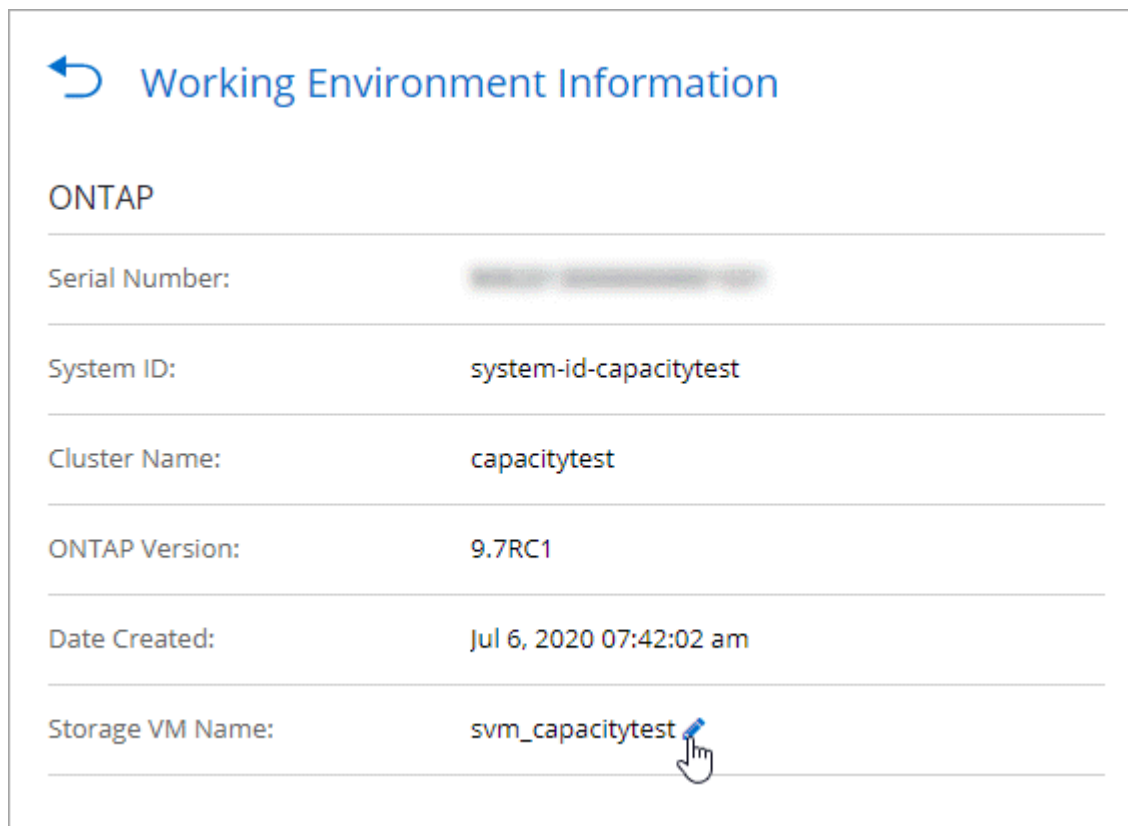
### Modifying the storage VM name

Cloud Manager automatically names the single storage VM (SVM) that it creates for Cloud Volumes ONTAP. You can modify the name of the SVM if you have strict naming standards. For example, you might want the name to match how you name the SVMs for your ONTAP clusters.

But if you created any additional SVMs for Cloud Volumes ONTAP, then you can't rename the SVMs from Cloud Manager. You'll need to do so directly from Cloud Volumes ONTAP by using System Manager or the CLI.

### Steps

1. From the working environment, click the menu icon, and then click **Information**.
2. Click the edit icon to the right of the storage VM name.



3. In the Modify SVM Name dialog box, change the name, and then click **Save**.

### Changing the password for Cloud Volumes ONTAP

Cloud Volumes ONTAP includes a cluster admin account. You can change the password for this account from Cloud Manager, if needed.



You should not change the password for the admin account through System Manager or the CLI. The password will not be reflected in Cloud Manager. As a result, Cloud Manager cannot monitor the instance properly.

### Steps

1. From the working environment, click the menu icon, and then click **Advanced > Set password**.
2. Enter the new password twice and then click **Save**.

The new password must be different than one of the last six passwords that you used.

### Changing the network MTU for c4.4xlarge and c4.8xlarge instances

By default, Cloud Volumes ONTAP is configured to use 9,000 MTU (also called jumbo frames) when you choose the c4.4xlarge instance or the c4.8xlarge instance in AWS. You can change the network MTU to 1,500 bytes if that is more appropriate for your network configuration.

#### About this task

A network maximum transmission unit (MTU) of 9,000 bytes can provide the highest maximum network throughput possible for specific configurations.

9,000 MTU is a good choice if clients in the same VPC communicate with the Cloud Volumes ONTAP system and some or all of those clients also support 9,000 MTU. If traffic leaves the VPC, packet fragmentation can occur, which degrades performance.

A network MTU of 1,500 bytes is a good choice if clients or systems outside of the VPC communicate with the Cloud Volumes ONTAP system.

### Steps

1. From the working environment, click the menu icon and then click **Advanced > Network Utilization**.
2. Select **Standard** or **Jumbo Frames**.
3. Click **Change**.

### Changing route tables associated with HA pairs in multiple AWS AZs

You can modify the AWS route tables that include routes to the floating IP addresses for an HA pair. You might do this if new NFS or CIFS clients need to access an HA pair in AWS.

### Steps

1. From the working environment, click the menu icon and then click **Information**.
2. Click **Route Tables**.
3. Modify the list of selected route tables and then click **Save**.

### Result

Cloud Manager sends an AWS request to modify the route tables.

### Managing the state of Cloud Volumes ONTAP

You can stop and start Cloud Volumes ONTAP from Cloud Manager to manage your cloud compute costs.

## Scheduling automatic shutdowns of Cloud Volumes ONTAP

You might want to shut down Cloud Volumes ONTAP during specific time intervals to lower your compute costs. Rather than do this manually, you can configure Cloud Manager to automatically shut down and then restart systems at specific times.

### About this task

When you schedule an automatic shutdown of your Cloud Volumes ONTAP system, Cloud Manager postpones the shutdown if an active data transfer is in progress. Cloud Manager shuts down the system after the transfer is complete.

This task schedules automatic shutdowns of both nodes in an HA pair.

### Steps

1. From the working environment, click the clock icon:



2. Specify the shutdown schedule:
  - a. Choose whether you want to shut down the system every day, every weekday, every weekend, or any combination of the three options.
  - b. Specify when you want to turn off the system and for how long you want it turned off.

### Example

The following image shows a schedule that instructs Cloud Manager to shut down the system every Saturday at 12:00 a.m. for 48 hours. Cloud Manager restarts the system every Monday at 12:00 a.m.

<input type="checkbox"/>	<b>Turn off every weekday</b> Mon, Tue, Wed, Thu, Fri	turn off at	08 : 00 PM	for	12	Hours (1-24)
<input checked="" type="checkbox"/>	<b>Turn off every weekend</b> Sat	turn off at	12 : 00 AM	for	48	Hours (1-48)

3. Click **Save**.

### Result

Cloud Manager saves the schedule. The clock icon changes to indicate that a schedule is set:

## Stopping Cloud Volumes ONTAP

Stopping Cloud Volumes ONTAP saves you from accruing compute costs and creates snapshots of the root and boot disks, which can be helpful for troubleshooting.

### About this task

When you stop an HA pair, Cloud Manager shuts down both nodes.

## Steps

1. From the working environment, click the **Turn off** icon.



2. Keep the option to create snapshots enabled because the snapshots can enable system recovery.
3. Click **Turn Off**.

It can take up to a few minutes to stop the system. You can restart systems at a later time from the working environment page.

## Monitoring AWS resource costs

Cloud Manager enables you to view the resource costs associated with running Cloud Volumes ONTAP in AWS. You can also see how much money you saved by using NetApp features that can reduce storage costs.

### About this task

Cloud Manager updates the costs when you refresh the page. You should refer to AWS for final cost details.

### Step

1. Verify that Cloud Manager can obtain cost information from AWS:
  - a. Ensure that the IAM policy that provides Cloud Manager with permissions includes the following actions:

```
"ce:GetReservationUtilization",  
"ce:GetDimensionValues",  
"ce:GetCostAndUsage",  
"ce:GetTags"
```

These actions are included in the latest [Cloud Manager policy](#). New systems deployed from NetApp Cloud Central automatically include these permissions.

- b. [Activate the WorkingEnvironmentId tag](#).


To track your AWS costs, Cloud Manager assigns a cost allocation tag to Cloud Volumes ONTAP instances. After you create your first working environment, activate the **WorkingEnvironmentId** tag. User-defined tags don't appear on AWS billing reports until you activate them in the Billing and Cost Management console.

2. On the Working Environments page, select a Cloud Volumes ONTAP working environment and then click **Cost**.

The Cost page displays costs for the current and previous months and shows your annual NetApp savings, if you enabled NetApp's cost-saving features on volumes.

The following image shows a sample Cost page:

Cloud Manager obtains AWS resource costs by using the AWS Cost Explorer service



## Connecting to Cloud Volumes ONTAP

If you need to perform advanced management of Cloud Volumes ONTAP, you can do so using OnCommand System Manager or the command line interface.

### Connecting to System Manager

You might need to perform some Cloud Volumes ONTAP tasks from System Manager, which is a browser-based management tool that runs on the Cloud Volumes ONTAP system. For example, you need to use System Manager if you want to create LUNs.

#### Before you begin

The computer from which you are accessing Cloud Manager must have a network connection to Cloud Volumes ONTAP. For example, you might need to log in to Cloud Manager from a jump host in AWS or Azure.



When deployed in multiple AWS Availability Zones, Cloud Volumes ONTAP HA configurations use a floating IP address for the cluster management interface, which means external routing is not available. You must connect from a host that is part of the same routing domain.

#### Steps

1. From the Working Environments page, double-click the Cloud Volumes ONTAP system that you want to manage with System Manager.
2. Click the menu icon, and then click **Advanced > System Manager**.
3. Click **Launch**.

System Manager loads in a new browser tab.

4. At the login screen, enter **admin** in the User Name field, enter the password that you specified when you created the working environment, and then click **Sign In**.

## Result

The System Manager console loads. You can now use it to manage Cloud Volumes ONTAP.

## Connecting to the Cloud Volumes ONTAP CLI

The Cloud Volumes ONTAP CLI enables you to execute all administrative commands and is a good choice for advanced tasks or if you are more comfortable using the CLI. You can connect to the CLI using Secure Shell (SSH).

### Before you begin

The host from which you use SSH to connect to Cloud Volumes ONTAP must have a network connection to Cloud Volumes ONTAP. For example, you might need to use SSH from a jump host in AWS or Azure.



When deployed in multiple AZs, Cloud Volumes ONTAP HA configurations use a floating IP address for the cluster management interface, which means external routing is not available. You must connect from a host that is part of the same routing domain.

### Steps

1. In Cloud Manager, identify the IP address of the cluster management interface:
  - a. On the Working Environments page, select the Cloud Volumes ONTAP system.
  - b. Copy the cluster management IP address that appears in the right pane.
2. Use SSH to connect to the cluster management interface IP address using the admin account.

### Example

The following image shows an example using PuTTY:



3. At the login prompt, enter the password for the admin account.

### Example

```
password: *****  
COT2:::>
```

## Adding existing Cloud Volumes ONTAP systems to Cloud Manager


You can discover and add existing Cloud Volumes ONTAP systems to Cloud Manager. You might do this if you deployed a new Cloud Manager system.

### Before you begin

You must know the password for the Cloud Volumes ONTAP admin user account.

## Steps

1. On the Working Environments page, click **Add Working Environment**.
2. Select the cloud provider in which the system resides.
3. Choose the type of Cloud Volumes ONTAP system.
4. Click the link to discover an existing system.



5. On the Region page, choose the region where the instances are running, and then select the instances.
6. On the Credentials page, enter the password for the Cloud Volumes ONTAP admin user, and then click **Go**.

## Result

Cloud Manager adds the Cloud Volumes ONTAP instances to the workspace.

## Deleting a Cloud Volumes ONTAP working environment

It is best to delete Cloud Volumes ONTAP systems from Cloud Manager, rather than from your cloud provider's console. For example, if you terminate a licensed Cloud Volumes ONTAP instance from AWS, then you can't use the license key for another instance. You must delete the working environment from Cloud Manager to release the license.

### About this task

When you delete a working environment, Cloud Manager terminates instances, deletes disks, and snapshots.



Cloud Volumes ONTAP instances have termination protection enabled to help prevent accidental termination from AWS. However, if you do terminate a Cloud Volumes ONTAP instance from AWS, you must go to the AWS CloudFormation console and delete the instance's stack. The stack name is the name of the working environment.

## Steps

1. From the working environment, click menu icon and then click **Delete**.
2. Type the name of the working environment and then click **Delete**.

It can take up to 5 minutes to delete the working environment.



# Provision volumes using a file service

## Azure NetApp Files

### Learn about Azure NetApp Files

Azure NetApp Files enables enterprises to migrate and run their performance-intensive and latency-sensitive core, business-critical applications in Azure with no need to refactor for the cloud.

#### Features

- Support for multiple protocols enables "lift & shift" of both Linux & Windows applications to run seamlessly in Azure.
- Multiple performance tiers allow for close alignment with workload performance requirements.
- Leading certifications including SAP HANA, GDPR, and HIPAA enables migration of the most demanding workloads to Azure.

#### Additional features in Cloud Manager

- Migrate NFS or SMB data to Azure NetApp Files directly from Cloud Manager. Data migrations are powered by NetApp's Cloud Sync service. [Learn more](#).
- Using Artificial Intelligence (AI) driven technology, Cloud Compliance can help you understand data context and identify sensitive data that resides in your Azure NetApp Files accounts. [Learn more](#).

#### Cost

[View Azure NetApp Files pricing](#).

Note that your subscription and charging are maintained by the Azure NetApp Files service and not by Cloud Manager.

#### Supported regions

[View supported Azure regions](#).

#### Requesting access

You need to be granted access to Azure NetApp Files by [submitting an online request](#). You'll need to wait for approval from the Azure NetApp Files team before you can proceed.

#### Getting help

For technical support issues associated with Azure NetApp Files, use the Azure portal to log a support request to Microsoft. Select your associated Microsoft subscription and select the **Azure NetApp Files** service name under **Storage**. Provide the remaining information required to create your Microsoft support request.

For issues related to Cloud Sync and Azure NetApp Files, you can start with NetApp using your Cloud Sync serial number directly from the Cloud Sync service. You will need to access the Cloud Sync service through the link in Cloud Manager. [View the process to enable Cloud Sync support](#).

## Related links

- [NetApp Cloud Central: Azure NetApp Files](#)
- [Azure NetApp Files documentation](#)
- [Cloud Sync documentation](#)

## Setting up Azure NetApp Files

Create an Azure NetApp Files working environment in Cloud Manager to create and manage NetApp accounts, capacity pools, volumes, and snapshots.

### Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



#### Request access

Submit an [online request](#) to be granted access to Azure NetApp Files.



#### Set up an Azure AD application

From Azure, grant permissions to an Azure AD application and copy the application (client) ID, the directory (tenant) ID, and the value of a client secret.



#### Create an Azure NetApp Files working environment

In Cloud Manager, click **Add Working Environment > Microsoft Azure > Azure NetApp Files** and then provide details about the AD application.

### Requesting access

You need to be granted access to Azure NetApp Files by [submitting an online request](#). You'll need to wait for approval from the Azure NetApp Files team before you can proceed.

### Setting up an Azure AD application

Cloud Manager needs permissions to set up and manage Azure NetApp Files. You can grant the required permissions to an Azure account by creating and setting up an Azure AD application and by obtaining the Azure credentials that Cloud Manager needs.

#### Creating the AD application

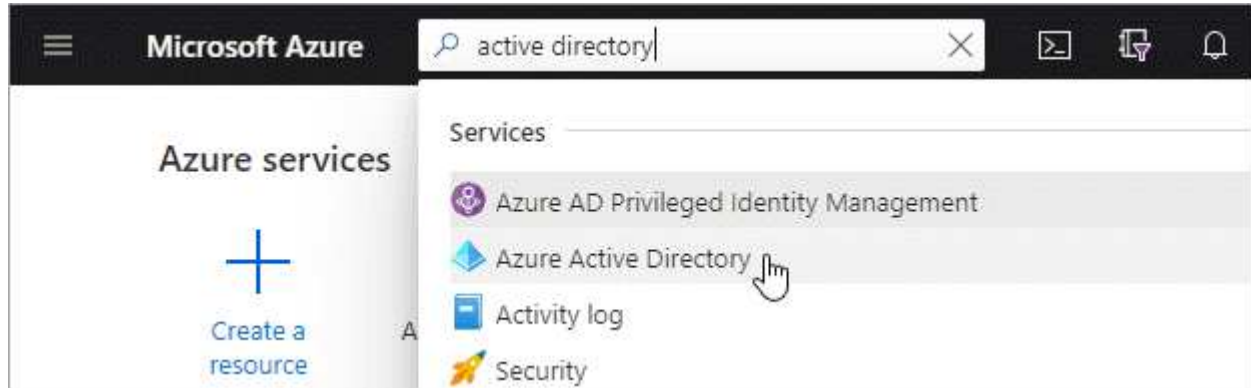
Create an Azure Active Directory (AD) application and service principal that Cloud Manager can use for role-based access control.

#### Before you begin

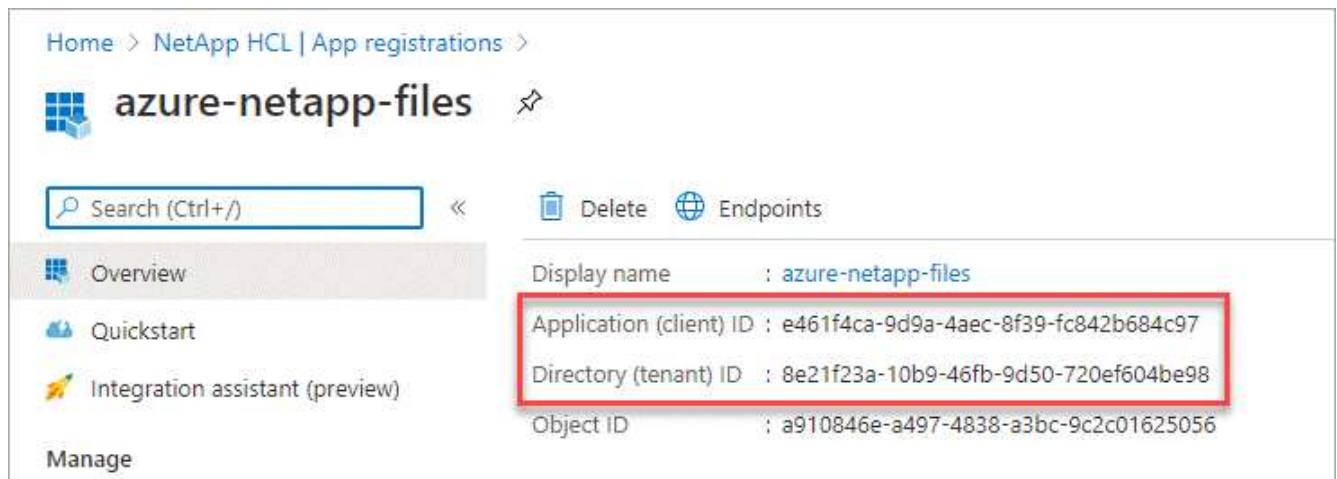
You must have the right permissions in Azure to create an Active Directory application and to assign the application to a role. For details, refer to [Microsoft Azure Documentation: Required permissions](#).

## Steps

1. From the Azure portal, open the **Azure Active Directory** service.



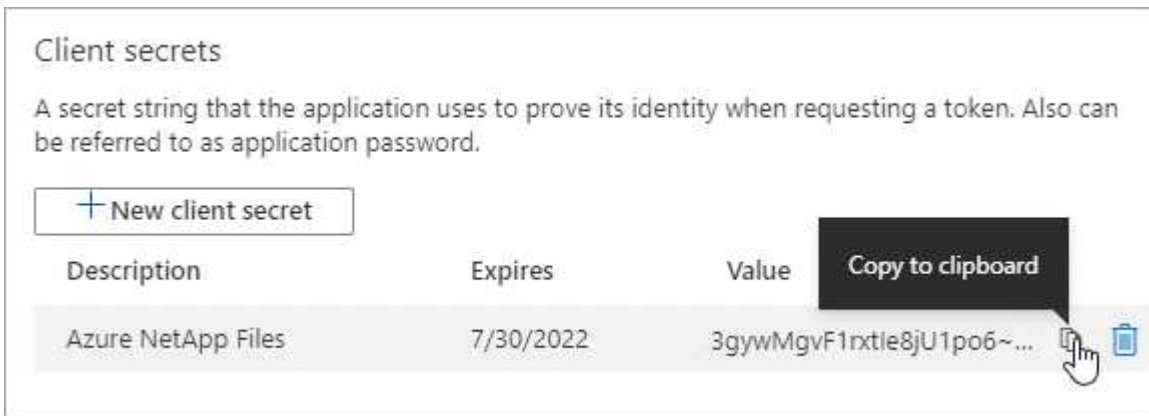
2. In the menu, click **App registrations**.
3. Create the application:
  - a. Click **New registration**.
  - b. Specify details about the application:
    - **Name**: Enter a name for the application.
    - **Account type**: Select an account type (any will work with Cloud Manager).
    - **Redirect URI**: You can leave this blank.
  - c. Click **Register**.
4. Copy the **Application (client) ID** and the **Directory (tenant) ID**.



When you create the Azure NetApp Files working environment in Cloud Manager, you need to provide the application (client) ID and the directory (tenant) ID for the application. Cloud Manager uses the IDs to programmatically sign in.

5. Create a client secret for the application so Cloud Manager can use it to authenticate with Azure AD:
  - a. Click **Certificates & secrets > New client secret**.
  - b. Provide a description of the secret and a duration.
  - c. Click **Add**.

d. Copy the value of the client secret.



### Result

Your AD application is now setup and you should have copied the application (client) ID, the directory (tenant) ID, and the value of the client secret. You need to enter this information in Cloud Manager when you add an Azure NetApp Files working environment.

### Assigning the app to a role


You must bind the service principal to your Azure subscription and assign it a custom role that has the required permissions.

### Steps

1. [Create a custom role in Azure.](#)

The following steps describe how to create the role from the Azure portal.

- a. Open the subscription and click **Access control (IAM)**.
- b. Click **Add > Add custom role**.



- c. In the **Basics** tab, enter a name and description for the role.
- d. Click **JSON** and click **Edit** which appears at the top right of the JSON format.
- e. Add the following permissions under *actions*:

```
"actions": [
  "Microsoft.NetApp/*",
  "Microsoft.Resources/resources/read",
  "Microsoft.Resources/subscriptions/resourceGroups/read",

  "Microsoft.Resources/subscriptions/resourcegroups/resources/read",
  "Microsoft.Resources/subscriptions/resourceGroups/write",
  "Microsoft.Network/virtualNetworks/read",
  "Microsoft.Insights/Metrics/Read"
],
```

- f. Click **Save**, click **Next**, and then click **Create**.
2. Now assign the application to the role that you just created:
    - a. From the Azure portal, open the subscription and click **Access control (IAM) > Add > Add role assignment**.
    - b. Select the custom role that you created.
    - c. Keep **Azure AD user, group, or service principal** selected.
    - d. Search for the name of the application (you can't find it in the list by scrolling).

**Add role assignment** [X]

Role ⓘ  
ANF 2.0 ⓘ

Assign access to ⓘ  
Azure AD user, group, or service principal

Select ⓘ  
azure-netapp-files

azure-netapp-files

e. Select the application and click **Save**.

The service principal for Cloud Manager now has the required Azure permissions for that subscription.

### Creating an Azure NetApp Files working environment

Set up an Azure NetApp Files working environment in Cloud Manager so you can start creating volumes.

1. From the Working Environments page, click **Add Working Environment**.
2. Select **Microsoft Azure** and then **Azure NetApp Files**.
3. Provide details about the AD application that you previously set up.

**Azure NetApp Files Credentials**

Working Environment Name

Application (client) ID

Client Secret

Directory (tenant) ID

4. Click **Add**.

**Result**

You should now have an Azure NetApp Files working environment.



**What's next?**

[Start creating and managing volumes.](#)

**Creating and managing volumes for Azure NetApp Files**

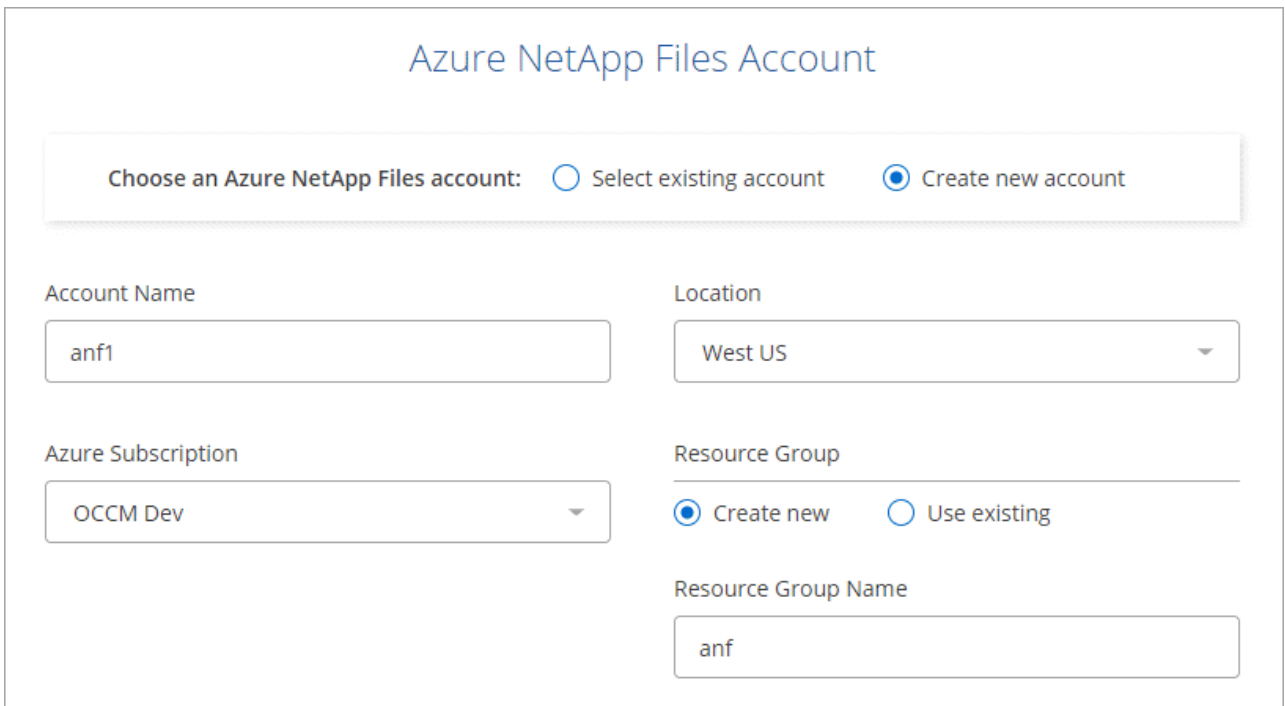
After you set up your working environment, you can create and manage Azure NetApp Files accounts, capacity pools, volumes, and snapshots.

## Creating volumes

You can create NFS or SMB volumes in a new or existing Azure NetApp Files account.

### Steps

1. Open the Azure NetApp Files working environment.
2. Click **Add New Volume**.
3. Provide the required information on each page:
  - **Azure NetApp Files Account:** Choose an existing Azure NetApp Files account or create a new account.



The screenshot shows the 'Azure NetApp Files Account' configuration page. At the top, the title 'Azure NetApp Files Account' is displayed. Below the title, there is a section for choosing an account type: 'Choose an Azure NetApp Files account:' with two radio buttons: 'Select existing account' (unselected) and 'Create new account' (selected). The form is divided into several sections: 'Account Name' with a text input field containing 'anf1'; 'Location' with a dropdown menu showing 'West US'; 'Azure Subscription' with a dropdown menu showing 'OCCM Dev'; 'Resource Group' with two radio buttons: 'Create new' (selected) and 'Use existing' (unselected); and 'Resource Group Name' with a text input field containing 'anf'.

- **Capacity Pool:** Select an existing capacity pool or create a new capacity pool.

If you create a new capacity pool, you need to specify a size and select a [service level](#).

The minimum size for the capacity pool is 4 TB. You can specify a size in multiples of 4 TB.

- **Details & Tags:** Enter a volume name and size, the VNet and subnet where the volume should reside, and optionally specify tags for the volume.
- **Protocol:** Choose the NFS or SMB protocol and enter the required information.

Here's an example of details for NFS.



Protocol

Select the volume's protocol:  NFS Protocol  SMB Protocol

Volume Path:

Select NFS Version:

NFSv3  NFSv4.1

Allowed Client & Access

Read & Write  Read Only

Read & Write  Read Only

Here's an example of details for SMB. You'll need to provide Active Directory information when you set up your first SMB volume.

Protocol

Select the volume's protocol:  NFS Protocol  SMB Protocol

Protocol:

Share Name:

Active Directory

Choose an Active Directory connection joined to your Azure NetApp Files account

Active Directory:


4. Click **Add Volume**.

## Mounting volumes

Access mounting instructions from within Cloud Manager so you can mount the volume to a host.

### Steps

1. Open the working environment.
2. Hover over the volume and select **Mount the volume**.



3. Follow the instructions to mount the volume.

### Editing a volume's size and tags

After you create a volume, you can modify its size and tags at any time.

#### Steps

1. Open the working environment.
2. Hover over the volume and select **Edit**.
3. Modify the size and tags as needed.
4. Click **Apply**.

### Managing Snapshot copies

Snapshot copies provide a point-in-time copy of your volume. Create Snapshot copies, restore the data to a new volume, and delete Snapshot copies.

#### Steps

1. Open the working environment.
2. Hover over the volume and choose one of the available options to manage Snapshot copies:
  - **Create a Snapshot copy**
  - **Restore to a new volume**
  - **Delete a Snapshot copy**
3. Follow the prompts to complete the selected action.

### Deleting volumes

Delete the volumes that you no longer need.

#### Steps

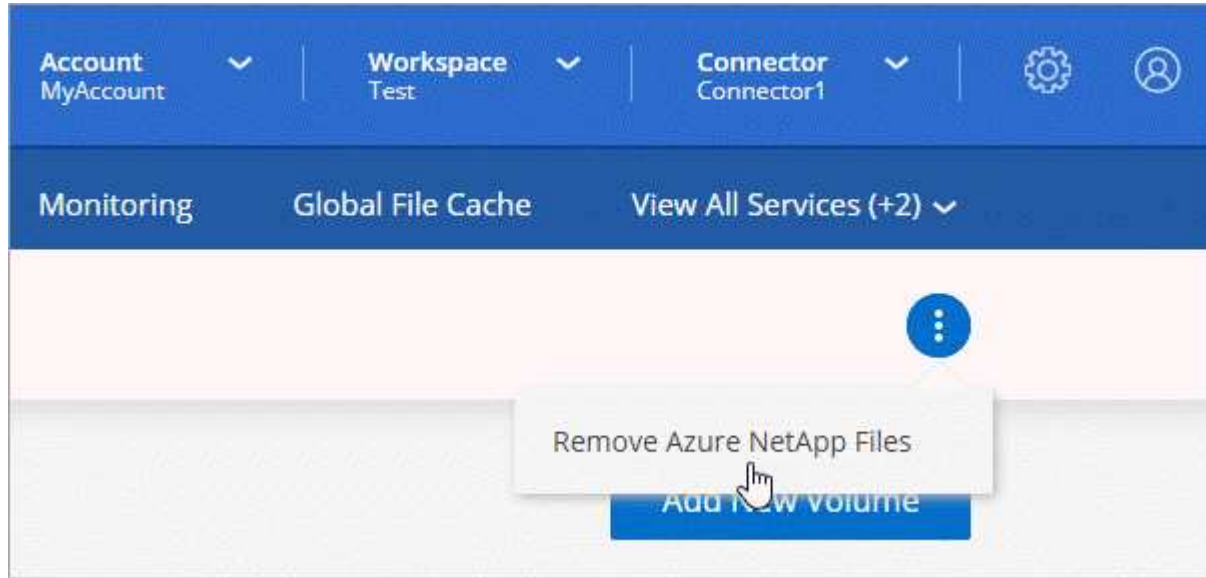
1. Open the working environment.
2. Hover over the volume and click **Delete**.
3. Confirm that you want to delete the volume.

## Removing Azure NetApp Files

This action removes Azure NetApp Files from Cloud Manager. It doesn't delete your Azure NetApp Files account or volumes. You can add Azure NetApp Files back to Cloud Manager at any time.

### Steps

1. Open the Azure NetApp Files working environment.
2. At the top right of the page, select the actions menu and click **Remove Azure NetApp Files**.



3. Click **Remove** to confirm.

## Cloud Volumes Service for AWS

### Learn about Cloud Volumes Service for AWS

NetApp Cloud Volumes Service for AWS is a cloud native file service that provides NAS volumes over NFS and SMB with all-flash performance. This service enables any workload, including legacy applications, to run in the AWS cloud.

### Benefits of using Cloud Volumes Service for AWS

Cloud Volumes Service for AWS provides the following benefits:

- Fully managed service, therefore no need to configure or manage storage devices
- Support for NFSv3 and NFSv4.1, and SMB 3.0 and 3.1.1 NAS protocols
- Secure access to Linux and Windows Elastic Container Service (ECS) instances, with support including the following:
  - Amazon Linux 2, Red Hat Enterprise Linux 7.5, SLES 12 SP3, and Ubuntu 16.04 LTS
  - Windows Server 2008 R2, Windows Server 2012 R2, and Windows Server 2016
- Choice of bundled and pay-as-you-go pricing

## Cost

Volumes created by the Cloud Volumes Service for AWS are charged based on your subscription to the service, not through Cloud Manager.

There are no charges to discover a Cloud Volumes Service for AWS region or volume from Cloud Manager.

## Before you get started

- Cloud Manager can discover existing Cloud Volumes Service for AWS subscriptions and volumes. See the [NetApp Cloud Volumes Service for AWS Account Setup Guide](#) if you haven't set up your subscription yet. You must follow this setup process for each region before you can add the AWS subscriptions and volumes in Cloud Manager.
- You need to obtain the Cloud Volumes API key and secret key so you can provide them to Cloud Manager. [For instructions, refer to Cloud Volumes Service for AWS documentation.](#)

## Quick start

Get started quickly by following these steps, or go to the next section for full details.



### Verify support for your configuration

You have set up AWS for Cloud Volumes Service and you must have subscribed to one of the [NetApp Cloud Volumes Service offerings on the AWS Marketplace](#).



### Add your Cloud Volumes Service for AWS subscription

You must create a working environment for the volumes based on your Cloud Volumes Service for AWS subscription.



### Create a cloud volumes

Cloud volumes that already exist for this subscription appear in the new working environment. Otherwise you create new volumes from Cloud Manager.



### Mount a cloud volume

Mount new cloud volumes to your AWS instance so that users can begin to use the storage.

## Getting help

Use the Cloud Manager chat for general service questions.

For technical support issues associated with your cloud volumes, use your 20 digit "930" serial number located in the "Support" tab of the Cloud Volumes Service user interface. Use this support ID when opening a web ticket or calling for support. Be sure to activate your Cloud Volumes Service serial number for support from the Cloud Volumes Service user interface. [Those steps are explained here.](#)

## Limitations

- Cloud Manager doesn't support data replication between working environments when using Cloud Volumes Service volumes.
- Removing your Cloud Volumes Service for AWS subscription from Cloud Manager isn't supported. You can do this only through the Cloud Volumes Service for AWS interface.

## Related links

- [NetApp Cloud Central: Cloud Volumes Service for AWS](#)
- [NetApp Cloud Volumes Service for AWS documentation](#)

## Managing Cloud Volumes Service for AWS

Cloud Manager enables you to create cloud volumes based on your [Cloud Volumes Service for AWS](#) subscription. You can also discover cloud volumes that you have already created from the Cloud Volumes Service interface and add them to a working environment.

### Add your Cloud Volumes Service for AWS subscription

Regardless of whether you have already created volumes from the Cloud Volumes Service user interface, or if you just signed up for Cloud Volumes Service for AWS and have no volumes yet, the first step is to create a working environment for the volumes based on your AWS subscription.

If cloud volumes already exist for this subscription, then the volumes are automatically added to the new working environment. If you haven't added any cloud volumes yet for the AWS subscription, then you do that after you create the new working environment.



If you have subscriptions and volumes in multiple AWS regions, you need to perform this task for each region.


### Before you begin

You must have the following information available when adding a subscription in each region:

- Cloud Volumes API key and Secret key: [See the Cloud Volumes Service for AWS documentation to get this information.](#)
- The AWS region where the subscription was created.

### Steps

1. In Cloud Manager, add a new Working Environment, select the location **Amazon Web Services**, and click **Continue**.
2. Select **Cloud Volumes Service** and click **Continue**.



3. Provide information about your Cloud Volumes Service subscription:
  - a. Enter the Working Environment Name you want to use.
  - b. Enter the Cloud Volumes Service API key and secret key.
  - c. Select the AWS region where your cloud volumes reside, or where they will be deployed.
  - d. Click **Add**.

### Cloud Volumes Service Credentials

Working Environment Name

Cloud Volumes Service API Key

Cloud Volumes Service Secret Key

AWS Region

## Result

Cloud Manager displays your Cloud Volumes Service for AWS configuration on the Working Environments page.



If cloud volumes already exist for this subscription, then the volumes are automatically added to the new working environment, as shown in the screenshot. You can add additional cloud volumes from Cloud Manager.

If no cloud volumes exist for this subscription, then you can create them now.

## Create cloud volumes

For configurations where volumes already exist in the Cloud Volumes Service working environment you can use these steps to add new volumes.

For configurations where no volumes exist, you can create your first volume directly from Cloud Manager after you have set up your Cloud Volumes Service for AWS subscription. In the past, the first volume had to be created directly in the Cloud Volumes Service user interface.

## Before you begin

- If you want to use SMB in AWS, you must have set up DNS and Active Directory.
- When planning to create an SMB volume, you must have a Windows Active Directory server available to which you can connect. You will enter this information when creating the volume. Also, make sure that the Admin user is able to create a machine account in the Organizational unit (OU) path specified.
- You will need this information when creating the first volume in a new region/working environment:
  - AWS account ID: A 12-digit Amazon account identifier with no dashes. To find your account ID, refer to this [AWS topic](#).
  - Classless Inter-Domain Routing (CIDR) Block: An unused IPv4 CIDR block. The network prefix must range between /16 and /28, and it must also fall within the ranges reserved for private networks (RFC 1918). Do not choose a network that overlaps your VPC CIDR allocations.

## Steps

1. Select the new working environment and click **Add New Volume**.
2. If you are adding the first volume to the working environment in the region, you have to add AWS networking information.
  - a. Enter the IPv4 range (CIDR) for the region.
  - b. Enter the 12-digit AWS account ID (with no dashes) to connect your Cloud Volumes account to your AWS account.
  - c. Click **Continue**.

**Network Setup**

! Your Cloud Volumes Service account isn't connected to your AWS account yet. Enter information about your AWS networking to connect the accounts. For details, see the [Cloud Volumes Service for AWS Account Setup document](#).

CIDR (IPv4) AWS Account ID

3. The Accepting Virtual Interfaces page describes some steps you will need to perform after you add the volume so that you are prepared to complete that step. Just click **Continue** again.
4. In the Details & Tags page, enter details about the volume:
  - a. Enter a name for the volume.
  - b. Specify a size within the range of 100 GiB to 90,000 GiB (equivalent to 88 TiBs).  
[Learn more about allocated capacity.](#)
  - c. Specify a service level: Standard, Premium, or Extreme.  
[Learn more about service levels.](#)
  - d. Enter one or more tag names to categorize the volume if you want.
  - e. Click **Continue**.
5. In the Protocol page, select NFS, SMB, or Dual Protocol and then define the details. Required entries for NFS and SMB are shown in separate sections below.
6. In the Volume Path field, specify the name of the volume export you will see when you mount the volume.
7. If you select Dual-protocol you can select the security style by selecting NTFS or UNIX. Security styles affect the file permission type used and how permissions can be modified.
  - UNIX uses NFSv3 mode bits, and only NFS clients can modify permissions.
  - NTFS uses NTFS ACLs, and only SMB clients can modify permissions.
8. For NFS:
  - a. In the NFS Version field, select NFSv3, NFSv4.1, or both depending on your requirements.
  - b. Optionally, you can create an export policy to identify the clients that can access the volume. Specify the:
    - Allowed clients by using an IP address or Classless Inter-Domain Routing (CIDR).
    - Access rights as Read & Write or Read Only.
    - Access protocol (or protocols if the volume allows both NFSv3 and NFSv4.1 access) used for users.
    - Click **+ Add Export Policy Rule** if you want to define additional export policy rules.

The following image shows the Volume page filled out for the NFS protocol:



**Protocol**

Select the volume's protocol:  NFS Protocol  SMB Protocol  Dual Protocol

**Volume Path** ⓘ

vol1

Select NFS Version:

NFSv3  NFSv4.1

**Export Policy**

**Allowed Client & Access** ⓘ

192.168.1.2/24  Read & Write  Read Only

Select NFS Version:  NFSv3  NFSv4.1

192.168.1.22/24  Read & Write  Read Only

Select NFS Version:  NFSv3  NFSv4.1

9. For SMB:

- a. You can enable SMB session encryption by checking the box for SMB Protocol Encryption.
- b. You can integrate the volume with an existing Windows Active Directory server by completing the fields in the Active directory section:

Field	Description
DNS Primary IP Address	The IP addresses of the DNS servers that provide name resolution for the SMB server. Use a comma to separate the IP addresses when referencing multiple servers, for example, 172.31.25.223, 172.31.2.74..
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the SMB server to join. When using AWS Managed Microsoft AD, use the value from the "Directory DNS name" field.
SMB Server NetBIOS name	A NetBIOS name for the SMB server that will be created.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
Organizational Unit	The organizational unit within the AD domain to associate with the SMB server. The default is CN=Computers for connections to your own Windows Active Directory server. If you configure AWS Managed Microsoft AD as the AD server for the Cloud Volumes Service, you should enter <b>OU=Computers,OU=corp</b> in this field.

The following image shows the Volume page filled out for the SMB protocol:

↑ SMB Connectivity Setup

DNS Primary IP Address	User Name
<input type="text" value="127.0.0.1"/>	<input type="text" value="administrator"/>
Active Directory Domain to Join	Password
<input type="text" value="yourdomain.com up to 107 characters"/>	<input type="password"/>
SMB Server NetBIOS Name	Organizational Unit
<input type="text" value="WEName"/>	<input type="text" value="CN=Computers"/>



You should follow the guidance on AWS security group settings to enable cloud volumes to integrate with Windows Active Directory servers correctly. See [AWS security group settings for Windows AD servers](#) for more information.

10. In the Volume from Snapshot page, if you want this volume to be created based on a snapshot of an existing volume, select the snapshot from the Snapshot Name drop-down list.
11. In the Snapshot Policy page, you can enable Cloud Volumes Service to create snapshot copies of your volumes based on a schedule. You can do this now or edit the volume later to define the snapshot policy.

See [Creating a snapshot policy](#) for more information about snapshot functionality.

12. Click **Add Volume**.

The new volume is added to the working environment.

### After you finish

If this is the first volume created in this AWS subscription, you need to launch the AWS Management Console to accept the two virtual interface that will be used in this AWS region to connect all your cloud volumes. See the [NetApp Cloud Volumes Service for AWS Account Setup Guide](#) for details.

You must accept the interfaces within 10 minutes after clicking the **Add Volume** button or the system may time out. If this happens, email [cvs-support@netapp.com](mailto:cvs-support@netapp.com) with your AWS Customer ID and NetApp Serial Number. Support will fix the issue and you can restart the onboarding process.

Then continue with [Mounting the cloud volume](#).

### Mount the cloud volume

You can mount a cloud volume to your AWS instance. Cloud volumes currently support NFSv3 and NFSv4.1 for Linux and UNIX clients, and SMB 3.0 and 3.1.1 for Windows clients.

**Note:** Please use the highlighted protocol/dialect supported by your client.

### Steps

1. Open the working environment.
2. Hover over the volume and click **Mount the volume**.

NFS and SMB volumes display mount instructions for that protocol. Dual-protocol volumes provide both sets of instructions.

3. Hover over the commands and copy them to your clipboard to make this process easier. Just add the destination directory/mount point at the end of the command.

### NFS example:

#### Mount the volume - testk

---

##### Setting up your instance

1. Open an SSH client and connect to your instance.
2. Install the nfs client on your instance.  
On Red Hat Enterprise Linux or SuSE Linux instance:  

```
$ sudo yum install -y nfs-utils
```

On an Ubuntu or Debian instance:  

```
$ sudo apt-get install nfs-common
```

##### Mounting your volume

1. Create a new directory on your instance:  

```
$ sudo mkdir /dir
```
2. Mount your NFSv3 volume using the command below:  

```
sudo mount -t nfs -o rw,hard,rsize=65536,wsiz=65536,vers=3,tc...
```
3. Mount your NFSv4.1 volume using the command below:  

```
sudo mount -t nfs -o rw,hard,rsize=65536,wsiz=65536,vers=4.1,t...
```

The maximum I/O size defined by the `rsize` and `wsiz` options is 1048576, however 65536 is the recommended default for most use cases.

Note that Linux clients will default to NFSv4.1 unless the version is specified with the `vers=<nfs_version>` option.

### SMB example:

### Mount the volume - <Volume Name>

**Mapping your network drive**

1. Click the **Start** button and then click on **Computer**.
2. Click **Map Network Drive**.
3. In the **Drive** list, click any available drive letter.
4. In the **Folder** box, type this:  

```
\\test.cv-pm.local\silly-condescending-mcnulty
```

To connect every time you log on to your computer, check the **Reconnect at logon** option.

5. Click **Finish**.

4. Connect to your Amazon Elastic Compute Cloud (EC2) instance by using an SSH or RDP client, and then follow the mount instructions for your instance.

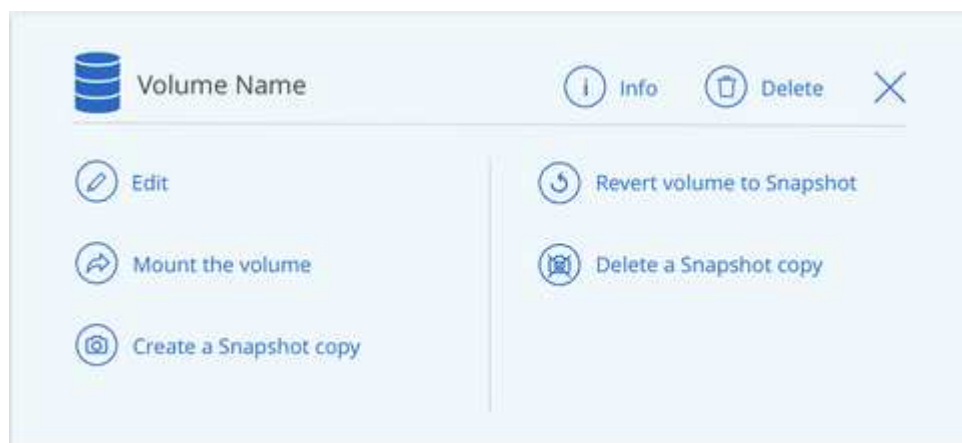
After completing the steps in the mount instructions, you have successfully mounted the cloud volume to your AWS instance.

## Managing existing volumes

You can manage existing volumes as your storage needs change. You can view, edit, restore, and delete volumes.

### Steps

1. Open the working environment.
2. Hover over the volume.



The screenshot shows a context menu for an AWS EBS volume. At the top, there is a volume icon, the text 'Volume Name', and three icons: 'Info', 'Delete', and a close 'X' button. Below this, the menu is divided into two columns. The left column contains three options: 'Edit' (with a pencil icon), 'Mount the volume' (with a hand icon), and 'Create a Snapshot copy' (with a camera icon). The right column contains two options: 'Revert volume to Snapshot' (with a circular arrow icon) and 'Delete a Snapshot copy' (with a trash can icon).

3. Manage your volumes:


Task	Action
View information about a volume	Select a volume, and then click <b>Info</b> .
Edit a volume (including snapshot policy)	<ol style="list-style-type: none"> <li>Select a volume, and then click <b>Edit</b>.</li> <li>Modify the volume's properties and then click <b>Update</b>.</li> </ol>
Get the NFS or SMB mount command	<ol style="list-style-type: none"> <li>Select a volume, and then click <b>Mount the volume</b>.</li> <li>Click <b>Copy</b> to copy the command(s).</li> </ol>
Create a Snapshot copy on demand	<ol style="list-style-type: none"> <li>Select a volume, and then click <b>Create a Snapshot copy</b>.</li> <li>Change the snapshot name, if needed, and then click <b>Create</b>.</li> </ol>
Replace the volume with the contents of a Snapshot copy	<ol style="list-style-type: none"> <li>Select a volume, and then click <b>Revert volume to Snapshot</b>.</li> <li>Select a Snapshot copy and click <b>Revert</b>.</li> </ol>
Delete a Snapshot copy	<ol style="list-style-type: none"> <li>Select a volume, and then click <b>Delete a Snapshot copy</b>.</li> <li>Select the Snapshot copy you want to delete and click <b>Delete</b>.</li> <li>Click <b>Delete</b> again to confirm.</li> </ol>
Delete a volume	<ol style="list-style-type: none"> <li>Unmount the volume from all clients: <ul style="list-style-type: none"> <li>On Linux clients, use the <code>umount</code> command.</li> <li>On Windows clients, click <b>Disconnect network drive</b>.</li> </ul> </li> <li>Select a volume, and then click <b>Delete</b>.</li> <li>Click <b>Delete</b> again to confirm.</li> </ol>


## Remove Cloud Volumes Service from Cloud Manager

You can remove a Cloud Volumes Service for AWS subscription and all existing volumes from Cloud Manager. The volumes are not deleted, they are just removed from the Cloud Manager interface.

### Steps

1. Open the working environment.





2. Click the  button at the top of the page and click **Remove Cloud Volumes Service**.
3. In the confirmation dialog box, click **Remove**.

## Manage Active Directory configuration

If you change your DNS servers or Active Directory domain, you need to modify the SMB server in Cloud Volumes Services so that it can continue to serve storage to clients.

You can also delete the link to an Active Directory if you no longer need it.

### Steps

1. Open the working environment.
2. Click the  button at the top of the page and click **Manage Active Directory**.
3. If no Active Directory is configured, you can add one now. If one is configured, you can modify the settings or delete it using the  button.
4. Specify the settings for the Active Directory that you want to join:

Field	Description
DNS Primary IP Address	The IP addresses of the DNS servers that provide name resolution for the SMB server. Use a comma to separate the IP addresses when referencing multiple servers, for example, 172.31.25.223, 172.31.2.74.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the SMB server to join. When using AWS Managed Microsoft AD, use the value from the "Directory DNS name" field.
SMB Server NetBIOS name	A NetBIOS name for the SMB server that will be created.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
Organizational Unit	The organizational unit within the AD domain to associate with the SMB server. The default is CN=Computers for connections to your own Windows Active Directory server. If you configure AWS Managed Microsoft AD as the AD server for the Cloud Volumes Service, you should enter <b>OU=Computers,OU=corp</b> in this field.

5. Click **Save** to save your settings.

## Manage cloud volumes snapshots

You can create a snapshot policy for each volume so that you can recover or restore the entire contents of a volume from an earlier time. You can also create an on-demand snapshot of a cloud volume when needed.

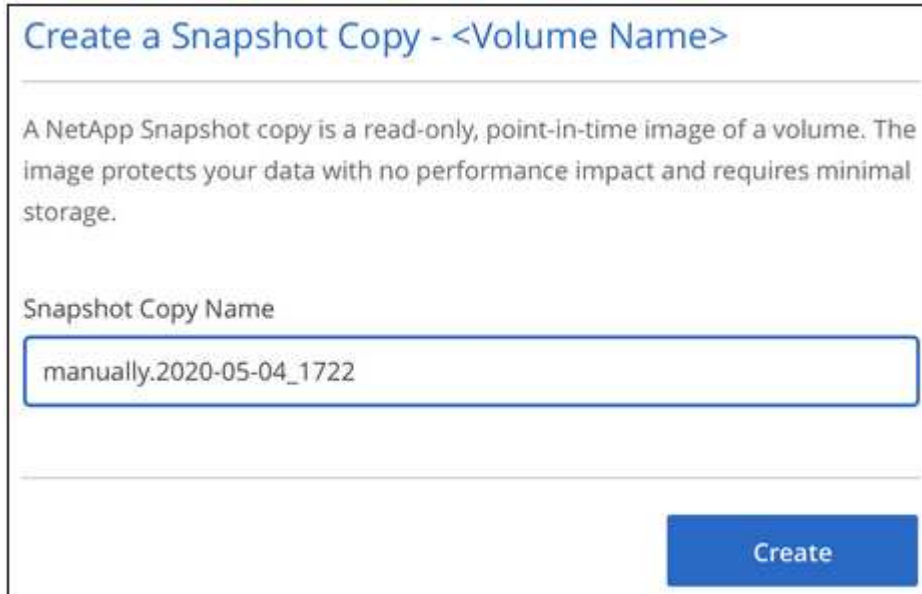
### Create an on-demand snapshot

You can create an on-demand snapshot of a cloud volume if you want to create a snapshot with the current

volume state.

### Steps

1. Open the working environment.
2. Hover over the volume and click **Create a snapshot copy**.
3. Enter a name for the snapshot, or use the automatically generated name, and click **Create**.



**Create a Snapshot Copy - <Volume Name>**

A NetApp Snapshot copy is a read-only, point-in-time image of a volume. The image protects your data with no performance impact and requires minimal storage.

Snapshot Copy Name

**Create**

### Create or modify a snapshot policy

You can create or modify a snapshot policy as necessary for a cloud volume. You define the snapshot policy from the *Snapshot Policy* tab either when creating a volume or when editing a volume.

### Steps

1. Open the working environment.
2. Hover over the volume and click **Edit**.
3. From the *Snapshot Policy* tab, move the enable snapshots slider to the right.
4. Define the schedule for snapshots:
  - a. Select the frequency: **Hourly**, **Daily**, **Weekly**, or **Monthly**
  - b. Select the number of snapshots you want to keep.
  - c. Select the day, hour, and minute when the snapshot should be taken.

**Schedule Snapshot Policies:**

<input checked="" type="checkbox"/> Hourly	Number of Snapshot to Keep	Minute		
	<input type="text" value="12"/>	<input type="text" value="30"/>		
<input type="checkbox"/> Daily	Number of Snapshot to Keep	Hour	Minute	
	<input type="text" value="0"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	
<input checked="" type="checkbox"/> Weekly	Number of Snapshot to Keep	Days	Hour	Minute
	<input type="text" value="3"/>	<input type="text" value="Sunday x"/>	<input type="text" value="0"/>	<input type="text" value="0"/>
		<input type="checkbox"/> Sunday		
		<input type="checkbox"/> Monday		
		<input type="checkbox"/> Tuesday		
<input type="checkbox"/> Monthly	Number of Snapshot to Keep		Hour	Minute
	<input type="text" value="0"/>		<input type="text" value="0"/>	<input type="text" value="0"/>


5. Click **Add volume** or **Update volume** to save your policy settings.

### Disable a snapshot policy

You can disable a snapshot policy to stop snapshots from being created for a short period of time while retaining your snapshot policy settings.

#### Steps

1. Open the working environment.
2. Hover over the volume and click **Edit**.
3. From the *Snapshot Policy* tab, move the enable snapshots slider to the left.



4. Click **Update volume**.

When you want to re-enable the snapshot policy, move the enable snapshots slider to the right and click **Update volume**.

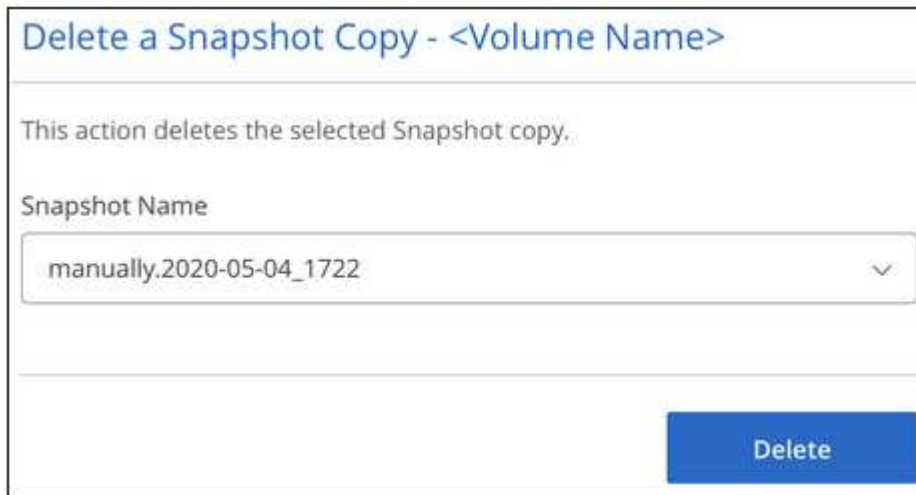
### Delete a snapshot

You can delete a snapshot from the Volumes page.



## Steps

1. Open the working environment.
2. Hover over the volume and click **Delete a Snapshot copy**.
3. Select the snapshot from the drop-down list and click **Delete**.



Delete a Snapshot Copy - <Volume Name>

This action deletes the selected Snapshot copy.

Snapshot Name

manually.2020-05-04\_1722

Delete

4. In the confirmation dialog box, click **Delete**.

## Revert a volume from a snapshot


You can revert a volume to an earlier point in time from an existing snapshot.

When you revert a volume, the content of the snapshot overwrites the existing volume configuration. Any changes that were made to the data in the volume after the snapshot was created are lost.

Note that clients do not need to remount the volume after the revert operation.

## Steps

1. Open the working environment.
2. Hover over the volume and click **Revert volume to Snapshot**.
3. Select the snapshot that you want to use to restore the existing volume from the drop-down list and click **Revert**.



## Reference

### Service levels and allocated capacity

The cost for Cloud Volumes Service for AWS is based on the *service level* and the *allocated capacity* that you select. Selecting the appropriate service level and capacity helps you meet your storage needs at the lowest cost.

### Considerations

Storage needs include two fundamental aspects:

- The storage *capacity* for holding data
- The storage *bandwidth* for interacting with data

If you consume more storage space than the capacity you selected for the volume, the following considerations apply:

- You will be billed for the additional storage capacity that you consume at the price defined by your service level.
- The amount of storage bandwidth available to the volume does not increase until you increase the allocated capacity size or change the service level.

### Service levels

Cloud Volumes Service for AWS supports three service levels. You specify your service level when you create or modify the volume.

The service levels are catered to different storage capacity and storage bandwidth needs:

- **Standard** (capacity)

If you want capacity at the lowest cost, and your bandwidth needs are limited, then the Standard service level might be most appropriate for you. An example is using the volume as a backup target.

- Bandwidth: 16 KB of bandwidth per GB provisioned capacity

- **Premium** (a balance of capacity and performance)

If your application has a balanced need for storage capacity and bandwidth, then the Premium service level might be most appropriate for you. This level is less expensive per MB/s than the Standard service level, and it is also less expensive per GB of storage capacity than the Extreme service level.

- Bandwidth: 64 KB of bandwidth per GB provisioned capacity

- **Extreme** (performance)

The Extreme service level is least expensive in terms of storage bandwidth. If your application demands storage bandwidth without the associated demand for lots of storage capacity, then the Extreme service level might be most appropriate for you.

- Bandwidth: 128 KB of bandwidth per GB provisioned capacity

### **Allocated capacity**

You specify your allocated capacity for the volume when you create or modify the volume.

While you would select your service level based on your general, high-level business needs, you should select your allocated capacity size based on the specific needs of applications, for example:

- How much storage space the applications need
- How much storage bandwidth per second the applications or the users require

Allocated capacity is specified in GBs. A volume's allocated capacity can be set within the range of 100 GB to 100,000 GB (equivalent to 100 TBs).

### **Number of inodes**

Volumes less than or equal to 1 TB can use up to 20 million inodes. The number of inodes increase by 20 million for each TB you allocate, up to a maximum of 100 million inodes.

- <= 1TB = 20 million inodes
- >1 TB to 2 TB = 40 million inodes
- >2 TB to 3 TB = 60 million inodes
- >3 TB to 4 TB = 80 million inodes
- >4 TB to 100 TB = 100 million inodes

### **Bandwidth**

The combination of both the service level and the allocated capacity you select determines the maximum bandwidth for the volume.

If your applications or users need more bandwidth than your selections, you can change the service level or increase the allocated capacity. The changes do not disrupt data access.

### **Selecting the service level and the allocated capacity**

To select the most appropriate service level and allocated capacity for your needs, you need to know how much capacity and bandwidth you require at the peak or the edge.

## List of service levels and allocated capacity

The leftmost column indicates the capacity, and the other columns define the MB/s available at each capacity point based on service level.

See [Contract subscription pricing](#) and [Metered subscription pricing](#) for complete details on pricing.

Capacity (TB)	Standard (MB/s)	Premium (MB/s)	Extreme (MB/s)
0.1 (100 GB)	1.6	6.4	12.8
1	16	64	128
2	32	128	256
3	48	192	384
4	64	256	512
5	80	320	640
6	96	384	768
7	112	448	896
8	128	512	1,024
9	144	576	1,152
10	160	640	1,280
11	176	704	1,408
12	192	768	1,536
13	208	832	1,664
14	224	896	1,792
15	240	960	1,920
16	256	1,024	2,048
17	272	1,088	2,176
18	288	1,152	2,304
19	304	1,216	2,432
20	320	1,280	2,560
21	336	1,344	2,688
22	352	1,408	2,816
23	368	1,472	2,944
24	384	1,536	3,072
25	400	1,600	3,200
26	416	1,664	3,328
27	432	1,728	3,456
28	448	1,792	3,584

<b>Capacity (TB)</b>	<b>Standard (MB/s)</b>	<b>Premium (MB/s)</b>	<b>Extreme (MB/s)</b>
29	464	1,856	3,712
30	480	1,920	3,840
31	496	1,984	3,968
32	512	2,048	4,096
33	528	2,112	4,224
34	544	2,176	4,352
35	560	2,240	4,480
36	576	2,304	4,500
37	592	2,368	4,500
38	608	2,432	4,500
39	624	2,496	4,500
40	640	2,560	4,500
41	656	2,624	4,500
42	672	2,688	4,500
43	688	2,752	4,500
44	704	2,816	4,500
45	720	2,880	4,500
46	736	2,944	4,500
47	752	3,008	4,500
48	768	3,072	4,500
49	784	3,136	4,500
50	800	3,200	4,500
51	816	3,264	4,500
52	832	3,328	4,500
53	848	3,392	4,500
54	864	3,456	4,500
55	880	3,520	4,500
56	896	3,584	4,500
57	912	3,648	4,500
58	928	3,712	4,500
59	944	3,776	4,500
60	960	3,840	4,500
61	976	3,904	4,500

<b>Capacity (TB)</b>	<b>Standard (MB/s)</b>	<b>Premium (MB/s)</b>	<b>Extreme (MB/s)</b>
62	992	3,968	4,500
63	1,008	4,032	4,500
64	1,024	4,096	4,500
65	1,040	4,160	4,500
66	1,056	4,224	4,500
67	1,072	4,288	4,500
68	1,088	4,352	4,500
69	1,104	4,416	4,500
70	1,120	4,480	4,500
71	1,136	4,500	4,500
72	1,152	4,500	4,500
73	1,168	4,500	4,500
74	1,184	4,500	4,500
75	1,200	4,500	4,500
76	1,216	4,500	4,500
77	1,232	4,500	4,500
78	1,248	4,500	4,500
79	1,264	4,500	4,500
80	1,280	4,500	4,500
81	1,296	4,500	4,500
82	1,312	4,500	4,500
83	1,328	4,500	4,500
84	1,344	4,500	4,500
85	1,360	4,500	4,500
86	1,376	4,500	4,500
87	1,392	4,500	4,500
88	1,408	4,500	4,500
89	1,424	4,500	4,500
90	1,440	4,500	4,500
91	1,456	4,500	4,500
92	1,472	4,500	4,500
93	1,488	4,500	4,500
94	1,504	4,500	4,500

Capacity (TB)	Standard (MB/s)	Premium (MB/s)	Extreme (MB/s)
95	1,520	4,500	4,500
96	1,536	4,500	4,500
97	1,552	4,500	4,500
98	1,568	4,500	4,500
99	1,584	4,500	4,500
100	1,600	4,500	4,500

### Example 1

For example, your application requires 25 TB of capacity and 100 MB/s of bandwidth. At 25 TB of capacity, the Standard service level would provide 400 MB/s of bandwidth at a cost of \$2,500 (estimate: see current pricing), making Standard the most suitable service level in this case.

### Example 2

For example, your application requires 12 TB of capacity and 800 MB/s of peak bandwidth. Although the Extreme service level can meet the demands of the application at the 12 TB mark, it is more cost-effective (estimate: see current pricing) to select 13 TB at the Premium service level.

### AWS security group settings for Windows AD servers

If you use Windows Active Directory (AD) servers with cloud volumes, you should familiarize yourself with the guidance on AWS security group settings. The settings enable cloud volumes to integrate with AD correctly.

By default, the AWS security group applied to an EC2 Windows instance does not contain inbound rules for any protocol except RDP. You must add rules to the security groups that are attached to each Windows AD instance to enable inbound communication from Cloud Volumes Service. The required ports are as follows:

Service	Port	Protocol
AD Web Services	9389	TCP
DNS	53	TCP
DNS	53	UDP
ICMPv4	N/A	Echo Reply
Kerberos	464	TCP
Kerberos	464	UDP
Kerberos	88	TCP
Kerberos	88	UDP
LDAP	389	TCP

Service	Port	Protocol
LDAP	389	UDP
LDAP	3268	TCP
NetBIOS name	138	UDP
SAM/LSA	445	TCP
SAM/LSA	445	UDP
Secure LDAP	636	TCP
Secure LDAP	3269	TCP
w32time	123	UDP

If you are deploying and managing your AD installation domain controllers and member servers on an AWS EC2 instance, you will require several security group rules to allow traffic for the Cloud Volumes Service. Below is an example of how to implement these rules for AD applications as part of the AWS CloudFormation template.

```
{
  "AWSTemplateFormatVersion" : "2010-09-09",
  "Description" : "Security Group for AD",
  "Parameters" :
  {
    "VPC" :
    {
      "Type" : "AWS::EC2::VPC::Id",
      "Description" : "VPC where the Security Group will belong:"
    },
    "Name" :
    {
      "Type" : "String",
      "Description" : "Name Tag of the Security Group:"
    },
    "Description" :
    {
      "Type" : "String",
      "Description" : "Description Tag of the Security Group:",
      "Default" : "Security Group for Active Directory for CVS "
    },
    "CIDRrangeforTCPandUDP" :
    {
      "Type" : "String",
      "Description" : "CIDR Range for the UDP ports
445,138,464,389,53,123 and for the TCP ports
464,339,3389,3268,88,636,9389,445 and 0-65535: *CIDR range format:
10.0.0.0/24"
    }
  }
}
```



```

    }
  },
  "Resources" :
  {
    "ADSGWest" :
    {
      "Type" : "AWS::EC2::SecurityGroup",
      "Properties" :
      {
        "GroupDescription" : {"Ref" : "Description"},
        "VpcId" : { "Ref" : "VPC" },
        "SecurityGroupIngress" : [
          {
            "IpProtocol" : "udp",
            "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
            "FromPort" : "445",
            "ToPort" : "445"
          },
          {
            "IpProtocol" : "udp",
            "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
            "FromPort" : "138",
            "ToPort" : "138"
          },
          {
            "IpProtocol" : "udp",
            "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
            "FromPort" : "464",
            "ToPort" : "464"
          },
          {
            "IpProtocol" : "tcp",
            "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
            "FromPort" : "464",
            "ToPort" : "464"
          },
          {
            "IpProtocol" : "udp",
            "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
            "FromPort" : "389",
            "ToPort" : "389"
          },
          {
            "IpProtocol" : "udp",
            "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
            "FromPort" : "53",

```

```

        "ToPort" : "53"
    },
    {
        "IpProtocol" : "tcp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "339",
        "ToPort" : "339"
    },
    {
        "IpProtocol" : "udp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "123",
        "ToPort" : "123"
    },
    {
        "IpProtocol" : "tcp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "3389",
        "ToPort" : "3389"
    },
    {
        "IpProtocol" : "tcp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "3268",
        "ToPort" : "3268"
    },
    {
        "IpProtocol" : "tcp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "88",
        "ToPort" : "88"
    },
    {
        "IpProtocol" : "tcp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "636",
        "ToPort" : "636"
    },
    {
        "IpProtocol" : "tcp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "3269",
        "ToPort" : "3269"
    },
    {
        "IpProtocol" : "tcp",

```

```

        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "53",
        "ToPort" : "53"
    },
    {
        "IpProtocol" : "tcp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "0",
        "ToPort" : "65535"
    },
    {
        "IpProtocol" : "tcp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "9389",
        "ToPort" : "9389"
    },
    {
        "IpProtocol" : "tcp",
        "CidrIp" : {"Ref" : "CIDRrangeforTCPandUDP"},
        "FromPort" : "445",
        "ToPort" : "445"
    }
]
}
}
},
"Outputs" :
{
    "SecurityGroupID" :
    {
        "Description" : "Security Group ID",
        "Value" : { "Ref" : "ADSGWest" }
    }
}
}

```

## Cloud Volumes Service for GCP

### Learn about Cloud Volumes Service for Google Cloud

NetApp Cloud Volumes Service for Google Cloud enables you to quickly add multi-protocol workloads as well as build and deploy both Windows-based and UNIX-based apps.

## Key features:

- Migrate data between on-premises and Google Cloud.
- Provision volumes from 1 to 100 TiB in seconds.
- Multiprotocol support (you can create an NFS or SMB volume).
- Protect data with automated, efficient snapshots.
- Accelerate app development with rapid cloning.

## Cost

Volumes created by the Cloud Volumes Service for Google Cloud are charged to your subscription to the service, not through Cloud Manager.

[View pricing](#)

There are no charges to discover a Cloud Volumes Service for Google Cloud region or volume from Cloud Manager.

## Supported regions

[View supported Google Cloud regions.](#)

## Before you get started

Cloud Manager can discover existing Cloud Volumes Service for GCP subscriptions and volumes. See the [NetApp Cloud Volumes Service for Google Cloud documentation](#) if you haven't set up your subscription yet.

## Getting help

Use the Cloud Manager chat for general questions about Cloud Volumes Service operation in Cloud Manager.

For general questions about Cloud Volumes Service for Google Cloud, email NetApp's Google Cloud Team at [gcinfo@netapp.com](mailto:gcinfo@netapp.com).

For technical issues associated with your cloud volumes, you can create a technical support case from the Google Cloud Console. See [obtaining support](#) for details.

## Limitations

- Cloud Manager doesn't support data replication between working environments when using Cloud Volumes Service volumes.
- Deleting your Cloud Volumes Service for Google Cloud subscription from Cloud Manager isn't supported. You can do this only through the Google Cloud Console.

## Related links

- [NetApp Cloud Central: Cloud Volumes Service for Google Cloud](#)
- [NetApp Cloud Volumes Service for Google Cloud documentation](#)

## Set up Cloud Volumes Service for Google Cloud

Create a Cloud Volumes Service for Google Cloud working environment in Cloud Manager to create and manage volumes and snapshots.

### Quick start

Get started quickly by following these steps, or go to the next section for full details.



#### Enable the Cloud Volumes Service API

From Google, enable the Cloud Volumes Service for GCP API so that Cloud Manager can manage the subscription and cloud volumes.



#### Create a GCP service account and download credentials

From Google, create a GCP service account and role so that Cloud Manager can access your Cloud Volumes Service for GCP account.



#### Create a Cloud Volumes Service for GCP working environment

In Cloud Manager, click **Add Working Environment > Google Cloud > Cloud Volumes Service** and then provide details about the service account and Google Cloud project.

### Enable the Cloud Volumes Service API

In Google Cloud Shell, run the following command to enable the Cloud Volumes Service API:

```
gcloud --project=<my-cvs-project> services enable cloudvolumesgcp-api.netapp.com
```

### Give Cloud Manager access to the Cloud Volumes Service for GCP account

You must complete the following tasks so that Cloud Manager can access your Google Cloud project:

- Create a new service account
- Add the new service account member to your project and assign it specific roles (permissions)
- Create and download a key pair for the service account that is used to authenticate to Google

### Steps

1. In the Google Cloud Console, go to the **Service Accounts** page.
2. Click **Select a project**, choose your project, and click **Open**.
3. Click **Create Service Account**, enter the service account name (friendly display name) and description, and click **Create**.
4. From the *IAM page* click **Add** and fill out the fields in the *Add Members* page:
  - a. In the New Members field, enter the full service account ID, for example, [user1-service-account-cvs@project1.iam.gserviceaccount.com](#).

b. Add these roles:

- *NetApp Cloud Volumes Admin*
- *Compute Network Viewer*
- *Folder Viewer*

c. Click **Save**.

5. From the *Service account details* page click **Add key > Create new key**.

6. Select **JSON** as the key type and click **Create**.

By clicking **Create** your new public/private key pair is generated and downloaded to your system. It serves as the only copy of the private key. Store this file securely because it can be used to authenticate as your service account.

For detailed steps, see the Google Cloud topics [Creating and managing service accounts](#), [Granting, changing, and revoking access to resources](#), and [Creating and managing service account keys](#).

## Create a Cloud Volumes Service for GCP working environment

Set up a Cloud Volumes Service for GCP working environment in Cloud Manager so you can start creating volumes.

Regardless of whether you have already created volumes from the Google Cloud Console, or if you just signed up for Cloud Volumes Service for GCP and have no volumes yet, the first step is to create a working environment for the volumes based on your GCP subscription.

If cloud volumes already exist for this subscription, then the volumes will appear in the new working environment. If you haven't added any cloud volumes yet for the GCP subscription, then you do that after you create the new working environment.



If you have subscriptions and volumes in multiple GCP projects, you need to perform this task for each project.

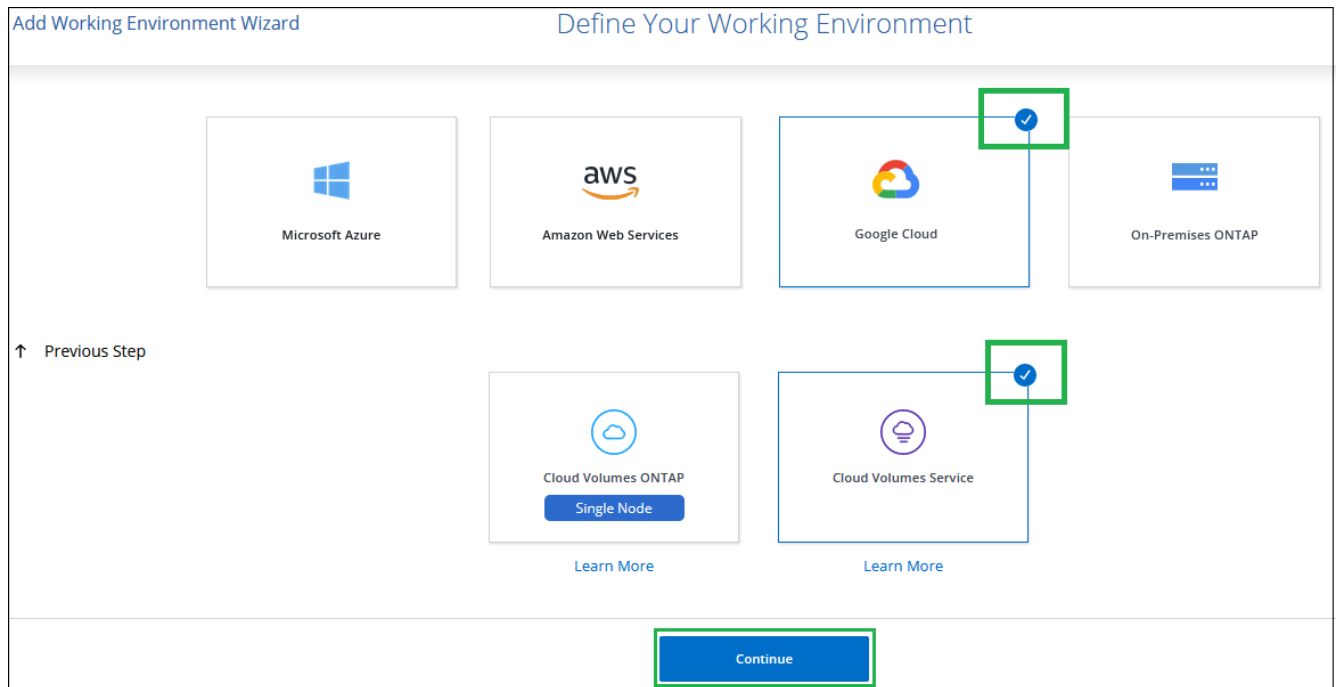
### Before you begin

You must have the following information available when adding a subscription for each project:

- Service account credentials (JSON private key you downloaded)
- Project name

### Steps

1. In Cloud Manager, add a new Working Environment, select the location **Google Cloud**, and click **Continue**.
2. Select **Cloud Volumes Service** and click **Continue**.



3. Provide information about your Cloud Volumes Service subscription:
  - a. Enter the Working Environment Name you want to use.
  - b. Copy/paste the JSON private key you downloaded in the previous steps.
  - c. Select the name of your Google Cloud project.
  - d. Click **Add**.

Cloud Volumes Service Credentials

Working Environment Name

Service Account Credentials

Paste the contents of the JSON file here:

**Apply**

Project

- Select project -

## Result

Cloud Manager displays your Cloud Volumes Service for Google Cloud working environment.



If cloud volumes already exist for this subscription, then the volumes appear in the new working environment, as shown in the screenshot. You can add additional cloud volumes from Cloud Manager.

If no cloud volumes exist for this subscription, create them now.

### What's next?

[Start creating and managing volumes.](#)

## Create and manage volumes for Cloud Volumes Service for Google Cloud

Cloud Manager enables you to create cloud volumes based on your [Cloud Volumes Service for Google Cloud](#) subscription. You can also edit certain attributes of a volume, get the relevant mount commands, create snapshot copies, and delete cloud volumes.

### Create cloud volumes

You can create NFS or SMB volumes in a new or existing Cloud Volumes Service for Google Cloud account. Cloud volumes currently support NFSv3 and NFSv4.1 for Linux and UNIX clients, and SMB 3.x for Windows clients.

### Before you begin

- If you want to use SMB in GCP, you must have set up DNS and Active Directory.
- When planning to create an SMB volume, you must have a Windows Active Directory server available to which you can connect. You will enter this information when creating the volume. Also, make sure that the Admin user is able to create a machine account in the Organizational unit (OU) path specified.

### Steps

1. Select the working environment and click **Add New Volume**.
2. In the Details & Location page, enter details about the volume:
  - a. Enter a name for the volume.
  - b. Specify a size within the range of 1 TiB (1024 GiB) to 100 TiB.  
[Learn more about allocated capacity.](#)
  - c. Specify a service level: Standard, Premium, or Extreme.  
[Learn more about service levels.](#)
  - d. Select the Google Cloud region.
  - e. Select the VPC Network from which the volume will be accessible. Note that the VPC cannot be changed or edited after the volume is created.
  - f. Click **Continue**.
3. In the Protocol page, select NFS or SMB and then define the details. Required entries for NFS and SMB are shown in separate sections below.
4. For NFS:
  - a. In the Volume Path field, specify the name of the volume export you will see when you mount the volume.



- b. Select NFSv3, NFSv4.1, or both depending on your requirements.
- c. Optionally, you can create an export policy to identify the clients that can access the volume. Specify the:
  - Allowed clients by using an IP address or Classless Inter-Domain Routing (CIDR).
  - Access rights as Read & Write or Read Only.
  - Access protocol (or protocols if the volume allows both NFSv3 and NFSv4.1 access) used for users.
  - Click **+ Add Export Policy Rule** if you want to define additional export policy rules.

The following image shows the Volume page filled out for the NFS protocol:

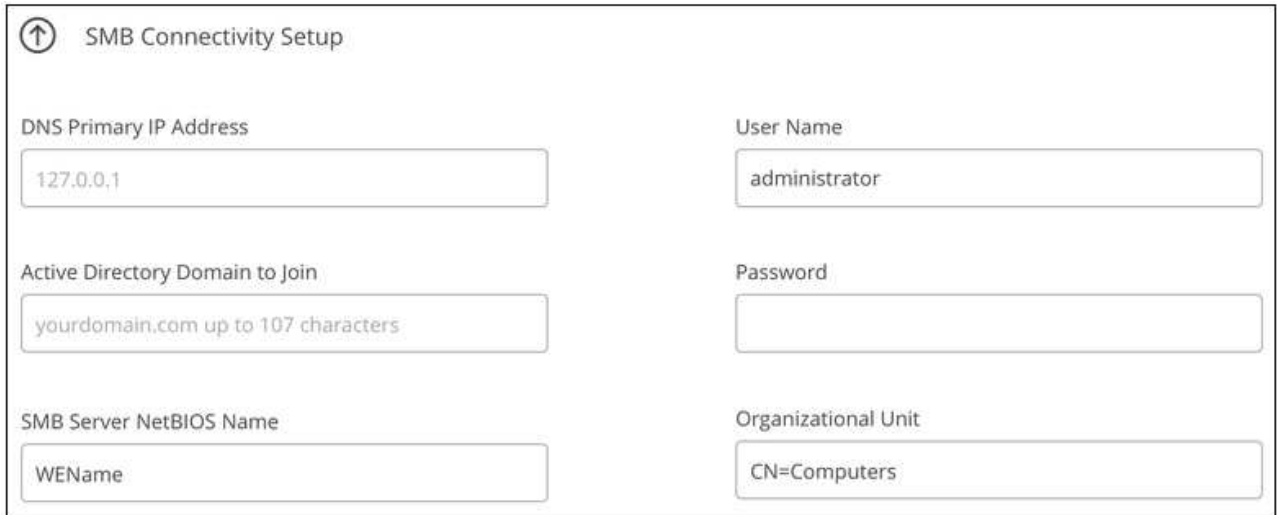
#### 5. For SMB:

- a. In the Volume Path field, specify the name of the volume export you will see when you mount the volume and click **Continue**.
- b. If Active Directory has been set up, you see the configuration. If it is the first volume being set up and no Active Directory has been set up, you can enable SMB session encryption in the SMB Connectivity Setup page:

Field	Description
DNS Primary IP Address	The IP addresses of the DNS servers that provide name resolution for the SMB server. Use a comma to separate the IP addresses when referencing multiple servers, for example, 172.31.25.223, 172.31.2.74..
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the SMB server to join.
SMB Server NetBIOS name	A NetBIOS name for the SMB server that will be created.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.

Field	Description
Organizational Unit	The organizational unit within the AD domain to associate with the SMB server. The default is CN=Computers for connections to your own Windows Active Directory server.

The following image shows the Volume page filled out for the SMB protocol:



SMB Connectivity Setup

DNS Primary IP Address: 127.0.0.1

User Name: administrator

Active Directory Domain to Join: yourdomain.com up to 107 characters

Password:

SMB Server NetBIOS Name: WEName

Organizational Unit: CN=Computers

6. Click **Continue**.
7. If you want to create the volume based on a snapshot of an existing volume, select the snapshot from the Snapshot Name drop-down list. Otherwise just click **Continue**.
8. In the Snapshot Policy page, you can enable Cloud Volumes Service to create snapshot copies of your volumes based on a schedule. You can do this now by moving the selector to the right, or you can edit the volume later to define the snapshot policy.

See [Creating a snapshot policy](#) for more information about snapshot functionality.

9. Click **Add Volume**.

The new volume is added to the working environment.

Continue with [Mounting the cloud volume](#).

## Mount cloud volumes

Access mounting instructions from within Cloud Manager so you can mount the volume to a host.

**Note:** Please use the highlighted protocol/dialect supported by your client.

### Steps

1. Open the working environment.
2. Hover over the volume and click **Mount the volume**.

NFS and SMB volumes display mount instructions for that protocol.

3. Hover over the commands and copy them to your clipboard to make this process easier. Just add the destination directory/mount point at the end of the command.

## NFS example:

### Mount the volume - testk

#### Setting up your instance

1. Open an SSH client and connect to your instance.
2. Install the nfs client on your instance.  
On Red Hat Enterprise Linux or SuSE Linux instance:

```
$ sudo yum install -y nfs-utils
```

On an Ubuntu or Debian instance:

```
$ sudo apt-get install nfs-common
```

#### Mounting your volume

1. Create a new directory on your instance:

```
$ sudo mkdir /dir
```
2. Mount your NFSv3 volume using the command below:

```
sudo mount -t nfs -o rw,hard,rsize=65536,wsiz=65536,vers=3,tc...
```
3. Mount your NFSv4.1 volume using the command below:

```
sudo mount -t nfs -o rw,hard,rsize=65536,wsiz=65536,vers=4,1,t...
```

The maximum I/O size defined by the `rsize` and `wsiz` options is 1048576, however 65536 is the recommended default for most use cases.

Note that Linux clients will default to NFSv4.1 unless the version is specified with the `vers=<nfs_version>` option.

## SMB example:

### Mount the volume - <Volume Name>

#### Mapping your network drive

1. Click the **Start** button and then click on **Computer**.
2. Click **Map Network Drive**.
3. In the **Drive** list, click any available drive letter.
4. In the **Folder** box, type this:  

```
\\test.cv-pm.local\silly-condescending-mcnulty
```

To connect every time you log on to your computer, check the **Reconnect at logon** option.

5. Click **Finish**.

4. Map your network drive by following the mount instructions for your instance.


After completing the steps in the mount instructions, you have successfully mounted the cloud volume to your GCP instance.

## Manage existing volumes

You can manage existing volumes as your storage needs change. You can view, edit, restore, and delete volumes.

### Steps

1. Open the working environment.
2. Hover over the volume.




3. Manage your volumes:

Task	Action
View information about a volume	Click <b>Info</b> .
Edit a volume (including snapshot policy)	<ol style="list-style-type: none"> <li>Click <b>Edit</b>.</li> <li>Modify the volume's properties and then click <b>Update</b>.</li> </ol>
Get the NFS or SMB mount command	<ol style="list-style-type: none"> <li>Click <b>Mount the volume</b>.</li> <li>Click <b>Copy</b> to copy the command(s).</li> </ol>
Create a Snapshot copy on demand	<ol style="list-style-type: none"> <li>Click <b>Create a Snapshot copy</b>.</li> <li>Change the name, if needed, and then click <b>Create</b>.</li> </ol>
Replace the volume with the contents of a Snapshot copy	<ol style="list-style-type: none"> <li>Click <b>Revert volume to snapshot</b>.</li> <li>Select a Snapshot copy and click <b>Restore</b>.</li> </ol>
Delete a Snapshot copy	<ol style="list-style-type: none"> <li>Click <b>Delete a Snapshot copy</b>.</li> <li>Select the snapshot and click <b>Delete</b>.</li> <li>Click <b>Delete</b> again when prompted to confirm.</li> </ol>
Delete a volume	<ol style="list-style-type: none"> <li>Unmount the volume from all clients: <ul style="list-style-type: none"> <li>◦ On Linux clients, use the <code>umount</code> command.</li> <li>◦ On Windows clients, click <b>Disconnect network drive</b>.</li> </ul> </li> <li>Select a volume, and then click <b>Delete</b>.</li> <li>Click <b>Delete</b> again to confirm.</li> </ol>

## Remove Cloud Volumes Service from Cloud Manager

You can remove a Cloud Volumes Service for Google Cloud subscription and all existing volumes from Cloud Manager. The volumes are not deleted, they are just removed from the Cloud Manager interface.

### Steps



1. Open the working environment.
2. Click the  button at the top of the page and click **Remove Cloud Volumes Service**.
3. In the confirmation dialog box, click **Remove**.

## Manage Active Directory configuration

If you change your DNS servers or Active Directory domain, you need to modify the SMB server in Cloud Volumes Services so that it can continue to serve storage to clients.

### Steps

1. Open the working environment.

2. Click the  button at the top of the page and click **Manage Active Directory**. If no Active Directory is configured, you can add one now. If one is configured, you can modify or delete the settings using the  button.
3. Specify the settings for the SMB server:

Field	Description
DNS Primary IP Address	The IP addresses of the DNS servers that provide name resolution for the SMB server. Use a comma to separate the IP addresses when referencing multiple servers, for example, 172.31.25.223, 172.31.2.74.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the SMB server to join.
SMB Server NetBIOS name	A NetBIOS name for the SMB server that will be created.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
Organizational Unit	The organizational unit within the AD domain to associate with the SMB server. The default is CN=Computers for connections to your own Windows Active Directory server.

4. Click **Save** to save your settings.

## Manage cloud volumes snapshots

You can create a snapshot policy for each volume so that you can recover or restore the entire contents of a volume from an earlier time. You can also create an on-demand snapshot of a cloud volume when needed.

### Create an on-demand snapshot

You can create an on-demand snapshot of a cloud volume if you want to create a snapshot with the current volume state.

#### Steps

1. Open the working environment.
2. Hover over the volume and click **Create a snapshot copy**.
3. Enter a name for the snapshot, or use the automatically generated name, and click **Create**.

### Create a Snapshot Copy - <Volume Name>

A NetApp Snapshot copy is a read-only, point-in-time image of a volume. The image protects your data with no performance impact and requires minimal storage.

Snapshot Copy Name

**Create**

The snapshot is created.

### Create or modify a snapshot policy

You can create or modify a snapshot policy as necessary for a cloud volume. You define the snapshot policy from the *Snapshot Policy* tab either when creating a volume or when editing a volume.

#### Steps

1. Open the working environment.
2. Hover over the volume and click **Edit**.
3. From the *Snapshot Policy* tab, move the enable snapshots slider to the right.
4. Define the schedule for snapshots:
  - a. Select the frequency: **Hourly**, **Daily**, **Weekly**, or **Monthly**
  - b. Select the number of snapshots you want to keep.
  - c. Select the day, hour, and minute when the snapshot should be taken.

**Schedule Snapshot Policies:**

<input checked="" type="checkbox"/> Hourly	Number of Snapshot to Keep	Minute		
	<input type="text" value="12"/>	<input type="text" value="30"/>		
<input type="checkbox"/> Daily	Number of Snapshot to Keep	Hour	Minute	
	<input type="text" value="0"/>	<input type="text" value="0"/>	<input type="text" value="0"/>	
<input checked="" type="checkbox"/> Weekly	Number of Snapshot to Keep	Days	Hour	Minute
	<input type="text" value="3"/>	<input type="text" value="Sunday x"/>	<input type="text" value="0"/>	<input type="text" value="0"/>
		<input type="checkbox"/> Sunday		
		<input type="checkbox"/> Monday		
		<input type="checkbox"/> Tuesday		
<input type="checkbox"/> Monthly	Number of Snapshot to Keep		Hour	Minute
	<input type="text" value="0"/>		<input type="text" value="0"/>	<input type="text" value="0"/>


5. Click **Add volume** or **Update volume** to save your policy settings.

### Disable a snapshot policy

You can disable a snapshot policy to stop snapshots from being created for a short period of time while retaining your snapshot policy settings.

#### Steps

1. Open the working environment.
2. Hover over the volume and click **Edit**.
3. From the *Snapshot Policy* tab, move the enable snapshots slider to the left.



4. Click **Update volume**.

When you want to re-enable the snapshot policy, move the enable snapshots slider to the right and click **Update volume**.

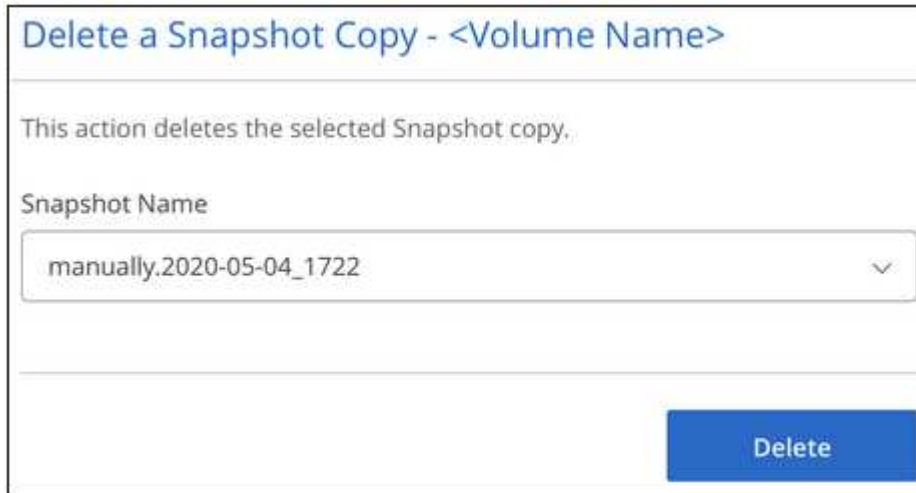
### Delete a snapshot

You can delete a snapshot if it is no longer needed.



### Steps

1. Open the working environment.
2. Hover over the volume and click **Delete a Snapshot copy**.
3. Select the snapshot from the drop-down list and click **Delete**.



Delete a Snapshot Copy - <Volume Name>

This action deletes the selected Snapshot copy.

Snapshot Name

manually.2020-05-04\_1722

Delete

4. In the confirmation dialog box, click **Delete**.

### Restore a snapshot to a new volume

You can restore a snapshot to a new volume as necessary.

### Steps

1. Open the working environment.
2. Hover over the volume and click **Restore to a new volume**.
3. Select the snapshot that you want to use to create the new volume from the drop-down list.
4. Enter a name for the new volume and click **Restore**.

### Restore to a new volume - <Volume Name>

This operation restores data from a Snapshot copy to a new volume.

Snapshot Name

manually.2020-05-04\_1722

Restored Volume Name:

vol\_restore

Restore

The volume is created in the working environment.

5. If you need to change any of the volume attributes, such as volume path or service level:
  - a. Hover over the volume and click **Edit**.
  - b. Make your changes and click **Update volume**.

#### **After you finish**

Continue with [Mounting the cloud volume](#).

# Manage ONTAP clusters

## Discovering ONTAP clusters

Cloud Manager can discover the ONTAP clusters in your on-premises environment, in a NetApp Private Storage configuration, and in the IBM Cloud. Discovering an ONTAP cluster enables you to provision storage, replicate data, back up data, and tier cold data from an on-prem cluster to the cloud.

### What you'll need

- A Connector installed in a cloud provider or on your premises.

If you want to tier cold data to the cloud, then you should review requirements for the Connector based on where you plan to tier cold data.

- [Learn about Connectors](#)
- [Switching between Connectors](#)
- [Learn about Cloud Tiering](#)
- The cluster management IP address and the password for the admin user account to add the cluster to Cloud Manager.

Cloud Manager discovers ONTAP clusters using HTTPS. If you use custom firewall policies, they must meet the following requirements:

- The Connector host must allow outbound HTTPS access through port 443.

If the Connector is in the cloud, all outbound communication is allowed by the predefined security group.

- The ONTAP cluster must allow inbound HTTPS access through port 443.

The default "mgmt" firewall policy allows inbound HTTPS access from all IP addresses. If you modified this default policy, or if you created your own firewall policy, you must associate the HTTPS protocol with that policy and enable access from the Connector host.

### Steps

1. On the Working Environments page, click **Add Working Environment** and select **On-Premises ONTAP**.
2. If you're prompted, create a Connector.

Refer to the links above for more details.

3. On the **ONTAP Cluster Details** page, enter the cluster management IP address, the password for the admin user account, and the location of the cluster.

## ONTAP Cluster Details

Provide a few details about your ONTAP cluster so Cloud Manager can discover it.

Cluster Management IP Address

User Name

Password

Add

4. On the Details page, enter a name and description for the working environment, and then click **Go**.

### Result

Cloud Manager discovers the cluster. You can now create volumes, replicate data to and from the cluster, set up data tiering to the cloud, back up volumes to the cloud, and launch System Manager to perform advanced tasks.

## Managing storage for ONTAP clusters

After you discover your ONTAP cluster from Cloud Manager, you can open the working environment to provision and manage storage.

### Creating volumes for ONTAP clusters

Cloud Manager enables you to provision NFS, CIFS, and iSCSI volumes on ONTAP clusters.

#### Before you begin

The data protocols must be set up on the cluster using System Manager or the CLI.

#### About this task

You can create volumes on existing aggregates. You can't create new aggregates from Cloud Manager.

#### Steps

1. On the Working Environments page, double-click the name of the ONTAP cluster on which you want to provision volumes.
2. Click **Add New Volume**.
3. On the Create New Volume page, enter details for the volume, and then click **Create**.

Some of the fields in this page are self-explanatory. The following table describes fields for which you might need guidance:

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, Cloud Manager enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bust adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, Cloud Manager automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, select it, click Target IQN, and then use the IQN to connect to the LUN from your hosts.</p>
Usage Profile	Usage profiles define the NetApp storage efficiency features that are enabled for a volume.

## Replicating data

You can replicate data between Cloud Volumes ONTAP systems and ONTAP clusters by choosing a one-time data replication, which can help you move data to and from the cloud, or a recurring schedule, which can help with disaster recovery or long-term retention.

[Click here for more details.](#)

## **Backing up data**

You can back up data from your on-premises ONTAP system to low-cost object storage in the cloud by using the Cloud Manager Backup to Cloud service. This service provides backup and restore capabilities for protection and long-term archive of your cloud data.

[Click here for more details.](#)

## **Tiering data to the cloud**

Extend your data center to the cloud by automatically tiering inactive data from ONTAP clusters to object storage.

[Click here for more details.](#)

# Back up to the cloud

## Learn about Backup to Cloud

Backup to Cloud is an add-on service for Cloud Volumes ONTAP and on-premises ONTAP clusters that delivers backup and restore capabilities for protection, and long-term archive of your cloud data. Backups are stored in an object store in your cloud account, independent of volume Snapshot copies used for near-term recovery or cloning.

Backup to Cloud is powered by the [Cloud Backup Service](#).



You must use Cloud Manager for all backup and restore operations. Any actions taken directly from ONTAP or from your cloud provider results in an unsupported configuration.

### Features

- Back up independent copies of your data volumes to low-cost object storage in the cloud.
- Backup data is secured with AES-256 bit encryption at-rest and TLS 1.2 HTTPS connections in-flight.
- Back up from cloud to cloud, and from on-premises ONTAP systems to cloud.
- Support for up to 1,019 backups of a single volume.
- Restore data from a specific point in time.
- Restore the data to a volume on the source system or to a different system.

### Supported working environments and object storage providers

Backup to Cloud is supported with the following types of working environments:

- Cloud Volumes ONTAP in AWS
- Cloud Volumes ONTAP in Azure
- On-premises ONTAP clusters

### Cost

Backup to Cloud is available in two pricing options: Bring Your Own License (BYOL) and Pay As You Go (PAYGO).

For BYOL you'll pay NetApp to use the service for a period of time, say 6 months, and for a maximum amount backup capacity, say 10 GB (before storage efficiencies), and you'll need to pay your cloud provider for object storage costs. You'll receive a serial number that you enter in the Cloud Manager Licensing page to enable the service. When either limit is reached you'll need to renew the license. See [Adding and updating your Backup BYOL license](#). The Backup BYOL license applies to all Cloud Volumes ONTAP systems associated with your [Cloud Central account](#).

For PAYGO you'll need to pay your cloud provider for object storage costs and NetApp for backup licensing costs. The licensing costs are based on used capacity (before storage efficiencies):

- AWS: [Go to the Cloud Manager Marketplace offering for pricing details](#).

- Azure: [Go to the Cloud Manager Marketplace offering for pricing details.](#)

## Free trial

A 30-day free trial is available. When using the trial version, you are notified about the number of free trial days that remain. At the end of your free trial, backups stop being created. You must subscribe to the service or purchase a license to continue using the service.

Backups are not deleted when the service is disabled. You'll continue to be charged by your cloud provider for object storage costs for the capacity that your backups use unless you delete the backups.

## How Backup to Cloud works

When you enable Backup to Cloud on a Cloud Volumes ONTAP or on-premises ONTAP system, the service performs a full backup of your data. Volume snapshots are not included in the backup image. After the initial backup, all additional backups are incremental, which means that only changed blocks and new blocks are backed up.

## Where backups reside

Backup copies are stored in an S3 bucket or Azure Blob container that Cloud Manager creates in your cloud account. For Cloud Volumes ONTAP systems the object store is created in the same region where the Cloud Volumes ONTAP system is located. For on-premises ONTAP systems you identify the region when you enable the service.

There's one object store per Cloud Volumes ONTAP or on-premises ONTAP system. Cloud Manager names the object store as follows: `netapp-backup-clusteruid`

Be sure not to delete this object store.

Notes:

- In AWS, Cloud Manager enables the [Amazon S3 Block Public Access feature](#) on the S3 bucket.
- In Azure, Cloud Manager uses a new or existing resource group with a storage account for the Blob container.

## Supported S3 storage classes

In Amazon S3, backups start in the *Standard* storage class and transition to the *Standard-Infrequent Access* storage class after 30 days.

## Supported Azure Blob access tiers

In Azure, each backup is associated with the *cold* access tier.

## Backup settings are system wide

When you enable Backup to Cloud, all the volumes you identify on the system are backed up to the cloud.

The schedule and number of backups to retain are defined at the system level. The backup settings affect all volumes on the system.



## The schedule is daily, weekly, monthly, or a combination

You can choose daily, or weekly, or monthly backups of all volumes. You can also select one of the system-defined policies that provide backups and retention for 3 months, 1 year, and 7 years. These policies are:

Policy Name	Backups per interval...			Max. Backups
	Daily	Weekly	Monthly	
Netapp3MonthsRetention	30	13	3	46
Netapp1YearRetention	30	13	12	55
Netapp7YearsRetention	30	53	84	167

Once you have reached the maximum number of backups for a category, or interval, older backups are removed so you always have the most current backups.

Note that the retention period for backups of data protection volumes is the same as defined in the source SnapMirror relationship. You can change this if you want by using the API.

### Backups are taken at midnight

- Daily backups start just after midnight each day.
- Weekly backups start just after midnight on Sunday mornings.
- Monthly backups start just after midnight on the first of each month.

At this time, you can't schedule backup operations at a user specified time.

### Backup copies are associated with your Cloud Central account

Backup copies are associated with the [Cloud Central account](#) in which Cloud Manager resides.

If you have multiple Cloud Manager systems in the same Cloud Central account, each Cloud Manager system will display the same list of backups. That includes the backups associated with Cloud Volumes ONTAP and on-premises ONTAP instances from other Cloud Manager systems.

### BYOL license considerations

When using a Backup to Cloud BYOL license, Cloud Manager notifies you when backups are nearing the capacity limit or nearing the license expiration date. You receive these notifications:

- when backups have reached 80% of licensed capacity, and again when you have reached the limit
- 30 days before a license is due to expire, and again when the license expires

Use the chat icon in the lower right of the Cloud Manager interface to renew your license when you receive these notifications.

Two things can happen when your license expires:

- If the account you are using for your ONTAP systems has a marketplace account, the backup service continues to run, but you are shifted over to a PAYGO licensing model. You are charged by your cloud

provider for object storage costs, and by NetApp for backup licensing costs, for the capacity that your backups are using.

- If the account you are using for your ONTAP systems does not have a marketplace account, the backup service continues to run, but you will continue to receive the expiration message.

Once you renew your BYOL subscription, Cloud Manager automatically obtains the new license from NetApp and installs it. If Cloud Manager can't access the license file over the secure internet connection, you can obtain the file yourself and manually upload it to Cloud Manager. For instructions, see [Adding and updating your Backup BYOL license](#).

Systems that were shifted over to a PAYGO license are returned to the BYOL license automatically. And systems that were running without a license will stop receiving the warning message and will be charged for backups that occurred while the license was expired.

## Supported volumes

Backup to Cloud supports read-write volumes and data protection (DP) volumes.

FlexGroup volumes aren't currently supported.

## Limitations

- WORM storage (SnapLock) is not supported on a Cloud Volumes ONTAP or on-premises system when Backup to Cloud is enabled.
- Backup to Cloud restrictions when making backups from on-premises ONTAP systems:
  - The on-prem cluster must be running ONTAP 9.7P5 or later.
  - Cloud Manager must be deployed on Azure. There is no support for on-premises Cloud Manager deployments.
  - The destination location for backups is only object storage on Azure.
  - Backups can be restored only to Cloud Volumes ONTAP systems deployed on Azure. You cannot restore a backup to an on-premises ONTAP system or to a Cloud Volumes ONTAP system that is using a different cloud provider.
- When backing up data protection (DP) volumes, the rule that is defined for the SnapMirror policy on the source volume must use a label that matches the allowed Backup to Cloud policy names of **daily**, **weekly**, or **monthly**. Otherwise the backup will fail for that DP volume.
- In Azure, if you enable Backup to Cloud when Cloud Volumes ONTAP is deployed, Cloud Manager creates the resource group for you and you cannot change it. If you want to pick your own resource group when enabling Backup to Cloud, **disable** Backup to Cloud when deploying Cloud Volumes ONTAP and then enable Backup to Cloud and choose the resource group from the Backup to Cloud Settings page.
- When backing up volumes from Cloud Volumes ONTAP systems, volumes that you create outside of Cloud Manager aren't automatically backed up.

For example, if you create a volume from the ONTAP CLI, ONTAP API, or System Manager, then the volume won't be automatically backed up.

If you want to back up these volumes, you would need to disable Backup to Cloud and then enable it again.

# Get started

## Backing up data to Amazon S3

Complete a few steps to get started backing up data from Cloud Volumes ONTAP to Amazon S3.

### Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.

1


### Verify support for your configuration

- You're running Cloud Volumes ONTAP 9.6 or later in AWS.
- You have subscribed to the [Cloud Manager Marketplace Backup offering](#), or you have purchased and [activated](#) a Backup to Cloud BYOL license from NetApp.
- The IAM role that provides Cloud Manager with permissions includes S3 permissions from the latest [Cloud Manager policy](#).

2

### Enable Backup to Cloud on your new or existing system

- New systems: Backup to Cloud is enabled by default in the working environment wizard. Be sure to keep the option enabled.
- Existing systems: Select the working environment and click **Activate** next to the Backup to Cloud service in the right-panel, and then follow the setup wizard.



3

### Define the backup policy

The default policy backs up volumes every day and retains the most recent 30 backup copies of each volume. Change to weekly or monthly backups, or select one of the system-defined policies that provide more options. You can also change the number of backup copies to retain.

### Define Policy

**Policy - Retention & Schedule**

Create a New Policy   
  Select an Existing Policy

Backup Every:    
 Number of backups to retain:

---

**DP Volumes**

Data protection volume backups use the same retention period as defined in the source SnapMirror relationship by default. Use the API if you want to change this value

---

**Information**

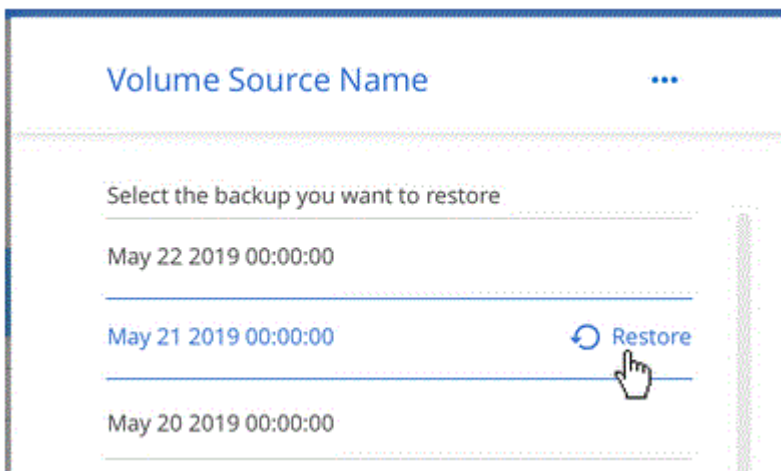
Backup\_Bucket\_Name  
Bucket Name

#### 4 Select the volumes that you want to back up

Identify which volumes you want to back up in the Select Volumes page.

#### 5 Restore your data, as needed

From the Backup List, select a volume, select a backup, and then restore data from the backup to a new volume.



#### Requirements

Read the following requirements to make sure that you have a supported configuration before you start backing up volumes to S3.

#### Supported ONTAP versions

Cloud Volumes ONTAP 9.6 and later.

## Supported AWS regions

Backup to Cloud is supported in all AWS regions [where Cloud Volumes ONTAP is supported](#).

## License requirements

For Backup to Cloud PAYGO licensing, a Cloud Manager subscription is available in the AWS Marketplace that enables deployments of Cloud Volumes ONTAP 9.6 and later (PAYGO) and Backup to Cloud. You need to [subscribe to this Cloud Manager subscription](#) before you enable Backup to Cloud. Billing for Backup to Cloud is done through this subscription.

For Backup to Cloud BYOL licensing, you do not need an AWS Backup to Cloud subscription. You need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. See [Adding and updating your Backup BYOL license](#).

And you need to have a AWS subscription for the storage space where your backups will be located.

## AWS permissions required

The IAM role that provides Cloud Manager with permissions must include S3 permissions from the latest [Cloud Manager policy](#).

Here are the specific permissions from the policy:

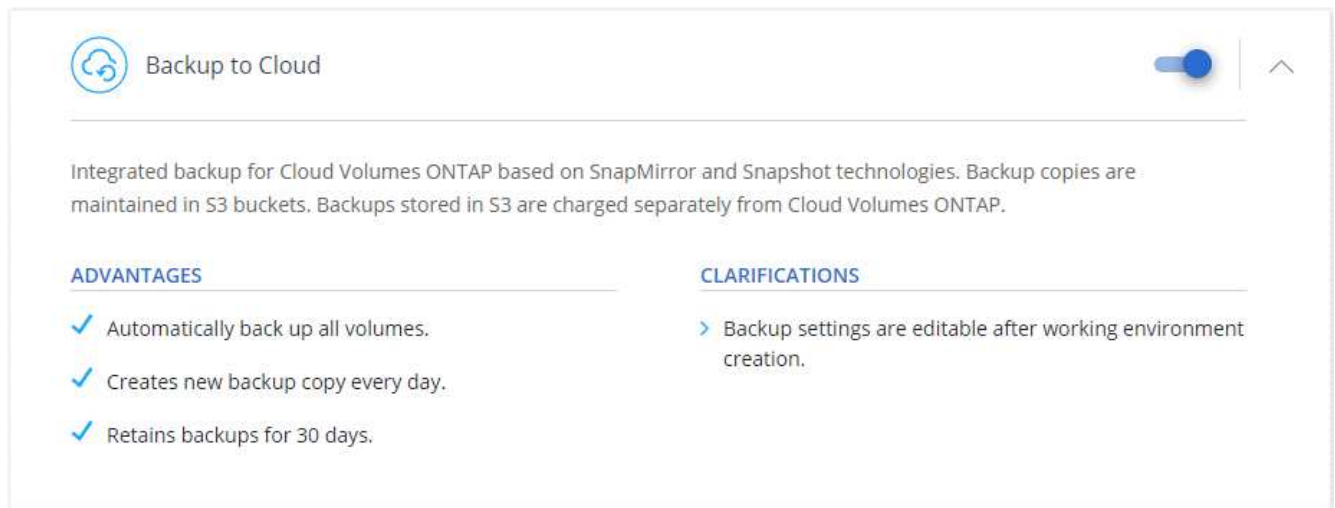
```
{
    "Sid": "backupPolicy",
    "Effect": "Allow",
    "Action": [
        "s3:DeleteBucket",
        "s3:GetLifecycleConfiguration",
        "s3:PutLifecycleConfiguration",
        "s3:PutBucketTagging",
        "s3:ListBucketVersions",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:ListAllMyBuckets",
        "s3:GetBucketTagging",
        "s3:GetBucketLocation",
        "s3:GetBucketPolicyStatus",
        "s3:GetBucketPublicAccessBlock",
        "s3:GetBucketAcl",
        "s3:GetBucketPolicy",
        "s3:PutBucketPublicAccessBlock"
    ],
    "Resource": [
        "arn:aws:s3:::netapp-backup-*"
    ]
}
```

## Enabling Backup to Cloud on a new system

Backup to Cloud is enabled by default in the working environment wizard. Be sure to keep the option enabled.

### Steps

1. Click **Create Cloud Volumes ONTAP**.
2. Select Amazon Web Services as the cloud provider and then choose a single node or HA system.
3. Fill out the Details & Credentials page.
4. On the Services page, leave the service enabled and click **Continue**.



5. Complete the pages in the wizard to deploy the system.

### Result

Backup to Cloud is enabled on the system and backs up volumes every day and retains the most recent 30 backup copies.

### What's next?

[You can manage backups by changing the backup schedule, restoring volumes, and more.](#)

## Enabling Backup to Cloud on an existing system

Enable Backup to Cloud at any time directly from the working environment.

### Steps

1. Select the working environment and click **Activate** next to the Backup to Cloud service in the right-panel.



2. Define the backup schedule and retention value and click **Continue**.

### Define Policy

**Policy - Retention & Schedule**

Create a New Policy   
  Select an Existing Policy

Backup Every:    
 Number of backups to retain:

---

**DP Volumes**

Data protection volume backups use the same retention period as defined in the source SnapMirror relationship by default. Use the API if you want to change this value

---

**Information**

Backup\_Bucket\_Name  
Bucket Name

See [the list of existing policies](#).

3. Select the volumes that you want to back up and click **Activate**.

### Select Volumes

57 Volumes 🔍

<input checked="" type="checkbox"/>	Volume Name	Volume Type	Disk Type	SVM Name	Used Capacity	Allocated Capacity	Volume Status
<input checked="" type="checkbox"/>	Volume_Name_1	RW	GP2	SVM_Name_1	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_2	RW	GP2	SVM_Name_2	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_3	RW	GP2	SVM_Name_3	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_4	DP	GP2	SVM_Name_4	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_5	RW	GP2	SVM_Name_5	2.25 TB	10 TB	Active

### Result

Backup to Cloud starts taking the initial backups of each selected volume.

### What's next?

[You can manage backups by changing the backup schedule, restoring volumes, and more.](#)

## Backing up data to Azure Blob storage

Complete a few steps to get started backing up data from Cloud Volumes ONTAP to Azure Blob storage.

### Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.

# 1

## Verify support for your configuration

- You're running Cloud Volumes ONTAP 9.7 or later in Azure.
- You have a valid cloud provider subscription for the storage space where your backups will be located.
- You have subscribed to the [Cloud Manager Marketplace Backup offering](#), or you have purchased and [activated](#) a Backup to Cloud BYOL license from NetApp.

# 2

## Enable Backup to Cloud on your new or existing system

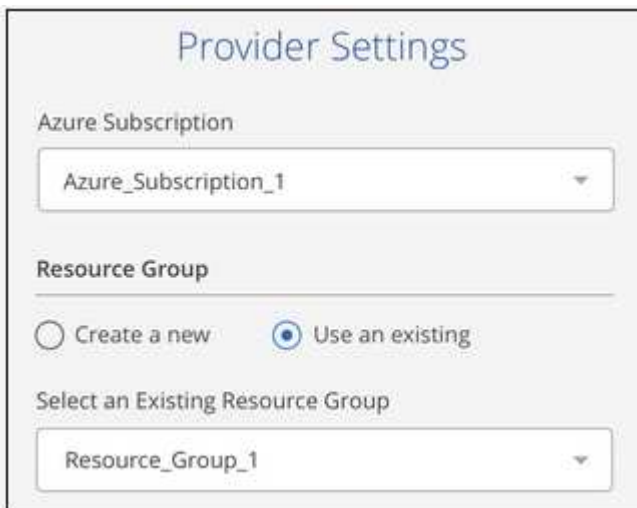
- New systems: Backup to Cloud is enabled by default in the working environment wizard. Be sure to keep the option enabled.
- Existing systems: Select the working environment and click **Activate** next to the Backup to Cloud service in the right-panel, and then follow the setup wizard.



# 3

## Enter the provider details

Select the provider subscription and choose whether you want to create a new resource group or use an already existing resource group.

A screenshot of a 'Provider Settings' configuration window. The title 'Provider Settings' is at the top. Below it, there are three sections: 1. 'Azure Subscription' with a dropdown menu showing 'Azure\_Subscription\_1'. 2. 'Resource Group' with two radio buttons: 'Create a new' (unselected) and 'Use an existing' (selected). 3. 'Select an Existing Resource Group' with a dropdown menu showing 'Resource\_Group\_1'.

# 4

## Define the backup policy

The default policy backs up volumes every day and retains the most recent 30 backup copies of each volume. Change to weekly or monthly backups, or select one of the system-defined policies that provide more options.



### Define Policy

**Policy - Retention & Schedule**

Create a New Policy   
  Select an Existing Policy

Backup Every:    
 Number of backups to retain:

**DP Volumes**

Data protection volume backups use the same retention period as defined in the source SnapMirror relationship by default. Use the API if you want to change this value

**Storage Account**

Cloud Manager will create the storage account after you complete the wizard

## 5

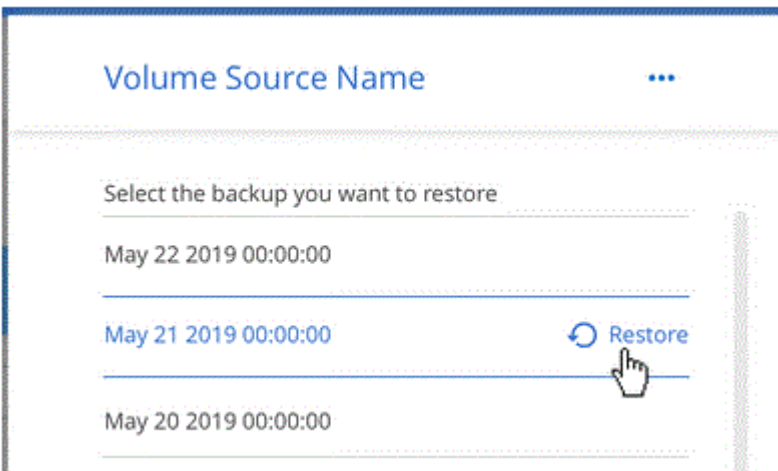
### Select the volumes that you want to back up

Identify which volumes you want to back up in the Select Volumes page.

## 6

### Restore your data, as needed

From the Backup List, select a volume, select a backup, and then restore data from the backup to a new volume.



### Requirements

Read the following requirements to make sure that you have a supported configuration before you start backing up volumes to Azure Blob storage.

### Supported ONTAP versions

Cloud Volumes ONTAP 9.7 and later.

### Supported Azure regions

Backup to Cloud is supported in all Azure regions [where Cloud Volumes ONTAP is supported](#).

## License requirements

For Backup to Cloud PAYGO licensing, a subscription through the Azure Marketplace is required before you enable Backup to Cloud. Billing for Backup to Cloud is done through this subscription. [You can subscribe from the Details & Credentials page of the working environment wizard.](#)

For Backup to Cloud BYOL licensing, you need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. See [Adding and updating your Backup BYOL license.](#)

And you need to have a Microsoft Azure subscription for the storage space where your backups will be located.

## Enabling Backup to Cloud on a new system

Backup to Cloud is enabled by default in the working environment wizard. Be sure to keep the option enabled.



If you want to pick the name of the resource group, **disable** Backup to Cloud when deploying Cloud Volumes ONTAP. Follow the steps for [enabling backup to cloud on an existing system](#) to enable Backup to Cloud and choose the resource group.

## Steps

1. Click **Create Cloud Volumes ONTAP**.
2. Select Microsoft Azure as the cloud provider and then choose a single node or HA system.
3. Fill out the Details & Credentials page and be sure that an Azure Marketplace subscription is in place.
4. On the Services page, leave the service enabled and click **Continue**.

The screenshot shows the 'Backup to Cloud' service configuration page. At the top left is a cloud icon with a checkmark, followed by the text 'Backup to Cloud'. To the right is a blue toggle switch in the 'on' position. Below this is a descriptive paragraph: 'Integrated backup for Cloud Volumes ONTAP based on SnapMirror and Snapshot technologies. Backup copies are maintained in Storage Accounts. Backups stored in Storage Accounts are charged separately from Cloud Volumes ONTAP.' Underneath are two columns: 'ADVANTAGES' and 'CLARIFICATIONS'. The 'ADVANTAGES' column lists three items with checkmarks: 'Automatically back up all volumes.', 'Creates new backup copy every day.', and 'Retains backups for 30 days.' The 'CLARIFICATIONS' column has one item with a right-pointing arrow: 'Backup settings are editable after working environment creation.'

5. Complete the pages in the wizard to deploy the system.

## Result

Backup to Cloud is enabled on the system and backs up volumes every day and retains the most recent 30 backup copies.

## What's next?

[You can manage backups by changing the backup schedule, restoring volumes, and more.](#)

## Enabling Backup to Cloud on an existing system

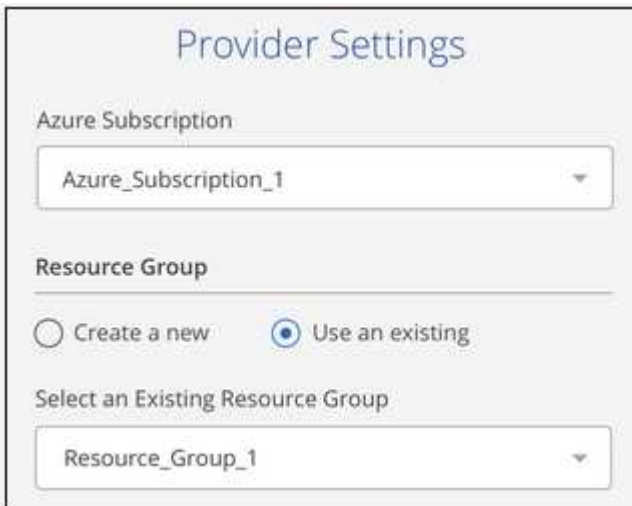
Enable Backup to Cloud at any time directly from the working environment.

### Steps

1. Select the working environment and click **Activate** next to the Backup to Cloud service in the right-panel.

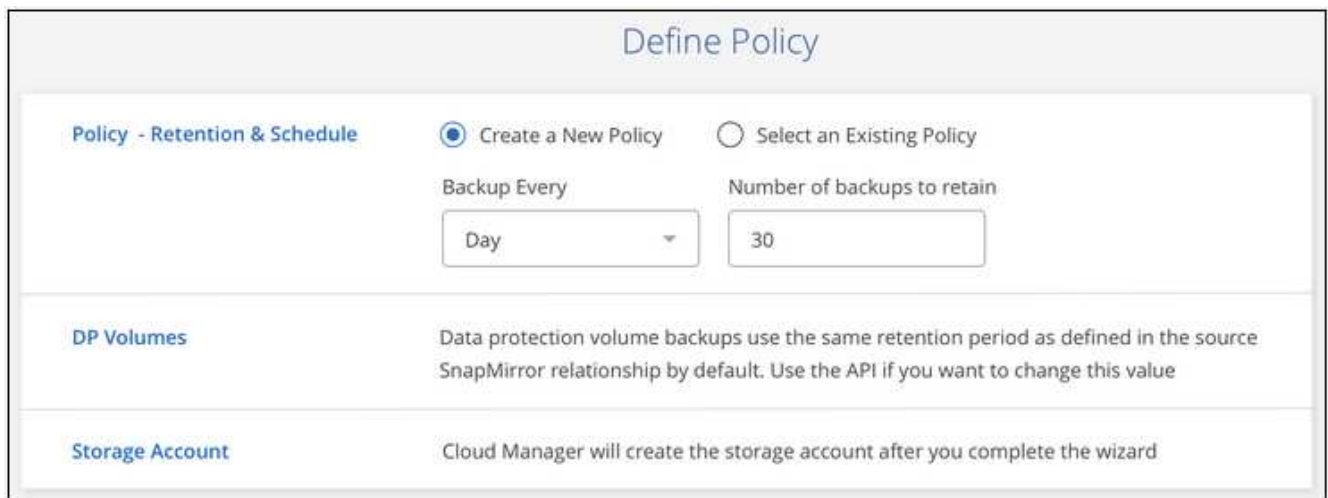


2. Select the provider details:
  - a. The Azure subscription used to store the backups.
  - b. The resource group - you can create a new resource group or select an existing resource group.
  - c. And then click **Continue**.

A screenshot of the 'Provider Settings' configuration page. The title 'Provider Settings' is at the top. Below it, there are three sections: 'Azure Subscription' with a dropdown menu showing 'Azure\_Subscription\_1'; 'Resource Group' with two radio buttons, 'Create a new' (unselected) and 'Use an existing' (selected); and 'Select an Existing Resource Group' with a dropdown menu showing 'Resource\_Group\_1'.

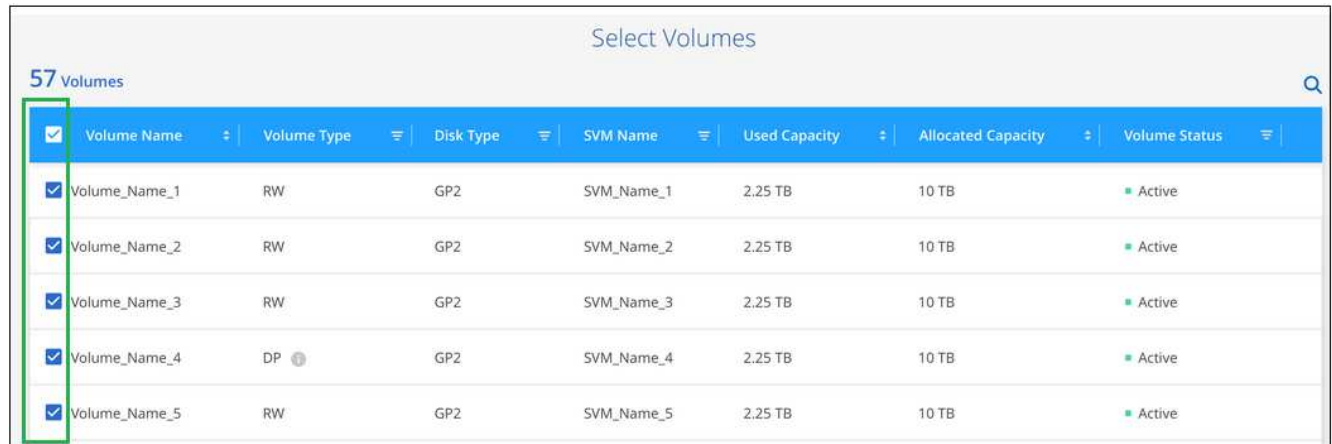
Note that you cannot change the subscription or the resource group after the services has started.

3. In the *Define Policy* page, select the backup schedule and retention value and click **Continue**.

A screenshot of the 'Define Policy' configuration page. The title 'Define Policy' is at the top. Below it, there are three sections: 'Policy - Retention & Schedule' with two radio buttons, 'Create a New Policy' (selected) and 'Select an Existing Policy' (unselected); 'Backup Every' with a dropdown menu showing 'Day' and 'Number of backups to retain' with a text input field containing '30'; 'DP Volumes' with the text 'Data protection volume backups use the same retention period as defined in the source SnapMirror relationship by default. Use the API if you want to change this value'; and 'Storage Account' with the text 'Cloud Manager will create the storage account after you complete the wizard'.

See [the list of existing policies](#).

4. Select the volumes that you want to back up and click **Activate**.



<input checked="" type="checkbox"/>	Volume Name	Volume Type	Disk Type	SVM Name	Used Capacity	Allocated Capacity	Volume Status
<input checked="" type="checkbox"/>	Volume_Name_1	RW	GP2	SVM_Name_1	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_2	RW	GP2	SVM_Name_2	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_3	RW	GP2	SVM_Name_3	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_4	DP	GP2	SVM_Name_4	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_5	RW	GP2	SVM_Name_5	2.25 TB	10 TB	Active

## Result

Backup to Cloud starts taking the initial backups of each selected volume.

## What's next?

You can [manage backups by changing the backup schedule, restoring volumes, and more](#).

## Backing up data from an on-premises ONTAP system to the cloud

Complete a few steps to get started backing up data from your on-premises ONTAP system to low-cost object storage in the cloud.

## Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.

### 1

#### Verify support for your configuration

- You have discovered the on-premises cluster and added it to a working environment in Cloud Manager. See [Discovering ONTAP clusters](#) for details.
- You're running ONTAP 9.7P5 or later on the cluster.
- You have a valid cloud provider subscription for the storage space where your backups will be located.
- You have subscribed to the [Cloud Manager Marketplace Backup offering](#), or you have purchased and [activated](#) a Backup to Cloud BYOL license from NetApp.

### 2

#### Enable Backup to Cloud on the system

Select the working environment and click **Activate** next to the Backup to Cloud service in the right-panel, and then follow the setup wizard.



### 3 Select the cloud provider and enter provider details

Select the provider and then select the provider subscription, the region, and resource group. You also need to specify the IPspace in the ONTAP cluster where the volumes reside.

**Provider Settings**

**Provider Information**

Azure Subscription  
Azure\_Subscription\_1

Region  
Default\_CM\_Region

IPspace  
IP\_Space\_1

**Resource Group**

Create a new  Use an existing

Select an Existing Resource Group  
Resource\_Group\_1

### 4 Define the backup policy

The default policy backs up volumes every day and retains the most recent 30 backup copies of each volume. Change to weekly or monthly backups, or select one of the system-defined policies that provide more options.

**Define Policy**

**Policy - Retention & Schedule**

Create a New Policy  Select an Existing Policy

Backup Every: Day

Number of backups to retain: 30

**DP Volumes**

Data protection volume backups use the same retention period as defined in the source SnapMirror relationship by default. Use the API if you want to change this value

**Storage Account**

Cloud Manager will create the storage account after you complete the wizard

## 5


### Select the volumes that you want to back up

Identify which volumes you want to back up from the cluster.

## 6

### Restore your data, as needed

From the Backup List, select a volume, select a backup, and then restore data from the backup to a new volume on a Cloud Volumes ONTAP system that is using the same cloud provider.



### Requirements

Read the following requirements to make sure you have a supported configuration before you start backing up volumes to Azure Blob storage.

#### Supported ONTAP versions

ONTAP 9.7P5 and later.

#### Cluster networking requirements

An intercluster LIF is required on each ONTAP node that hosts the volumes you want to back up. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage. The Admin SVM must reside on the IPspace. [Learn more about IPspaces.](#)

When you set up backup to cloud, you are prompted for the IPspace to use. You should choose the IPspace that each LIF is associated with. That might be the "Default" IPspace or a custom IPspace that you created.

#### Supported Azure regions

Backup to Cloud is supported in all Azure regions [where cloud volumes are supported.](#)

#### License requirements

For Backup to Cloud PAYGO licensing, a subscription to the [Azure Marketplace Cloud Manager Backup offering](#) is required before you enable Backup to Cloud. Billing for Backup to Cloud is done through this subscription.

For Backup to Cloud BYOL licensing, you need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. See [Adding and updating your Backup BYOL license.](#)

And you need to have a Microsoft Azure subscription for the storage space where your backups will be located.

## Enabling Backup to Cloud

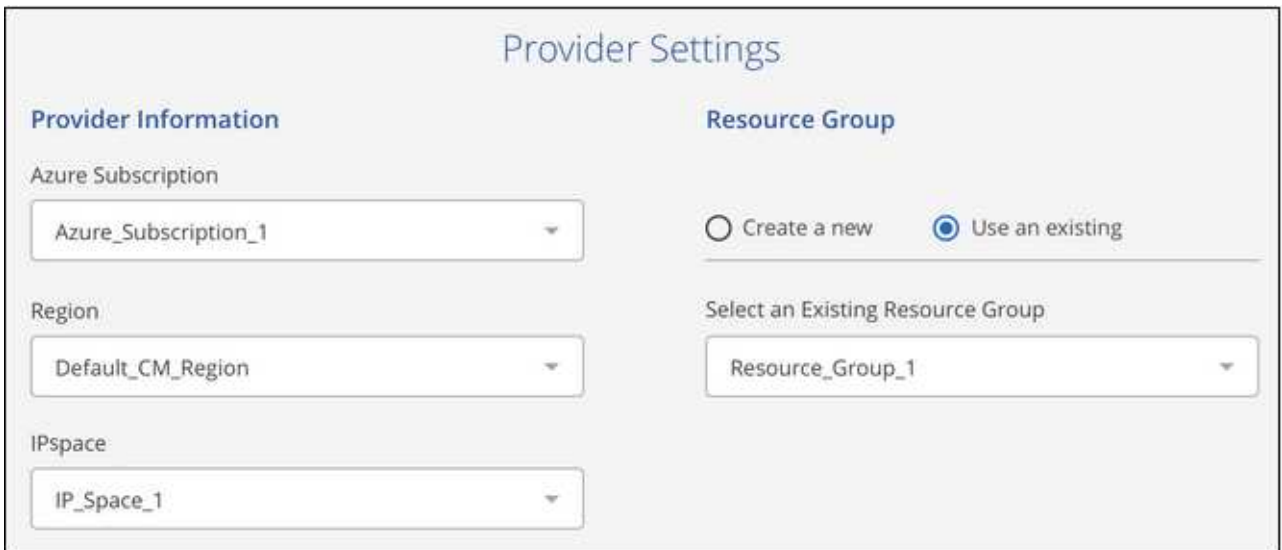
Enable Backup to Cloud at any time directly from the working environment.

### Steps

1. Select the working environment and click **Activate** next to the Backup to Cloud service in the right-panel.



2. Select the provider, and then enter the provider details:
  - a. The Azure subscription used to store the backups.
  - b. The Azure region.
  - c. The resource group - you can create a new resource group or select an existing resource group.
  - d. The IPspace in the ONTAP cluster where the volumes you want to back up reside.
  - e. And then click **Continue**.



Note that you cannot change the subscription or the resource group after the services has started.

3. In the *Define Policy* page, select the backup schedule and retention value and click **Continue**.

### Define Policy

**Policy - Retention & Schedule**

Create a New Policy   
  Select an Existing Policy

Backup Every: 
     
 Number of backups to retain:

---

**DP Volumes**

Data protection volume backups use the same retention period as defined in the source SnapMirror relationship by default. Use the API if you want to change this value

---

**Storage Account**

Cloud Manager will create the storage account after you complete the wizard

See [the list of existing policies](#).

- Select the volumes that you want to back up and click **Activate**.

### Select Volumes

57 Volumes 🔍

<input checked="" type="checkbox"/>	Volume Name	Volume Type	SVM Name	Used Capacity	Allocated Capacity	Volume Status
<input checked="" type="checkbox"/>	Volume_Name_1	RW	SVM_Name_1	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_2	RW	SVM_Name_2	2.25 TB	10 TB	Active
<input checked="" type="checkbox"/>	Volume_Name_3	RW	SVM_Name_3	2.25 TB	10 TB	Active

### Result

Backup to Cloud starts taking the initial backups of each selected volume.

### What's next?

You can [manage backups by changing the backup schedule, restoring volumes, and more](#).

## Managing backups for Cloud Volumes ONTAP and on-premises ONTAP systems

Manage backups for Cloud Volumes ONTAP and on-premises ONTAP systems by changing the backup schedule, restoring volumes, deleting backups, and more.

### Changing the schedule and backup retention


The default policy backs up volumes every day and retains the most recent 30 backup copies of each volume. You can change to weekly or monthly backups and you can change the number of backup copies to retain. You can also select one of the system-defined policies that provide scheduled backups for 3 months, 1 year, and 7 years.






Changing the backup policy affects only new volumes created after you change the schedule. It doesn't affect the schedule for any existing volumes.

## Steps


1. Select the working environment.
2. Click  and select **Backup Settings**.



3. From the *Backup Settings* page, click  for the working environment and select **Modify Backup Policy**.




4. From the *Modify Backup Policy* page, change the schedule and backup retention and then click **Save**.




## Starting and stopping backups of volumes

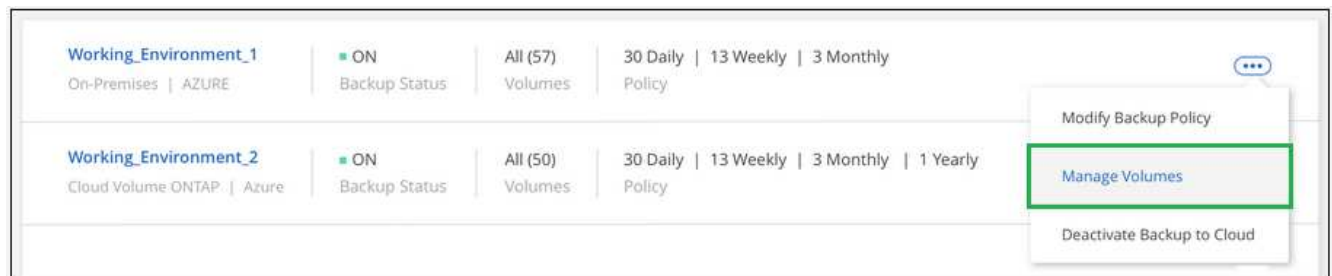
You can stop backing up a volume if you do not need backup copies of that volume and you do not want to pay for the cost to store the backups. You can also add a new volume to the backup list if it is not currently being backed up.

### Steps

1. Select the working environment.
2. Click  and select **Backup Settings**.



3. From the *Backup Settings* page, click  for the working environment and select **Manage Volumes**.



4. Select the checkbox for volumes that you want to start backing up, and deselect the checkbox for volumes that you want to stop backing up.



**Note:** When stopping a volume from being backed up you'll continue to be charged by your cloud provider for object storage costs for the capacity that the backups use unless you [delete the backups](#).


## Restoring a volume from a backup

When you restore data from a backup, Cloud Manager creates a *new* volume using the data from the backup. You can restore the data to a volume in the same working environment or to a different working environment that's located in the same cloud account as the source working environment. Because the backup does not contain any snapshots, the newly restored volume does not either.



Backups created from on-premises ONTAP systems can be restored only to Cloud Volumes ONTAP systems that use the same cloud provider as where the backup resides.

### Steps


1. Select the working environment.
2. Click  and select **View Backups**.



3. Select the row for the volume that you want to restore and click **View Backup List**.

6 of 16 Volumes						
Working Environment	Source Volume	Last Backup	Policy & Retention	Relationship Status		
gfcDevQaSaCvo (On)	cifsvol9 (Available)	Aug 13, 2020 02:00:12 PM UTC	30 Daily	Active (Idle)	<a href="#">View Backup List</a>	
gfcDevQaSaCvo (On)	smbvol (Available)	Aug 13, 2020 02:00:33 PM UTC	30 Daily	Active (Idle)	<a href="#">View Backup List</a>	

4. Find the backup that you want to restore and click the **Restore** icon.



5. Fill out the *Restore Backup to new volume* page:

- a. Select the working environment to which you want to restore the volume.
- b. Enter a name for the volume.
- c. Click **Restore**.

## Result

Cloud Manager creates a new volume based on the backup you selected. You can [manage this new volume](#) as required.

## Deleting backups

Backup to Cloud enables you to delete *all* backups of a specific volume. You can't delete *individual* backups.


You might do this if you no longer need the backups or if you deleted the source volume and want to remove all backups.



If you plan to delete a Cloud Volumes ONTAP or on-premises ONTAP system that has backups, you must delete the backups **before** deleting the system. Backup to Cloud doesn't automatically delete backups when you delete a system, and there is no current support in the UI to delete the backups after the system has been deleted.

## Steps

1. At the top of Cloud Manager, click **Backup**.
2. From the volume list, find the volume and click **View Backup List**.
3. Click **...** and select **Delete All Backups**.




4. In the confirmation dialog box, click **Delete**.


## Disabling Backup to Cloud


Disabling Backup to Cloud for a working environment disables backups of each volume on the system, and it also disables the ability to restore a volume. Any existing backups will not be deleted.

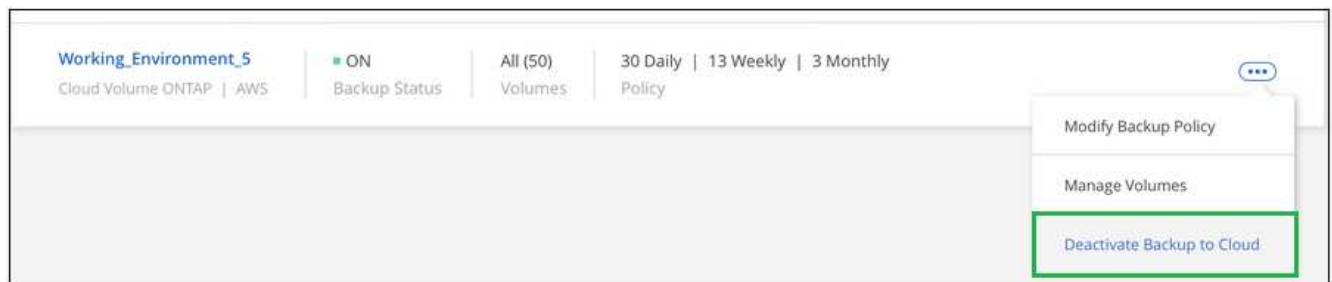
Note that you'll continue to be charged by your cloud provider for object storage costs for the capacity that your backups use unless you delete the backups.

### Steps

1. Select the working environment.
2. Click  and select **Backup Settings**.



3. From the *Backup Settings* page, click  for the working environment and select **Deactivate Backup to Cloud**.



4. In the confirmation dialog box, click **Deactivate**.


# Copy and synchronize data

## Cloud Sync overview

The NetApp Cloud Sync service offers a simple, secure, and automated way to migrate your data to any target, in the cloud or on your premises. Whether it's a file-based NAS dataset (NFS or SMB), Amazon Simple Storage Service (S3) object format, a NetApp StorageGRID® appliance, or any other cloud provider object store, Cloud Sync can convert and move it for you.

### Features


Watch the following video for an overview of Cloud Sync:



### How Cloud Sync works

Cloud Sync is a software-as-a-service (SaaS) platform that consists of a data broker, a cloud-based interface available through Cloud Manager, and a source and target.

The following image shows the relationship between Cloud Sync components:



The NetApp data broker software syncs data from a source to a target (this is called a *sync relationship*). You can run the data broker in AWS, Azure, Google Cloud Platform, or on your premises. The data broker needs an outbound internet connection over port 443 so it can communicate with the Cloud Sync service and contact a few other services and repositories. [View the list of endpoints.](#)

After the initial copy, the service syncs any changed data based on the schedule that you set.

## Supported storage types

Cloud Sync supports the following storage types:

- Any NFS server
- Any SMB server
- AWS EFS
- AWS S3
- Azure Blob
- Azure NetApp Files
- Cloud Volumes Service
- Cloud Volumes ONTAP
- Google Cloud Storage
- IBM Cloud Object Storage
- On-premises ONTAP cluster
- ONTAP S3 Storage

- [StorageGRID](#)

[Review the supported sync relationships.](#)

## Cost

There are two types of costs associated with using Cloud Sync: resource charges and service charges.

### Resource charges

Resource charges are related to the compute and storage costs for running the data broker in the cloud.

### Service charges

There are two ways to pay for sync relationships after your 14-day free trial ends. The first option is to subscribe from AWS or Azure, which enables you to pay hourly or annually. The second option is to purchase licenses directly from NetApp. Read the following sections for more details.

### Marketplace subscription

Subscribing to the Cloud Sync service from AWS or Azure enables you to pay at an hourly rate, or to pay annually. [You can subscribe through either AWS or Azure](#), depending on where you want to be billed.

### Hourly subscriptions

With an hourly pay-as-you-go subscription, the Cloud Sync service charges hourly based on the number of sync relationships that you create.

- [View pricing in Azure](#)
- [View pay-as-you-go pricing in AWS](#)

### Annual subscriptions

An annual subscription provides a license for 20 sync relationships that you pay for up front. If you go above 20 sync relationships and you've subscribed through Azure, you pay for the additional relationships by the hour.

[View annual pricing in AWS](#)

### Licenses from NetApp

Another way to pay for sync relationships up front is by purchasing licenses directly from NetApp. Each license enables you to create up to 20 sync relationships.

You can use these licenses with an AWS or Azure subscription. For example, if you have 25 sync relationships, you can pay for the first 20 sync relationships using a license and then pay-as-you-go from AWS or Azure with the remaining 5 sync relationships.

[Learn how to purchase licenses and add them to Cloud Sync.](#)

### License terms

Customers who purchase a Bring Your Own License (BYOL) to the Cloud Sync service should be aware of limitations associated with the license entitlement.

- Customers are entitled to leverage the BYOL license for a term not to exceed one year from the date of delivery.



- Customers are entitled to leverage the BYOL license to establish and not to exceed a total of 20 individual connections between a source and a target (each a “sync relationship”).
- A customer’s entitlement expires at the conclusion of the one-year license term, irrespective as to whether Customer has reached the 20 sync relationship limitation.
- In the event the Customer chooses to renew its license, unused sync relationships associated from the previous license grant DO NOT roll over to the license renewal.

## Data privacy

NetApp doesn’t have access to any credentials that you provide while using the Cloud Sync service. The credentials are stored directly on the data broker machine, which resides in your network.

Depending on the configuration that you choose, Cloud Sync might prompt you for credentials when you create a new relationship. For example, when setting up a relationship that includes an SMB server, or when deploying the data broker in AWS.

These credentials are always saved directly to the data broker itself. The data broker resides on a machine in your network, whether it’s on premises or in your cloud account. The credentials are never made available to NetApp.

The credentials are locally encrypted on the data broker machine using HashiCorp Vault.

## Limitations

- Cloud Sync is not supported in China.
- In addition to China, the Cloud Sync data broker is not supported in the following regions:
  - AWS GovCloud (US)
  - Azure US Gov
  - Azure US DoD

## Get started

### Quick start for Cloud Sync

Getting started with the Cloud Sync service includes a few steps.



#### Prepare your source and target

Verify that your source and target are supported and set up. The most important requirement is to verify connectivity between the data broker and the source and target locations. [Learn more](#).



#### Prepare a location for the NetApp data broker

The NetApp data broker software syncs data from a source to a target (this is called a *sync relationship*). You can run the data broker in AWS, Azure, Google Cloud Platform, or on your premises. The data broker needs an outbound internet connection over port 443 so it can communicate with the Cloud Sync service and contact a few other services and repositories. [View the list of endpoints](#).

Cloud Sync guides you through the installation process when you create a sync relationship, at which point you can deploy the data broker in the cloud or download an install script for your own Linux host.

- [Review AWS installation](#)
- [Review Azure installation](#)
- [Review GCP installation](#)
- [Review Linux host installation](#)



### Create your first sync relationship

Log in to [Cloud Manager](#), click **Sync**, and then drag and drop your selections for the source and target. Follow the prompts to complete the setup. [Learn more](#).



### Pay for your sync relationships after your free trial ends

Subscribe from AWS or Azure to pay-as-you-go or to pay annually. Or purchase licenses directly from NetApp. Just go to the License Settings page in Cloud Sync to set it up. [Learn more](#).

## Preparing the source and target

Prepare to sync data by verifying that your source and target are supported and setup.

### Supported sync relationships

Cloud Sync enables you to sync data from a source to a target (this is called a *sync relationship*). You should understand the supported relationships before you get started.

Source location	Supported target locations
AWS EFS	<ul style="list-style-type: none"><li>• AWS EFS</li><li>• AWS S3</li><li>• Azure Blob</li><li>• Azure NetApp Files (NFS)</li><li>• Cloud Volumes ONTAP (NFS)</li><li>• Cloud Volumes Service (NFS)</li><li>• IBM Cloud Object Storage</li><li>• Google Cloud Storage</li><li>• NFS server</li><li>• On-premises ONTAP cluster</li><li>• StorageGRID</li></ul>

Source location	Supported target locations
AWS S3	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
Azure Blob	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
Azure NetApp Files (NFS)	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>
Azure NetApp Files (SMB)	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (SMB)</li> <li>• Cloud Volumes Service (SMB)</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
Cloud Volumes ONTAP (NFS)	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
Cloud Volumes ONTAP (SMB)	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (SMB)</li> <li>• Cloud Volumes Service (SMB)</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
Cloud Volumes Service (NFS)	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>
Cloud Volumes Service (SMB)	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (SMB)</li> <li>• Cloud Volumes Service (SMB)</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
Google Cloud Storage	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
IBM Cloud Object Storage	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
NFS server	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>
On-prem ONTAP cluster (NFS)	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (NFS)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• StorageGRID</li> </ul>
On-prem ONTAP cluster (SMB)	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (SMB)</li> <li>• Cloud Volumes Service (SMB)</li> <li>• Google Cloud Storage</li> <li>• IBM Cloud Object Storage</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
ONTAP S3 Storage	<ul style="list-style-type: none"> <li>• StorageGRID</li> </ul>

Source location	Supported target locations
SMB server	<ul style="list-style-type: none"> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files (SMB)</li> <li>• Cloud Volumes ONTAP (NFS)</li> <li>• Cloud Volumes Service (NFS)</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• On-premises ONTAP cluster</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>
StorageGRID	<ul style="list-style-type: none"> <li>• AWS EFS</li> <li>• AWS S3</li> <li>• Azure Blob</li> <li>• Azure NetApp Files</li> <li>• Cloud Volumes ONTAP</li> <li>• Cloud Volumes Service</li> <li>• IBM Cloud Object Storage</li> <li>• Google Cloud Storage</li> <li>• NFS server</li> <li>• On-premises ONTAP cluster</li> <li>• ONTAP S3 Storage</li> <li>• SMB Server</li> <li>• StorageGRID</li> </ul>

**Notes:**

1. You can choose a specific Azure Blob storage tier when a Blob container is the target:
  - Hot storage
  - Cool storage
2. You can choose a specific S3 storage class when AWS S3 is the target:
  - Standard (this is the default class)
  - Intelligent-Tiering
  - Standard-Infrequent Access
  - One Zone-Infrequent Access
  - Glacier
  - Glacier Deep Archive



## Networking for the source and target

- The source and target must have a network connection to the data broker.

For example, if an NFS server is in your data center and the data broker is in AWS, then you need a network connection (VPN or Direct Connect) from your network to the VPC.

- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

## Source and target requirements

Verify that your source and targets meet the following requirements.

### AWS S3 bucket requirements

Make sure that your AWS S3 bucket meets the following requirements.

## Supported data broker locations for AWS S3

Sync relationships that include S3 storage require a data broker deployed in AWS or on your premises. In either case, Cloud Sync prompts you to associate the data broker with an AWS account during installation.

- [Learn how to deploy the AWS data broker](#)
- [Learn how to install the data broker on a Linux host](#)

## Supported AWS regions

All regions are supported except for the China and GovCloud (US) regions.

## Permissions required for S3 buckets in other AWS accounts

When setting up a sync relationship, you can specify an S3 bucket that resides in an AWS account that isn't associated with the data broker.

[The permissions included in this JSON file](#) must be applied to that S3 bucket so the data broker can access it. These permissions enable the data broker to copy data to and from the bucket and to list the objects in the bucket.

Note the following about the permissions included in the JSON file:

1. *<BucketName>* is the name of the bucket that resides in the AWS account that isn't associated with the data broker.
2. *<RoleARN>* should be replaced with one of the following:
  - If the data broker was manually installed on a Linux host, *RoleARN* should be the ARN of the AWS user for which you provided AWS credentials when deploying the data broker.
  - If the data broker was deployed in AWS using the CloudFormation template, *RoleARN* should be the ARN of the IAM role created by the template.

You can find the Role ARN by going to the EC2 console, selecting the data broker instance, and clicking the IAM role from the Description tab. You should then see the Summary page in the IAM console that contains the Role ARN.

# Summary

Delete role

**Role ARN** `arn:aws:iam::[account ID]:role/tanyaBroker0304-DataBrokerIamRole-1VMHXXMW3AQ05`

Role description [Edit](#)

## Azure Blob storage requirements

Make sure that your Azure Blob storage meets the following requirements.

## Supported data broker locations for Azure Blob

The data broker can reside in any location when a sync relationship includes Azure Blob storage.

## Supported Azure regions

All regions are supported except for the China, US Gov, and US DoD regions.

## Connection string required for relationships that include Azure Blob and NFS/SMB

When creating a sync relationship between an Azure Blob container and an NFS or SMB server, you need to provide Cloud Sync with the storage account connection string:

**a63cde60b553020** - Access keys

Storage account

Search (Ctrl+/)

- Overview
- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems
- Storage Explorer (preview)
- Settings
  - Access keys**
  - CORS
  - Configuration
  - Encryption

Use access keys to authenticate your applications when making requests to this Azure storage account. Store your access keys securely - for example, using Azure Key Vault - and don't share them. We recommend regenerating your access keys regularly. You are provided two access keys so that you can maintain connections using one key while regenerating the other.

When you regenerate your access keys, you must update any Azure resources and applications that access this storage account to use the new keys. This action will not interrupt access to disks from your virtual machines. [Learn more](#)

Storage account name: a63cde60b553020

**key1**

Key: vScjFdvVZqIPyO/

**Connection string**: DefaultEndpoints

If you want to sync data between two Azure Blob containers, then the connection string must include a [shared access signature \(SAS\)](#). You also have the option to use a SAS when syncing between a Blob container and an NFS or SMB server.

The SAS must allow access to the Blob service and all resource types (Service, Container, and Object). The SAS must also include the following permissions:

- For the source Blob container: Read and List
- For the target Blob container: Read, Write, List, Add, and Create

Search (Ctrl+/)

Overview  
Activity log  
Access control (IAM)  
Tags  
Diagnose and solve problems  
Storage Explorer (preview)

Settings

Access keys  
CORS  
Configuration  
Encryption  
**Shared access signature**  
Firewalls and virtual networks  
Advanced Threat Protection (pr...  
Properties  
Locks

Allowed services ⓘ  
 Blob  File  Queue  Table

Allowed resource types ⓘ  
 Service  Container  Object

Allowed permissions ⓘ  
 Read  Write  Delete  List  Add  Create  Update  Process

Start and expiry date/time ⓘ  
Start  
2018-10-23 10:07:32 AM  
End  
2019-10-23 6:07:32 PM  
(UTC-04:00) --- Current Time Zone ---

Allowed IP addresses ⓘ  
for example, 168.1.5.65 or 168.1.5.65-168.1.5.70

Allowed protocols ⓘ  
 HTTPS only  HTTPS and HTTP

Signing key ⓘ  
key1

**Generate SAS and connection string**

### Azure NetApp Files requirement

Use the Premium or Ultra service level when you sync data to or from Azure NetApp Files. You might experience failures and performance issues if the disk service level is Standard.



Consult a solutions architect if you need help determining the right service level. The volume size and volume tier determines the throughput that you can get.

[Learn more about Azure NetApp Files service levels and throughput.](#)

### Google Cloud Storage bucket requirements

Make sure that your Google Cloud Storage bucket meets the following requirements.

### Supported data broker locations for Google Cloud Storage

Sync relationships that include Google Cloud Storage require a data broker deployed in GCP or on your premises. Cloud Sync guides you through the data broker installation process when you create a sync relationship.

- [Learn how to deploy the GCP data broker](#)

- [Learn how to install the data broker on a Linux host](#)

## Supported GCP regions

All regions are supported.

## NFS server requirements

- The NFS server can be a NetApp system or a non-NetApp system.
- The file server must allow the data broker host to access the exports.
- NFS versions 3, 4.0, 4.1, and 4.2 are supported.

The desired version must be enabled on the server.

- If you want to sync NFS data from an ONTAP system, ensure that access to the NFS export list for an SVM is enabled (`vserver nfs modify -vserver svm_name -showmount enabled`).



The default setting for showmount is *enabled* starting with ONTAP 9.2.

## ONTAP S3 Storage requirements

ONTAP 9.7 supports the Amazon Simple Storage Service (Amazon S3) as a public preview. [Learn more about ONTAP support for Amazon S3.](#)

When you set up a sync relationship that includes ONTAP S3 Storage, you'll need to provide the following:

- The IP address of the LIF that's connected to ONTAP S3
- The access key and secret key that ONTAP is configured to use

## SMB server requirements

- The SMB server can be a NetApp system or a non-NetApp system.
- The file server must allow the data broker host to access the exports.
- SMB versions 1.0, 2.0, 2.1, 3.0 and 3.11 are supported.
- Grant the "Administrators" group with "Full Control" permissions to the source and target folders.

If you don't grant this permission, then the data broker might not have sufficient permissions to get the ACLs on a file or directory. If this occurs, you'll receive the following error: "getxattr error 95"

## SMB limitation for hidden directories and files

An SMB limitation affects hidden directories and files when syncing data between SMB servers. If any of the directories or files on the source SMB server were hidden through Windows, the hidden attribute isn't copied to the target SMB server.

## SMB sync behavior due to case-insensitivity limitation

The SMB protocol is case-insensitive, which means uppercase and lowercase letters are treated as being the same. This behavior can result in overwritten files and directory copy errors, if a sync relationship includes an SMB server and data already exists on the target.

For example, let's say that there's a file named "a" on the source and a file named "A" on the target. When Cloud Sync copies the file named "a" to the target, file "A" is overwritten by file "a" from the source.

In the case of directories, let's say that there's a directory named "b" on the source and a directory named "B" on the target. When Cloud Sync tries to copy the directory named "b" to the target, Cloud Sync receives an error that says the directory already exists. As a result, Cloud Sync always fails to copy the directory named "b."

The best way to avoid this limitation is to ensure that you sync data to an empty directory.

### **Permissions for a SnapMirror destination**

If the source for a sync relationship is a SnapMirror destination (which is read-only), "read/list" permissions are sufficient to sync data from the source to a target.

## **Networking overview for Cloud Sync**

Networking for Cloud Sync includes connectivity between the data broker and the source and target locations, and an outbound internet connection from the data broker over port 443.

### **Data broker location**


You can install the data broker in the cloud or on your premises.

#### **Data broker in the cloud**

The following image shows the data broker running in the cloud, in either AWS, GCP, or Azure. The source and target can be in any location, as long as there's a connection to the data broker. For example, you might have a VPN connection from your data center to your cloud provider.




When Cloud Sync deploys the data broker in AWS, Azure, or GCP, it creates a security group that enables the required outbound communication.



#### Data broker on your premises

The following image shows the data broker running on-prem, in a data center. Again, the source and target can be in any location, as long as there's a connection to the data broker.



## Networking requirements

- The source and target must have a network connection to the data broker.

For example, if an NFS server is in your data center and the data broker is in AWS, then you need a network connection (VPN or Direct Connect) from your network to the VPC.

- The data broker needs an outbound internet connection so it can poll the Cloud Sync service for tasks over port 443.
- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

## Networking endpoints

The NetApp data broker requires outbound internet access over port 443 to communicate with the Cloud Sync service and to contact a few other services and repositories. Your local web browser also requires access to endpoints for certain actions. If you need to limit outbound connectivity, refer to the following list of endpoints when configuring your firewall for outbound traffic.

## Data broker endpoints

The data broker contacts the following endpoints:

Endpoints	Purpose
olcentgbl.trafficmanager.net:443	To contact a repository for updating CentOS packages for the data broker host. This endpoint is contacted only if you manually install the data broker on a CentOS host.
rpm.nodesource.com:443 registry.npmjs.org:443 nodejs.org:443	To contact repositories for updating Node.js, npm, and other 3rd party packages used in development.
tgz.pm2.io:443	To access a repository for updating PM2, which is a 3rd party package used to monitor Cloud Sync.
sqs.us-east-1.amazonaws.com:443 kinesis.us-east-1.amazonaws.com:443	To contact the AWS services that Cloud Sync uses for operations (queuing files, registering actions, and delivering updates to the data broker).
s3.region.amazonaws.com:443  For example: s3.us-east-2.amazonaws.com:443 <a href="#">See AWS documentation for a list of S3 endpoints</a>	To contact Amazon S3 when a sync relationship includes an S3 bucket.
cf.cloudsync.netapp.com:443 repo.cloudsync.netapp.com:443	To contact the Cloud Sync service.
support.netapp.com:443	To contact NetApp support when using a BYOL license for sync relationships.
fedoraproject.org:443	To install 7z on the data broker virtual machine during installation and updates. 7z is needed to send AutoSupport messages to NetApp technical support.

## Web browser endpoints

Your web browser needs access to the following endpoint to download logs for troubleshooting purposes:

logs.cloudsync.netapp.com:443

## How to install a data broker

### Installing the data broker in AWS

When you create a sync relationship, choose the AWS Data Broker option to deploy the data broker software on a new EC2 instance in a VPC. Cloud Sync guides you through the installation process, but the requirements and steps are repeated on this page to help you prepare for installation.

You also have the option to install the data broker on an existing Linux host in the cloud or on your premises. [Learn more.](#)



## Supported AWS regions

All regions are supported except for the China and GovCloud (US) regions.

## Networking requirements

- The data broker needs an outbound internet connection so it can poll the Cloud Sync service for tasks over port 443.

When Cloud Sync deploys the data broker in AWS, it creates a security group that enables the required outbound communication. Note that you can configure the data broker to use a proxy server during the installation process.

If you need to limit outbound connectivity, see [the list of endpoints that the data broker contacts](#).

- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

## Permissions required to deploy the data broker in AWS

The AWS user account that you use to deploy the data broker must have the permissions included in [this NetApp-provided policy](#).

### Requirements to use your own IAM role with the AWS data broker

When Cloud Sync deploys the data broker, it creates an IAM role for the data broker instance. You can deploy the data broker using your own IAM role, if you prefer. You might use this option if your organization has strict security policies.

The IAM role must meet the following requirements:

- The EC2 service must be allowed to assume the IAM role as a trusted entity.
- [The permissions defined in this JSON file](#) must be attached to the IAM role so the data broker can function properly.

Follow the steps below to specify the IAM role when deploying the data broker.

## Installing the data broker


You can install a data broker in AWS when you create a sync relationship.

### Steps

1. Click **Create New Sync**.
2. On the **Define Sync Relationship** page, choose a source and target and click **Continue**.

Complete the steps until you reach the **Data Broker** page.

3. On the **Data Broker** page, click **Create Data Broker** and then select **Amazon Web Services**.

If you already have a data broker, you'll need to click the  icon first.

4. Enter a name for the data broker and click **Continue**.

5. Enter an AWS access key so Cloud Sync can create the data broker in AWS on your behalf.

The keys aren't saved or used for any other purposes.

If you'd rather not provide access keys, click the link at the bottom of the page to use a CloudFormation template instead. When you use this option, you don't need to provide credentials because you are logging in directly to AWS.

The following video shows how to launch the data broker instance using a CloudFormation template:

▶ [https://docs.netapp.com/us-en/occm38//media/video\\_cloud\\_sync.mp4](https://docs.netapp.com/us-en/occm38//media/video_cloud_sync.mp4) (video)

6. If you entered an AWS access key, select a location for the instance, select a key pair, choose whether to enable a public IP address, and then select an existing IAM role, or leave the field blank so Cloud Sync creates the role for you.

If you choose your own IAM role, you'll need to provide the required permissions.

**Basic Settings**

Location	Connectivity
Region US West   Oregon	Key Pair newKey
VPC vpc-3c46c059 - 10.60.21.0/25	Enable Public IP? <input checked="" type="radio"/> Enable <input type="radio"/> Disable
Subnet 10.60.21.0/25	IAM Role (optional) ⓘ

7. After the data broker is available, click **Continue** in Cloud Sync.

The following image shows a successfully deployed instance in AWS:

Select a NetApp Data Broker

aws	name	Active	
US West (Oregon) Region	10.60.21.0/25   vpc-3c46c059 VPC	10.60.21.5 Private IP	5f5002eecf378e000a560988 Broker ID
us-west-2c Availability Zone	10.60.21.0/25   subnet-e7f526be Subnet	i-0fc5c97e2f5f22c20 Instance ID	

8. Complete the pages in the wizard to create the new sync relationship.

## Result

You have deployed a data broker in AWS and created a new sync relationship. You can use this data broker with additional sync relationships.

## Installing the data broker in Azure

When you create a sync relationship, choose the Azure Data Broker option to deploy the data broker software on a new virtual machine in a VNet. Cloud Sync guides you through the installation process, but the requirements and steps are repeated on this page to help you prepare for installation.

You also have the option to install the data broker on an existing Linux host in the cloud or on your premises. [Learn more.](#)

## Supported Azure regions

All regions are supported except for the China, US Gov, and US DoD regions.

## Networking requirements

- The data broker needs an outbound internet connection so it can poll the Cloud Sync service for tasks over port 443.

When Cloud Sync deploys the data broker in Azure, it creates a security group that enables the required outbound communication.

If you need to limit outbound connectivity, see [the list of endpoints that the data broker contacts](#).

- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

## Authentication method

When you deploy the data broker, you'll need to choose an authentication method: a password or an SSH public-private key pair.

For help with creating a key pair, refer to [Azure Documentation: Create and use an SSH public-private key pair for Linux VMs in Azure](#).

## Installing the data broker


You can install a data broker in Azure when you create a sync relationship.


## Steps

1. Click **Create New Sync**.
2. On the **Define Sync Relationship** page, choose a source and target and click **Continue**.

Complete the pages until you reach the **Data Broker** page.

3. On the **Data Broker** page, click **Create Data Broker** and then select **Microsoft Azure**.

If you already have a data broker, you'll need to click the  icon first.



4. Enter a name for the data broker and click **Continue**.
5. If you're prompted, log in to your Microsoft account. If you're not prompted, click **Log in to Azure**.

The form is owned and hosted by Microsoft. Your credentials are not provided to NetApp.

6. Choose a location for the data broker and enter basic details about the virtual machine.

The screenshot shows two columns of configuration options. The left column is titled 'Location' and contains four dropdown menus: 'Subscription' (OCCM Dev), 'Azure Region' (West US 2), 'VNet' (Vnet1), and 'Subnet' (Subnet1). The right column is titled 'Virtual Machine' and contains: 'VM Name' (netappdatabroker), 'User Name' (databroker), 'Authentication Method' (radio buttons for Password and Public Key, with Password selected), 'Enter Password' (a masked field with asterisks), and 'Resource Group' (radio buttons for 'Generate a new group' and 'Use an existing group', with 'Generate a new group' selected).

7. Click **Continue** and keep the page open until the deployment is complete.

The process can take up to 7 minutes.

8. In Cloud Sync, click **Continue** once the data broker is available.
9. Complete the pages in the wizard to create the new sync relationship.

## Result

You have deployed a data broker in Azure and created a new sync relationship. You can use this data broker with additional sync relationships.

## Getting a message about needing admin consent?

If Microsoft notifies you that admin approval is required because Cloud Sync needs permission to access resources in your organization on your behalf, then you have two options:

1. Ask your AD admin to provide you with the following permission:

In Azure, go to **Admin Centers > Azure AD > Users and Groups > User Settings** and enable **Users can consent to apps accessing company data on their behalf**.

2. Ask your AD admin to consent on your behalf to **CloudSync-AzureDataBrokerCreator** using the following URL (this is the admin consent endpoint):

```
https://login.microsoftonline.com/{FILL HERE YOUR TENANT ID}/v2.0/adminconsent?client_id=8ee4ca3a-bafa-4831-97cc-5a38923cab85&redirect_uri=https://cloudsync.netapp.com&scope=https://management.azure.com/user_impersonationhttps://graph.microsoft.com/User.Read
```

As shown in the URL, our app URL is `https://cloudsync.netapp.com` and the application client ID is `8ee4ca3a-bafa-4831-97cc-5a38923cab85`.

## Installing the data broker in Google Cloud Platform

When you create a sync relationship, choose the GCP Data Broker option to deploy the data broker software on a new virtual machine instance in a VPC. Cloud Sync guides you through the installation process, but the requirements and steps are repeated on this page to help you prepare for installation.

You also have the option to install the data broker on an existing Linux host in the cloud or on your premises. [Learn more](#).

### Supported GCP regions

All regions are supported.

### Networking requirements

- The data broker needs an outbound internet connection so it can poll the Cloud Sync service for tasks over port 443.

When Cloud Sync deploys the data broker in GCP, it creates a security group that enables the required outbound communication.

If you need to limit outbound connectivity, see [the list of endpoints that the data broker contacts](#).

- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

## Permissions required to deploy the data broker in GCP

Ensure that the GCP user who deploys the data broker has the following permissions:

- `compute.networks.list`
- `compute.regions.list`
- `deploymentmanager.deployments.create`
- `deploymentmanager.deployments.delete`
- `deploymentmanager.operations.get`
- `iam.serviceAccounts.list`

## Permissions required for the service account

When you deploy the data broker, you need to select a service account that has the following permissions:

- `logging.logEntries.create`
- `resourcemanager.projects.get`
- `storage.buckets.get`
- `storage.buckets.list`
- `storage.objects.*`

## Installing the data broker


You can install a data broker in GCP when you create a sync relationship.

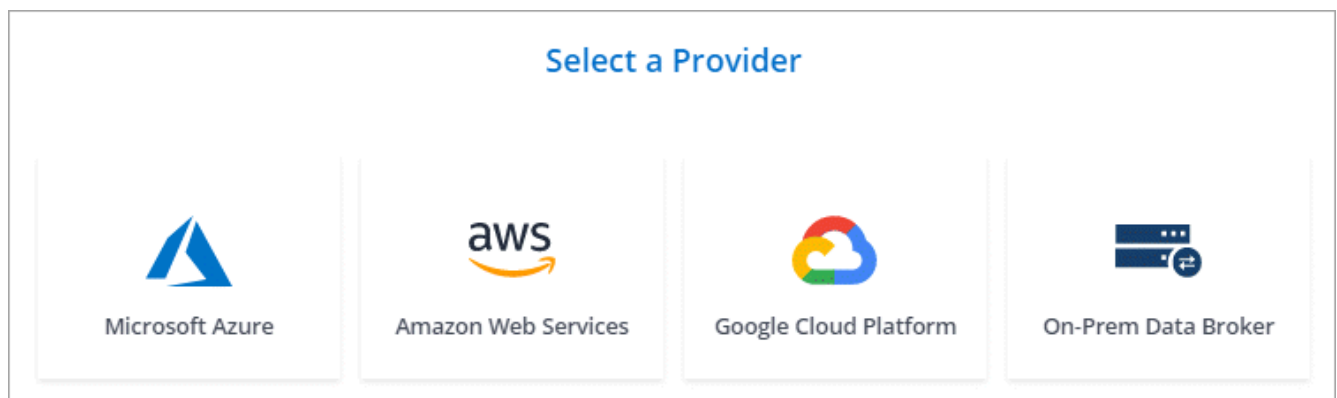
### Steps

1. Click **Create New Sync**.
2. On the **Define Sync Relationship** page, choose a source and target and click **Continue**.

Complete the steps until you reach the **Data Broker** page.

3. On the **Data Broker** page, click **Create Data Broker** and then select **Google Cloud Platform**.

If you already have a data broker, you'll need to click the  icon first.



4. Enter a name for the data broker and click **Continue**.
5. If you're prompted, log in with your Google account.

The form is owned and hosted by Google. Your credentials are not provided to NetApp.

6. Select a project and service account and then choose a location for the data broker.

### Basic Settings

<b>Project</b>	<b>Location</b>
Project <input type="text" value="OCCM-Dev"/>	Region <input type="text" value="us-west1"/>
Service Account <input type="text" value="test"/>	Zone <input type="text" value="us-west1-a"/>
Select a Service Account that includes <a href="#">these permissions</a>	VPC <input type="text" value="default"/>
	Subnet <input type="text" value="default"/>

7. Once the data broker is available, click **Continue** in Cloud Sync.

The instance takes approximately 5 to 10 minutes to deploy. You can monitor the progress from the Cloud Sync service, which automatically refreshes when the instance is available.

8. Complete the pages in the wizard to create the new sync relationship.

## Result

You've deployed a data broker in GCP and created a new sync relationship. You can use this data broker with additional sync relationships.

## Installing the data broker on a Linux host

When you create a sync relationship, choose the On-Prem Data Broker option to install the data broker software on an on-premises Linux host, or on an existing Linux host in the cloud. Cloud Sync guides you through the installation process, but the requirements and steps are repeated on this page to help you prepare for installation.

### Linux host requirements

- **Operating system:**
  - CentOS 7.0, 7.7, and 8.0
  - Red Hat Enterprise Linux 7.7 and 8.0

- Ubuntu Server 18.04 LTS
- SUSE Linux Enterprise Server 15 SP1

The command `yum update all` must be run on the host before you install the data broker.

A Red Hat Enterprise Linux system must be registered with Red Hat Subscription Management. If it is not registered, the system cannot access repositories to update required 3rd party software during installation.

- **RAM:** 16 GB
- **CPU:** 4 cores
- **Free disk space:** 10 GB
- **SELinux:** We recommend that you disable [SELinux](#) on the host.

SELinux enforces a policy that blocks data broker software updates and can block the data broker from contacting endpoints required for normal operation.

- **OpenSSL:** OpenSSL must be installed on the Linux host.

### Networking requirements

- The Linux host must have a connection to the source and target.
- The file server must allow the Linux host to access the exports.
- Port 443 must be open on the Linux host for outbound traffic to AWS (the data broker constantly communicates with the Amazon SQS service).
- NetApp recommends configuring the source, target, and data broker to use a Network Time Protocol (NTP) service. The time difference between the three components should not exceed 5 minutes.

### Enabling access to AWS

If you plan to use the data broker with a sync relationship that includes an S3 bucket, then you should prepare the Linux host for AWS access. When you install the data broker, you'll need to provide AWS keys for an AWS user that has programmatic access and specific permissions.

#### Steps

1. Create an IAM policy using [this NetApp-provided policy](#). [View AWS instructions](#).
2. Create an IAM user that has programmatic access. [View AWS instructions](#).

Be sure to copy the AWS keys because you need to specify them when you install the data broker software.

### Enabling access to Google Cloud

If you plan to use the data broker with a sync relationship that includes a Google Cloud Storage bucket, then you should prepare the Linux host for GCP access. When you install the data broker, you'll need to provide a key for a service account that has specific permissions.

#### Steps

1. Create a GCP service account that has Storage Admin permissions, if you don't already have one.
2. Create a service account key saved in JSON format. [View GCP instructions](#).



The file should contain at least the following properties: "project\_id", "private\_key", and "client\_email"



When you create a key, the file gets generated and downloaded to your machine.

3. Save the JSON file to the Linux host.

### Enabling access to Microsoft Azure

Access to Azure is defined per relationship by providing a storage account and a connection string in the Sync Relationship wizard.

### Installing the data broker


You can install a data broker on a Linux host when you create a sync relationship.


#### Steps

1. Click **Create New Sync**.
2. On the **Define Sync Relationship** page, choose a source and target and click **Continue**.

Complete the steps until you reach the **Data Broker** page.

3. On the **Data Broker** page, click **Create Data Broker** and then select **On-Prem Data Broker**.

If you already have a data broker, you'll need to click the  icon first.



Even though the option is labeled **On-Prem Data Broker**, it applies to a Linux host on your premises or in the cloud.

4. Enter a name for the data broker and click **Continue**.

The instructions page loads shortly. You'll need to follow these instructions—they include a unique link to download the installer.

5. On the instructions page:
  - a. Select whether to enable access to **AWS**, **Google Cloud**, or both.
  - b. Select an installation option: **No proxy**, **Use proxy server**, or **Use proxy server with authentication**.
  - c. Use the commands to download and install the data broker.

The following steps provide details about each possible installation option. Follow the instructions page to get the exact command based on your installation option.

d. Download the installer:

- No proxy:

```
curl <URI> -o data_broker_installer.sh
```

- Use proxy server:

```
curl <URI> -o data_broker_installer.sh -x <proxy_host>:<proxy_port>
```

- Use proxy server with authentication:

```
curl <URI> -o data_broker_installer.sh -x  
<proxy_username>:<proxy_password>@<proxy_host>:<proxy_port>
```

### URI

Cloud Sync displays the URI of the installation file on the instructions page, which loads when you follow the prompts to deploy the On-Prem Data Broker. That URI isn't repeated here because the link is generated dynamically and can be used only once. [Follow these steps to obtain the URI from Cloud Sync.](#)

e. Switch to superuser, make the installer executable and install the software:



Each command listed below includes parameters for AWS access and GCP access. Follow the instructions page to get the exact command based on your installation option.

- No proxy configuration:

```
sudo -s  
chmod +x data_broker_installer.sh  
./data_broker_installer.sh -a <aws_access_key> -s <aws_secret_key> -g  
<absolute_path_to_the_json_file>
```

- Proxy configuration:

```
sudo -s  
chmod +x data_broker_installer.sh  
./data_broker_installer.sh -a <aws_access_key> -s <aws_secret_key> -g  
<absolute_path_to_the_json_file> -h <proxy_host> -p <proxy_port>
```

- Proxy configuration with authentication:

```
sudo -s  
chmod +x data_broker_installer.sh  
./data_broker_installer.sh -a <aws_access_key> -s <aws_secret_key> -g  
<absolute_path_to_the_json_file> -h <proxy_host> -p <proxy_port> -u  
<proxy_username> -w <proxy_password>
```

### AWS keys

These are the keys for the user that you should have prepared [following these steps](#). The AWS keys are stored on the data broker, which runs in your on-premises or cloud network. NetApp

doesn't use the keys outside of the data broker.

### JSON file

This is the JSON file that contains a service account key that you should have prepared [following these steps](#).

6. Once the data broker is available, click **Continue** in Cloud Sync.
7. Complete the pages in the wizard to create the new sync relationship.

## Creating a sync relationship

When you create a sync relationship, the Cloud Sync service copies files from the source to the target. After the initial copy, the service syncs any changed data every 24 hours.

The steps below provide an example that shows how to set up a sync relationship from an NFS server to an S3 bucket.

### Steps

1. In Cloud Manager, click **Sync**.
2. On the **Define Sync Relationship** page, choose a source and target.

The following steps provide an example of how to create a sync relationship from an NFS server to an S3 bucket.



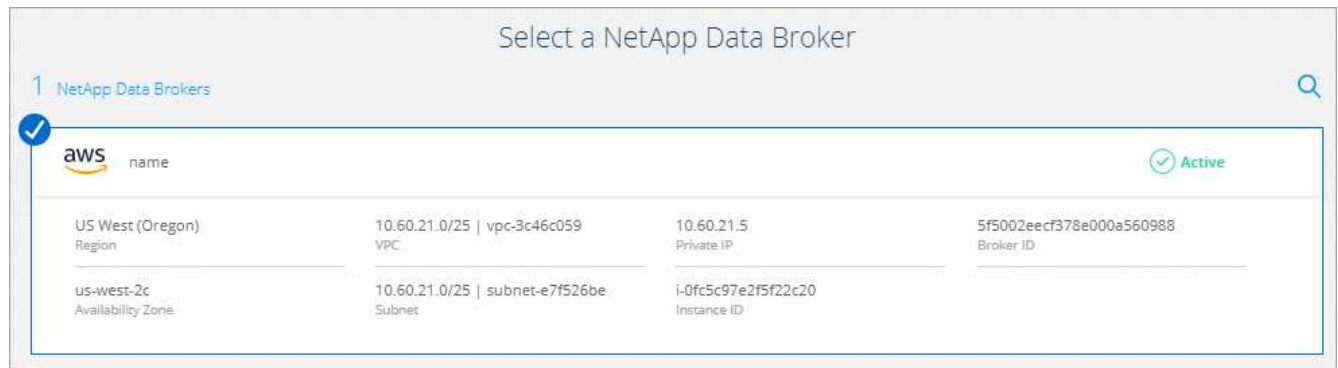
3. On the **NFS Server** page, enter the IP address or fully qualified domain name of the NFS server that you want to sync to AWS.
4. On the **Data Broker** page, follow the prompts to create a data broker virtual machine in AWS, Azure, or Google Cloud Platform, or to install the data broker software on an existing Linux host.

For more details, refer to the following pages:

- [Installing the data broker in AWS](#)
- [Installing the data broker in Azure](#)
- [Installing the data broker in GCP](#)
- [Installing the data broker on a Linux host](#)

5. After you install the data broker, click **Continue**.

The following image shows a successfully deployed data broker in AWS:



6. On the **Directories** page, select a top-level directory or subdirectory.

If Cloud Sync is unable to retrieve the exports, click **Add Export Manually** and enter the name of an NFS export.




If you want to sync more than one directory on the NFS server, then you must create additional sync relationships after you are done.

7. On the **AWS S3 Bucket** page, select a bucket:

- Drill down to select an existing folder within the bucket or to select a new folder that you create inside the bucket.
- Click **Add to the list** to select an S3 bucket that is not associated with your AWS account. [Specific permissions must be applied to the S3 bucket.](#)

8. On the **Bucket Setup** page, set up the bucket:

- Choose whether to enable S3 bucket encryption and then select an AWS KMS key, enter the ARN of a KMS key, or select AES-256 encryption.
- Select an S3 storage class. [View the supported storage classes.](#)



9. On the **Settings** page, define how source files and folders are synced and maintained in the target location:

### Schedule

Choose a recurring schedule for future syncs or turn off the sync schedule. You can schedule a relationship to sync data as often as every 1 minute.

### Retries

Define the number of times that Cloud Sync should retry to sync a file before skipping it.

### Recently Modified Files

Choose to exclude files that were recently modified prior to the scheduled sync.

### Delete Files on Source

Choose to delete files from the source location after Cloud Sync copies the files to the target location. This option includes the risk of data loss because the source files are deleted after they're copied.

If you enable this option, you also need to change a parameter in the `local.json` file on the data broker. Open the file and change the parameter named `workers.transferrer.delete-on-source` to **true**.

### Delete Files on Target

Choose to delete files from the target location, if they were deleted from the source. The default is to never delete files from the target location.

### Object tagging

When AWS S3 is the target in a sync relationship, Cloud Sync tags S3 objects with metadata that's relevant to the sync operation. You can disable tagging of S3 objects, if it's not desired in your environment. There's no impact to Cloud Sync if you disable tagging—Cloud Sync just stores the sync metadata in a different way.

### File Types

Define the file types to include in each sync: files, directories, and symbolic links.

### Exclude File Extensions

Specify file extensions to exclude from the sync by typing the file extension and pressing **Enter**. For example, type `log` or `.log` to exclude `*.log` files. A separator isn't required for multiple extensions. The following video provides a short demo:

► [https://docs.netapp.com/us-en/occm38//media/video\\_file\\_extensions.mp4](https://docs.netapp.com/us-en/occm38//media/video_file_extensions.mp4) (video)

### File Size

Choose to sync all files regardless of their size or just files that are in a specific size range.

### Date Modified

Choose all files regardless of their last modified date, files modified after a specific date, before a specific date, or between a time range.

10. On the **Relationship Tags** page, enter up to 9 relationship tags and then click **Continue**.

The Cloud Sync service assigns the tags to each object that it syncs to the S3 bucket.

11. Review the details of the sync relationship and then click **Create Relationship**.

## Result

Cloud Sync starts syncing data between the source and target.

## Paying for sync relationships after your free trial ends

There are two ways to pay for sync relationships after your 14-day free trial ends. The first option is to subscribe from AWS or Azure to pay-as-you-go or to pay annually. The second option is to purchase licenses directly from NetApp.

You can use licenses from NetApp with an AWS or Azure subscription. For example, if you have 25 sync relationships, you can pay for the first 20 sync relationships using a license and then pay-as-you-go from AWS or Azure with the remaining 5 sync relationships.

[Learn more about how licenses work.](#)

### What if I don't immediately pay after my free trial ends?

You won't be able to create any additional relationships. Existing relationships are not deleted, but you cannot make any changes to them until you subscribe or enter a license.

### Subscribing from AWS

AWS enables you to pay-as-you-go or to pay annually.

#### Steps to pay-as-you-go

1. Click **Sync > Licensing**.
2. Select **AWS**
3. Click **Subscribe** and then click **Continue**.
4. Subscribe from the AWS Marketplace, and then log back in to the Cloud Sync service to complete the registration.

The following video shows the process:

▶ [https://docs.netapp.com/us-en/occm38//media/video\\_cloud\\_sync\\_registering.mp4](https://docs.netapp.com/us-en/occm38//media/video_cloud_sync_registering.mp4) (video)

#### Steps to pay annually

1. [Go to the AWS Marketplace page.](#)
2. Click **Continue to Subscribe**.
3. Select your contract options and click **Create contract**.

### Subscribing from Azure

Azure enables you to pay-as-you-go or to pay annually.

#### What you'll need

An Azure user account that has Contributor or Owner permissions in the relevant subscription.

#### Steps

1. Click **Sync > Licensing**.
2. Select **Azure**.
3. Click **Subscribe** and then click **Continue**.
4. In the Azure portal, click **Create**, select your options, and click **Subscribe**.

Select **Monthly** to pay by the hour, or **Yearly** to pay for a year up front.

5. When deployment is complete, click the name of the SaaS resource in the notification pop-up.
6. Click **Configure Account** to return to Cloud Sync.

The following video shows the process:

▶ [https://docs.netapp.com/us-en/occm38//media/video\\_cloud\\_sync\\_registering\\_azure.mp4](https://docs.netapp.com/us-en/occm38//media/video_cloud_sync_registering_azure.mp4) (video)

### **Purchasing licenses from NetApp and adding them to Cloud Sync**

To pay for your sync relationships up front, you must purchase one or more licenses and add them to the Cloud Sync service.

#### **Steps**

1. Purchase a license by [contacting NetApp](#).
2. In Cloud Manager, click **Sync > Licensing**.
3. Click **Add License** and add the license.

## **Tutorials**

### **Copying ACLs between SMB shares**

Cloud Sync can copy access control lists (ACLs) between a source SMB share and a target SMB share. If needed, you can manually preserve the ACLs yourself by using robocopy.

#### **Choices**

- [Set up Cloud Sync to automatically copy ACLs](#)
- [Manually copy the ACLs yourself](#)

### **Setting up Cloud Sync to copy ACLs between SMB servers**

Copy ACLs between SMB servers by enabling a setting when you create a relationship or after you create a relationship.

Note that this feature is available for new sync relationships created after the 23 Feb 2020 release. If you'd like to use this feature with existing relationships created prior to that date, then you'll need to recreate the relationship.

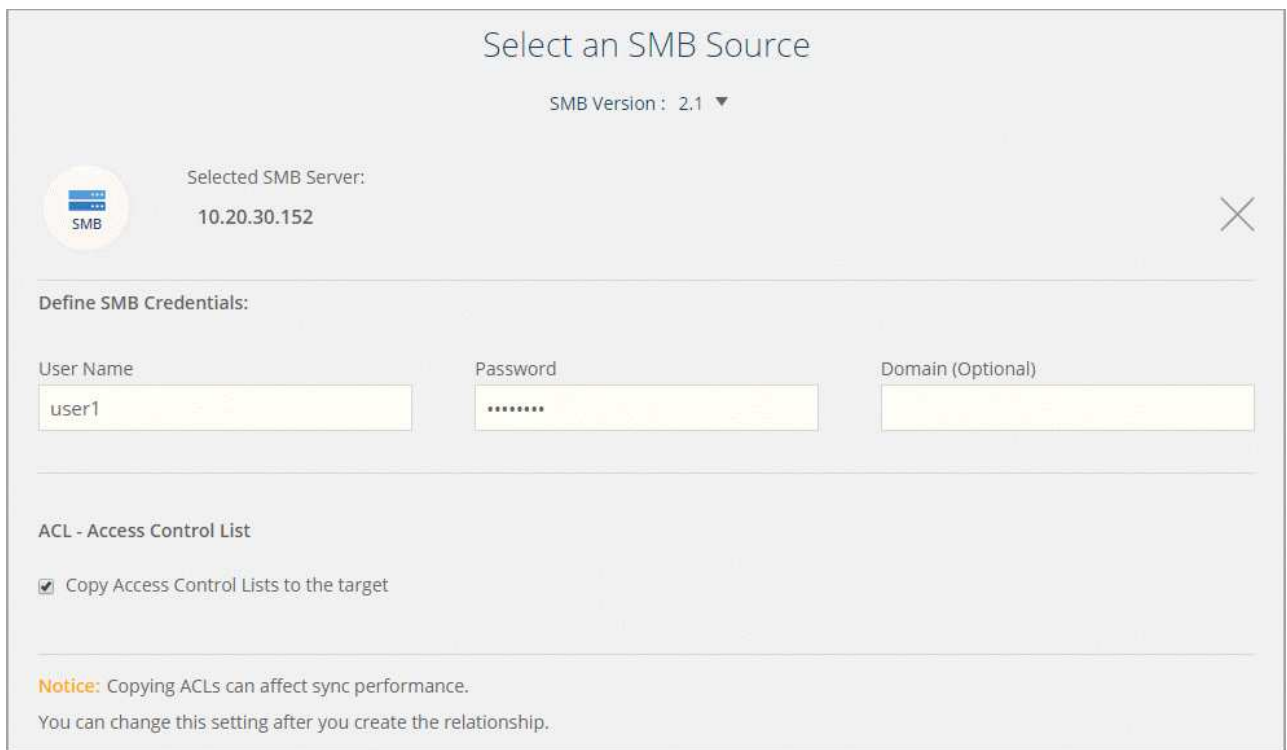
#### **What you'll need**

- A new sync relationship or an existing sync relationship created after the 23 Feb 2020 release.
- Any type of data broker.

This feature works with *any* type of data broker: the AWS, Azure, Google Cloud Platform, or on-prem data broker. The on-prem data broker can run [any supported operating system](#).

### Steps for a new relationship

1. From Cloud Sync, click **Create New Sync**.
2. Drag and drop **SMB Server** to the source and target and click **Continue**.
3. On the **SMB Server** page:
  - a. Enter a new SMB server or select an existing server and click **Continue**.
  - b. Enter credentials for the SMB server.
  - c. Select **Copy Access Control Lists to the target** and click **Continue**.



Select an SMB Source

SMB Version : 2.1 ▼

Selected SMB Server:  
10.20.30.152

Define SMB Credentials:

User Name: user1 Password: \*\*\*\*\* Domain (Optional):

ACL - Access Control List

Copy Access Control Lists to the target

**Notice:** Copying ACLs can affect sync performance.  
You can change this setting after you create the relationship.

4. Follow the remaining prompts to create the sync relationship.

### Steps for an existing relationship

1. Hover over the sync relationship and click the action menu.
2. Click **Settings**.
3. Select **Copy Access Control Lists to the target**.
4. Click **Save Settings**.

### Result

When syncing data, Cloud Sync preserves the ACLs between the source and target SMB shares.

### Manually copying ACLs

You can manually preserve ACLs between SMB shares by using the Windows robocopy command.

### Steps



1. Identify a Windows host that has full access to both SMB shares.
2. If either of the endpoints require authentication, use the **net use** command to connect to the endpoints from the Windows host.

You must perform this step before you use robocopy.

3. From Cloud Sync, create a new relationship between the source and target SMB shares or sync an existing relationship.
4. After the data sync is complete, run the following command from the Windows host to sync the ACLs and ownership:

```
robocopy /E /COPY:SOU /secfix [source] [target] /w:0 /r:0 /XD ~snapshots /UNILOG:"[logfilepath]
```

Both *source* and *target* should be specified using the UNC format. For example: \\<server>\<share>\<path>


## Syncing NFS data using data-in-flight encryption

If your business has strict security policies, you can sync NFS data using data-in-flight encryption. This feature is supported from an NFS server to another NFS server and from Azure NetApp Files to Azure NetApp Files.

For example, you might want to sync data between two NFS servers that are in different networks. Or you might need to securely transfer data on Azure NetApp Files across subnets or regions.

### How data-in-flight encryption works

Data-in-flight encryption encrypts NFS data when it's sent over the network between two data brokers. The following image shows a relationship between two NFS servers and two data brokers:



One data broker functions as the *initiator*. When it's time to sync data, it sends a connection request to the other data broker, which is the *listener*. That data broker listens for requests on port 443. You can use a different port, if needed, but be sure to check that the port is not in use by another service.

For example, if you sync data from an on-premises NFS server to a cloud-based NFS server, you can choose which data broker listens for the connection requests and which sends them.

Here's how in-flight encryption works:

1. After you create the sync relationship, the initiator starts an encrypted connection with the other data broker.
2. The source data broker encrypts data from the source using TLS 1.3.
3. It then sends the data over the network to the target data broker.
4. The target data broker decrypts the data before sending it to the target.
5. After the initial copy, the service syncs any changed data every 24 hours. If there is data to sync, the process starts with the initiator opening an encrypted connection with the other data broker.

If you prefer to sync data more frequently, [you can change the schedule after you create the relationship](#).

## Supported NFS versions

- For NFS servers, data-in-flight encryption is supported with NFS versions 3, 4.0, 4.1, and 4.2.
- For Azure NetApp Files, data-in-flight encryption is supported with NFS versions 3 and 4.1.

## What you'll need to get started

Be sure to have the following:

- Two NFS servers that meet [source and target requirements](#) or Azure NetApp Files in two subnets or regions.
- The IP addresses or fully qualified domain names of the servers.
- Network locations for two data brokers.

You can select an existing data broker but it must function as the initiator. The listener data broker must be a *new* data broker.

If you have not yet deployed a data broker, review the data broker requirements. Because you have strict security policies, be sure to review the networking requirements, which includes outbound traffic from port 443 and the [internet endpoints](#) that the data broker contacts.

- [Review AWS installation](#)
- [Review Azure installation](#)
- [Review GCP installation](#)
- [Review Linux host installation](#)

## Syncing NFS data using data-in-flight encryption

Create a new sync relationship between two NFS servers or between Azure NetApp Files, enable the in-flight encryption option, and follow the prompts.


### Steps

1. Click **Create New Sync**.
2. Drag and drop **NFS Server** to the source and target locations or **Azure NetApp Files** to the source and target locations and select **Yes** to enable data-in-flight encryption.

The following image shows what you'd select to sync data between two NFS servers:



The following image shows what you'd select to sync data between Azure NetApp Files:




3. Follow the prompts to create the relationship:

- a. **NFS Server/Azure NetApp Files:** Choose the NFS version and then specify a new NFS source or select an existing server.
- b. **Define Data Broker Functionality:** Define which data broker *listens* for connection requests on a port and which one *initiates* the connection. Make your choice based on your networking requirements.
- c. **Data Broker:** Follow the prompts to add a new source data broker or select an existing data broker.

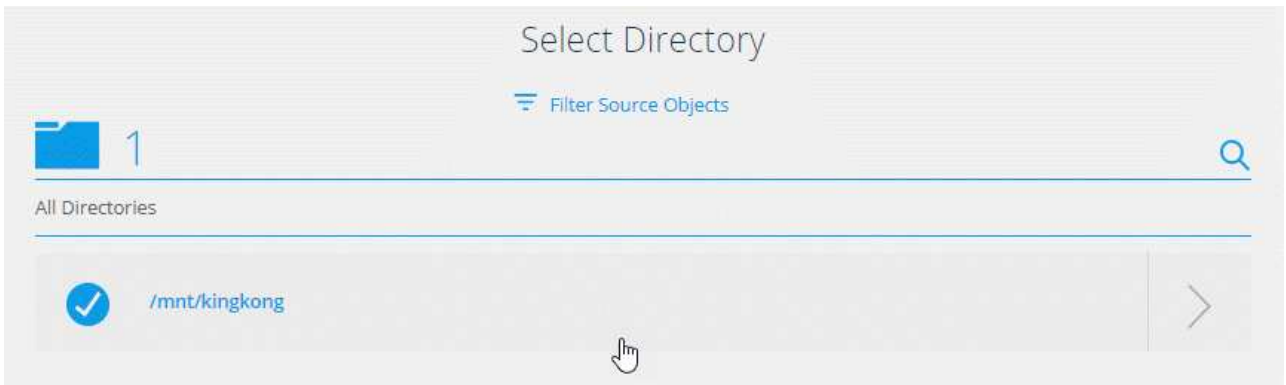
If the source data broker acts as the listener, then it must be a new data broker.

If you need a new data broker, Cloud Sync prompts you with the installation instructions. You can deploy the data broker in the cloud or download an installation script for your own Linux host.



- d. **Directories:** Choose the directories that you want to sync by selecting all directories, or by drilling down and selecting a subdirectory.


Click **Filter Source Objects** to modify settings that define how source files and folders are synced and maintained in the target location.




- e. **Target NFS Server/Target Azure NetApp Files:** Choose the NFS version and then enter a new NFS target or select an existing server.
- f. **Target Data Broker:** Follow the prompts to add a new source data broker or select an existing data broker.

If the target data broker acts as the listener, then it must be a new data broker.

Here's an example of the prompt when the target data broker functions as the listener. Notice the option to specify the port.



- g. **Target Directories:** Select a top-level directory, or drill down to select an existing subdirectory or to create a new folder inside an export.
- h. **Settings:** Define how source files and folders are synced and maintained in the target location.
- i. **Review:** Review the details of the sync relationship and then click **Create Relationship**.



## Result

Cloud Sync starts creating the new sync relationship. When it's done, click **View in Dashboard** to view details about the new relationship.

## Managing sync relationships

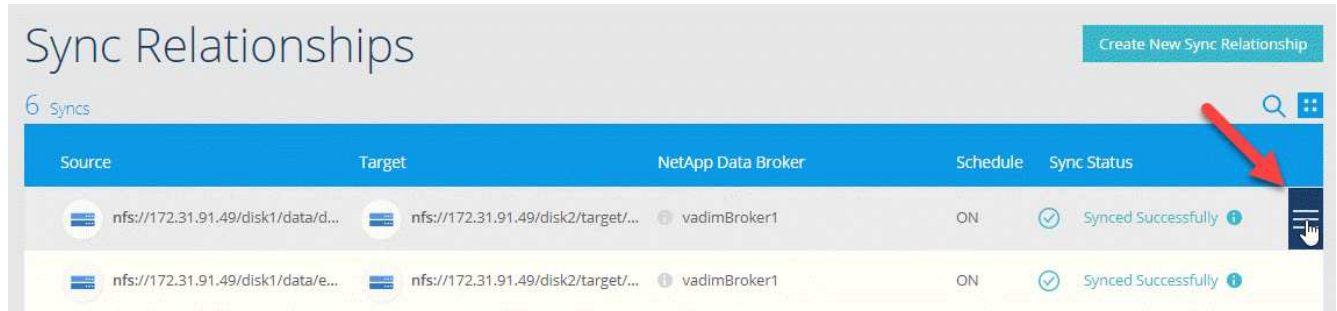
You can manage sync relationships at any time by immediately syncing data, changing schedules, and more.

## Performing an immediate data sync


Rather than wait for the next scheduled sync, you can press a button to immediately sync data between the source and target.

### Steps

1. From the **Sync Dashboard**, hover over the sync relationship and click the action menu.



2. Click **Sync Now** and then click **Sync** to confirm.



### Result

Cloud Sync starts the data sync process for the relationship.

## Accelerating sync performance

Accelerate the performance of a sync relationship by adding an additional data broker to the relationship. The additional data broker must be a *new* data broker.

### How this works

If the existing data brokers in the relationship are used in other sync relationships, then Cloud Sync automatically adds the new data broker to those relationships, as well.

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


- Relationship 1 uses data broker A
- Relationship 2 uses data broker B
- Relationship 3 uses data broker A

You want to accelerate the performance of relationship 1 so you add a new data broker to that relationship (data broker C). Because data broker A is also used in relationship 3, the new data broker is automatically added to relationship 3, as well.

### Steps

1. Ensure that at least one of the existing data brokers in the relationship are online.

2. Hover over the sync relationship and click the action menu.
3. Click **Accelerate**.



4. Follow the prompts to create a new data broker.

### Result

Cloud Sync adds the new data broker to the sync relationships. The performance of the next data sync should be accelerated.

## Changing the settings for a sync relationship

Modify settings that define how source files and folders are synced and maintained in the target location.

1. Hover over the sync relationship and click the action menu.
2. Click **Settings**.
3. Modify any of the settings.

**General**

Schedule	ON   Every 1 Day	▼
Retries	Retry 3 times before skipping file	▼

**Files and Directories**

Recently Modified Files	Exclude files that are modified up to 30 Seconds before a scheduled sync	▼
Delete Files On Source	Never delete files from the source location	▼
Delete Files On Target	Never delete files from the target location	▼
Object Tagging	Allow Cloud Sync to tag S3 objects	▼
File Types	Include All: Files, Directories, Symbolic Links	▼
Exclude File Extensions	None	▼
File Size	All	▼
Date Modified	All	▼

[Reset to defaults](#)

Here's a brief description of each setting:

### Schedule

Choose a recurring schedule for future syncs or turn off the sync schedule. You can schedule a relationship to sync data as often as every 1 minute.

### Retries

Define the number of times that Cloud Sync should retry to sync a file before skipping it.

### Recently Modified Files

Choose to exclude files that were recently modified prior to the scheduled sync.

### Delete Files on Source

Choose to delete files from the source location after Cloud Sync copies the files to the target location. This option includes the risk of data loss because the source files are deleted after they're copied.

If you enable this option, you also need to change a parameter in the local.json file on the data broker. Open the file and change the parameter named `workers.transferrer.delete-on-source` to **true**.



### Delete Files on Target

Choose to delete files from the target location, if they were deleted from the source. The default is to never delete files from the target location.

### Object tagging

When AWS S3 is the target in a sync relationship, Cloud Sync tags S3 objects with metadata that's relevant to the sync operation. You can disable tagging of S3 objects, if it's not desired in your environment. There's no impact to Cloud Sync if you disable tagging—Cloud Sync just stores the sync metadata in a different way.

### File Types

Define the file types to include in each sync: files, directories, and symbolic links.

### Exclude File Extensions

Specify file extensions to exclude from the sync by typing the file extension and pressing **Enter**. For example, type *log* or *.log* to exclude \*.log files. A separator isn't required for multiple extensions. The following video provides a short demo:

► [https://docs.netapp.com/us-en/occm38//media/video\\_file\\_extensions.mp4](https://docs.netapp.com/us-en/occm38//media/video_file_extensions.mp4) (video)

### File Size

Choose to sync all files regardless of their size or just files that are in a specific size range.

### Date Modified

Choose all files regardless of their last modified date, files modified after a specific date, before a specific date, or between a time range.

### Copy Access Control Lists to the target

Choose to copy access control lists (ACLs) between source SMB shares and target SMB shares. Note that this option is only available for sync relationships created after the 23 Feb 2020 release.

4. Click **Save Settings**.

### Result

Cloud Sync modifies the sync relationship with the new settings.

## Deleting relationships

You can delete a sync relationship, if you no longer need to sync data between the source and target. This action does not delete the data broker instance and it does not delete data from the target.

### Steps

1. Hover over the sync relationship and click the action menu.
2. Click **Delete** and then click **Delete** again to confirm.

### Result

Cloud Sync deletes the sync relationship.

## Cloud Sync APIs

The Cloud Sync capabilities that are available through the web UI are also available

through RESTful APIs.

## Getting started

To get started with the Cloud Sync APIs, you need to obtain a user token and your Cloud Central account ID. You'll need to add the token and account ID to the Authorization header when making API calls.

### Steps

1. Obtain a user token from NetApp Cloud Central.

```
POST https://netapp-cloud-account.auth0.com/oauth/token
Header: Content-Type: application/json
Body:
{
  "username": "<user_email>",
  "scope": "profile",
  "audience": "https://api.cloud.netapp.com",
  "client_id": "UaVhOIXMWQs5ilWdDxauXe5Mqkb34NJQ",
  "grant_type": "password",
  "password": "<user_password>"
}
```

2. Obtain your Cloud Central account ID.

```
GET https://cloudsync.netapp.com/api/accounts
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
```

This API will return a response like the following:

```
[
  {
    "accountId": "account-JeL97Ry3",
    "name": "Test"
  }
]
```

3. Add the user token and account ID in the Authorization header of each API call.

### Example

The following example shows an API call to create a data broker in Microsoft Azure. You would simply replace <user\_token> and <accountId> with the token and ID that you obtained in the previous steps.

```
POST https://cloudsync.netapp.com/api/data-brokers
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
x-account-id: <accountId>
Body: { "name": "databroker1", "type": "AZURE" }
```

## What should I do when the token expires?

The user token from NetApp Cloud Central has an expiration date. To refresh the token, you need to call the API from step 1 again.

The API response includes an "expires\_in" field that states when the token expires.

## API reference

Documentation for each Cloud Sync API is available from [NetApp Cloud Central](#).

## Using list APIs

List APIs are asynchronous APIs, so the result does not return immediately (for example: `GET /data-brokers/{id}/list-nfs-export-folders` and `GET /data-brokers/{id}/list-s3-buckets`). The only response from the server is HTTP status 202. To get the actual result, you must use the `GET /messages/client` API.

### Steps

1. Call the list API that you want to use.
2. Use the `GET /messages/client` API to view the result of the operation.
3. Use the same API by appending it with the ID that you just received: `GET http://cloudsync.netapp.com/api/messages/client?last=<id_from_step_2>`

Note that the ID changes each time that you call the `GET /messages/client` API.

### Example

When you call the `list-s3-buckets` API, a result is not immediately returned:

```
GET http://cloudsync.netapp.com/api/data-brokers/<data-broker-id>/list-s3-
buckets
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
x-account-id: <accountId>
```

The result is HTTP status code 202, which means the message was accepted, but was not processed yet.

To get the result of the operation, you need to use the following API:

```
GET http://cloudsync.netapp.com/api/messages/client
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
x-account-id: <accountId>
```

The result is an array with one object that includes an ID field. The ID field represents the last message that the server sent. For example:

```
[
  {
    "header": {
      "requestId": "init",
      "clientId": "init",
      "agentId": "init"
    },
    "payload": {
      "init": {}
    },
    "id": "5801"
  }
]
```

You would now make the following API call using the ID that you just received:

```
GET http://cloudsync.netapp.com/api/messages/client?last=<id_from_step_2>
Headers: Authorization: Bearer <user_token>
Content-Type: application/json
x-account-id: <accountId>
```

The result is an array of messages. Inside each message is a payload object, which consists of the name of the operation (as key) and its result (as value). For example:

```
[
  {
    "payload": {
      "list-s3-buckets": [
        {
          "tags": [
            {
              "Value": "100$",
              "Key": "price"
            }
          ],
          "region": {
            "displayName": "US West (Oregon)",
            "name": "us-west-2"
          },
          "name": "small"
        }
      ]
    },
    "header": {
      "requestId": "f687ac55-2f0c-40e3-9fa6-57fb8c4094a3",
      "clientId": "5beb032f548e6e35f4ed1ba9",
      "agentId": "5bed61f4489fb04e34a9aac6"
    },
    "id": "5802"
  }
]
```

## Cloud Sync technical FAQ

This FAQ can help if you're just looking for a quick answer to a question.

### Getting started

The following questions relate to getting started with Cloud sync.

#### How does Cloud Sync work?

Cloud Sync uses the NetApp data broker software to sync data from a source to a target (this is called a *sync relationship*).

The data broker controls the sync relationships between your sources and targets. After you set up a sync relationship, Cloud Sync analyzes your source system and breaks it up into multiple replication streams to push to your selected target data.

After the initial copy, the service syncs any changed data based on the schedule that you set.

## How does the 14-day free trial work?

The 14-day free trial starts when you sign up for the Cloud Sync service. You're not subject to NetApp charges for Cloud Sync relationships you create for 14 days. However, all resource charges for any data broker that you deploy still applies.

## How much does Cloud Sync cost?

There are two types of costs associated with using Cloud Sync: service charges and resource charges.

### Service charges

For pay-as-you-go pricing, Cloud Sync service charges are hourly, based on the number of sync relationships that you create.

- [View pay-as-you-go pricing in AWS](#)
- [View annual pricing in AWS](#)
- [View pricing in Azure](#)

Cloud Sync licenses are also available through your NetApp representative. Each license enables 20 sync relationships for 12 months.

[Learn more about licenses.](#)

### Resource charges

The resource charges are related to the compute and storage costs for running the data broker in the cloud.

## How is Cloud Sync billed?

There are two ways to pay for sync relationships after your 14-day free trial ends. The first option is to subscribe from AWS or Azure, which enables you to pay-as-you-go or to pay annually. The second option is to purchase licenses directly from NetApp.

## Can I use Cloud Sync outside the cloud?

Yes, you can use Cloud Sync in a non-cloud architecture. The source and target can reside on-premises and so can the data broker.

Note the following key points about using Cloud Sync outside of the cloud:

- For on-premises synchronization, a private Amazon S3 bucket is available through NetApp StorageGRID.
- The data broker does need an internet connection to communicate with the Cloud Sync service.
- If you don't purchase a license directly from NetApp, you will need an AWS or Azure account for the PAYGO Cloud Sync service billing.

## How do I access Cloud Sync?

Cloud Sync is available from Cloud Manager in the **Sync** tab.

## Supported sources and targets

The following questions related to the source and targets that are supported in a sync relationship.

### **Which sources and targets does Cloud Sync support?**

Cloud Sync supports many different types of sync relationships. [View the entire list.](#)

### **What versions of NFS and SMB does Cloud Sync support?**

Cloud Sync supports NFS version 3 and later, and SMB version 1 and later.

[Learn more about sync requirements.](#)

### **When Amazon S3 is the target, can the data be tiered to a specific S3 storage class?**

Yes, you can choose a specific S3 storage class when AWS S3 is the target:

- Standard (this is the default class)
- Intelligent-Tiering
- Standard-Infrequent Access
- One Zone-Infrequent Access
- Glacier
- Glacier Deep Archive

### **What about storage tiers for Azure Blob storage?**

You can choose a specific Azure Blob storage tier when a Blob container is the target:

- Hot storage
- Cool storage

## **Networking**

The following questions relate to networking requirements for Cloud Sync.

### **What are the networking requirements for Cloud Sync?**

The Cloud Sync environment requires that the data broker is connected with the source and the target through the selected protocol (NFS, SMB, EFS) or object storage API (Amazon S3, Azure Blob, IBM Cloud Object Storage).

In addition, the data broker needs an outbound internet connection over port 443 so it can communicate with the Cloud Sync service and contact a few other services and repositories.

For more details, [review networking requirements.](#)

### **Are there networking limitations related to data broker connectivity?**

Data brokers require internet access. We don't support a proxy server when deploying the data broker in Azure or in Google Cloud Platform.

## **Data synchronization**

The following questions relate to how data synchronization works.

## How often does synchronization occur?

The default schedule is set for daily synchronization. After the initial synchronization, you can:

- Modify the sync schedule to your desired number of days, hours, or minutes
- Disable the sync schedule
- Delete the sync schedule (no data will be lost; only the sync relationship will be removed)

## What is the minimum sync schedule?

You can schedule a relationship to sync data as often as every 1 minute.

## Does the data broker retry when a file fails to sync? Or does it timeout?

The data broker doesn't timeout when a single file fails to transfer. Instead, the data broker retries 3 times before skipping the file. The retry value is configurable in the settings for a sync relationship.

[Learn how to change the settings for a sync relationship.](#)

## What if I have a very large dataset?

If a single directory contains 600,000 files or more, [contact us](#) so we can help you configure the data broker to handle the payload. We might need to add additional memory to the data broker machine.

## Security

The following questions related to security.

### Is Cloud Sync secure?

Yes. All Cloud Sync service networking connectivity is done using [Amazon Simple Queue Service \(SQS\)](#).

All communication between the data broker and Amazon S3, Azure Blob, Google Cloud Storage, and IBM Cloud Object Storage is done through the HTTPS protocol.

If you're using Cloud Sync with on-premises (source or destination) systems, here's a few recommended connectivity options:

- An AWS Direct Connect, Azure ExpressRoute, or Google Cloud Interconnect connection, which is non-internet routed (and can only communicate with the cloud networks that you specify)
- A VPN connection between your on-premises gateway device and your cloud networks
- For extra secure data transfer with S3 buckets, Azure Blob storage, or Google Cloud Storage, an Amazon Private S3 Endpoint, Azure Virtual Network service endpoints, or Private Google Access may be established.

Any of these methods establishes a secure connection between your on-premises NAS servers and a Cloud Sync data broker.

### Is data encrypted by Cloud Sync?

- Cloud Sync supports data-in-flight encryption between source and target NFS servers. [Learn more.](#)
- Encryption is not supported with SMB.



- When an Amazon S3 bucket is the target in a sync relationship, you can choose whether to enable data encryption using AWS KMS encryption or AES-256 encryption.

## Permissions

The following questions relate to data permissions.

### **Are SMB data permissions synced to the target location?**

You can set up Cloud Sync to preserve access control lists (ACLs) between a source SMB share and a target SMB share. Or you can manually copy the ACLs yourself. [Learn how to copy ACLs between SMB shares.](#)

### **Are NFS data permissions synced to the target location?**

Cloud Sync automatically copies NFS permissions between NFS servers as follows:

- NFS version 3: Cloud Sync copies the permissions and the user group owner.
- NFS version 4: Cloud Sync copies the ACLs.

## Performance

The following questions relate to Cloud Sync performance.

### **What does the progress indicator for a sync relationship represent?**

The sync relationship shows the throughput of the data broker's network adapter. If you accelerated sync performance by using multiple data brokers, then the throughput is the sum of all traffic. This throughput refreshes every 20 seconds.

### **I'm experiencing performance issues. Can we limit the number of concurrent transfers?**

The data broker can sync 4 files at a time. If you have very large files (multiple TBs each), it can take a long time to complete the transfer process and performance might be impacted.

Limiting the number of concurrent transfers can help. [Contact us for help.](#)

### **Why am I experiencing low performance with Azure NetApp Files?**

When you sync data to or from Azure NetApp Files, you might experience failures and performance issues if the disk service level is Standard.

Change the service level to Premium or Ultra to enhance the sync performance.

[Learn more about Azure NetApp Files service levels and throughput.](#)

### **Why am I experiencing low performance with Cloud Volumes Service for AWS?**

When you sync data to or from a cloud volume, you might experience failures and performance issues if the level of performance for the cloud volume is Standard.

Change the Service level to Premium or Extreme to enhance the sync performance.

## How many data brokers are required?

When you create a new relationship, you start with a single data broker (unless you selected an existing data broker that belongs to an accelerated sync relationship). In many cases, a single data broker can meet the performance requirements for a sync relationship. If it doesn't, you can accelerate sync performance by adding additional data brokers. But you should first check other factors that can impact sync performance.

Multiple factors can impact data transfer performance. The overall sync performance might be impacted due to network bandwidth, latency, and network topology, as well as the data broker VM specs and storage system performance. For example, a single data broker in a sync relationship can reach 100 MB/s, while disk throughput on the target might only allow 64 MB/s. As a result, the data broker keeps trying to copy the data, but the target can't meet the performance of the data broker.

So be sure to check the performance of your networking and the disk throughput on the target.

Then you can consider accelerating sync performance by adding an additional data broker to share the load of that relationship. [Learn how to accelerate sync performance.](#)

## Deleting things

The following questions relate to deleting sync relationships and data from sources and targets.

### What happens if I delete my Cloud Sync relationship?

Deleting a relationship stops all future data syncs and terminates payment. Any data that was synced to the target remains as-is.

### What happens if I delete something from my source server? Is it removed from the target too?

By default, if you have an active sync relationship, the item deleted on the source server is not deleted from the target during the next synchronization. But there is an option in the sync settings for each relationship, where you can define that Cloud Sync will delete files in the target location if they were deleted from the source.

[Learn how to change the settings for a sync relationship.](#)

### What happens if I delete something from my target? Is it removed from my source too?

If an item is deleted from the target, it will not be removed from the source. The relationship is one-way—from source to target. On the next sync cycle, Cloud Sync compares the source to the target, identifies that the item is missing, and Cloud Sync copies it again from the source to the target.

## Troubleshooting

[NetApp Knowledgebase: Cloud Sync FAQ: Support and Troubleshooting](#)

## Data broker deep dive

The following question relates to the data broker.

### Can you explain the architecture of the data broker?

Sure. Here are the most important points:

- The data broker is a node.js application running on a Linux host.

- Cloud Sync deploys the data broker as follows:
  - AWS: From an AWS CloudFormation template
  - Azure: From Azure Resource Manager
  - Google: From Google Cloud Deployment Manager
  - If you use your own Linux host, you need to manually install the software
- The data broker software automatically upgrades itself to the latest version.
- The data broker uses AWS SQS as a reliable and secure communication channel and for control and monitoring. SQS also provides a persistency layer.
- You can add additional data brokers to a relationship to increase transfer speed and add high availability. There is service resiliency if one data broker fails.

# Gain insight into data privacy

## Learn about Cloud Compliance

Cloud Compliance is a data privacy and compliance service for Cloud Manager that scans your volumes, Amazon S3 buckets, and databases to identify the personal and sensitive data that resides in those files. Using Artificial Intelligence (AI) driven technology, Cloud Compliance helps organizations understand data context and identify sensitive data.

[Learn about the use cases for Cloud Compliance.](#)

### Features

Cloud Compliance provides several tools that can help you with your compliance efforts. You can use Cloud Compliance to:

- Identify Personal Identifiable Information (PII)
- Identify a wide scope of sensitive information as required by GDPR, CCPA, PCI, and HIPAA privacy regulations
- Respond to Data Subject Access Requests (DSAR)

### Supported working environments and data sources

Cloud Compliance can scan data from the following types of data sources:

- Cloud Volumes ONTAP in AWS
- Cloud Volumes ONTAP in Azure
- Azure NetApp Files
- Amazon S3
- Databases that reside anywhere (there is no requirement that the database resides in a working environment)

**Note:** For Azure NetApp Files, Cloud Compliance can scan any volumes that are in the same region as Cloud Manager.

### Cost

- The cost to use Cloud Compliance depends on the amount of data that you're scanning. As of October 7th, 2020, the first 1 TB of data that Cloud Compliance scans in a Cloud Manager workspace is free. This includes data from Cloud Volumes ONTAP volumes, Azure NetApp Files volumes, Amazon S3 buckets, and database schemas. A subscription to the AWS or Azure Marketplace is required to continue scanning data after that point. See [pricing](#) for details.

[Learn how to subscribe.](#)

- Installing Cloud Compliance requires deploying a cloud instance, which results in charges from the cloud provider where it is deployed. See the [the type of instance that is deployed for each cloud provider](#)

- Cloud Compliance requires that you have deployed a Connector. In many cases you already have a Connector because of other storage and services you are using in Cloud Manager. The Connector instance results in charges from the cloud provider where it is deployed. See the [type of instance that is deployed for each cloud provider](#).

### Data transfer costs

Data transfer costs depend on your setup. If the Cloud Compliance instance and data source are in the same Availability Zone and region, then there are no data transfer costs. But if the data source, such as a Cloud Volumes ONTAP cluster or S3 Bucket, is in a *different* Availability Zone or region, then you'll be charged by your cloud provider for data transfer costs. See these links for more details:

- [AWS: Amazon EC2 Pricing](#)
- [Microsoft Azure: Bandwidth Pricing Details](#)

### How Cloud Compliance works

At a high-level, Cloud Compliance works like this:

1. You deploy an instance of Cloud Compliance in Cloud Manager.
2. You enable it on one or more working environments, or on your databases.
3. Cloud Compliance scans the data using an AI learning process.
4. In Cloud Manager, you click **Compliance** and use the provided dashboard and reporting tools to help in your compliance efforts.

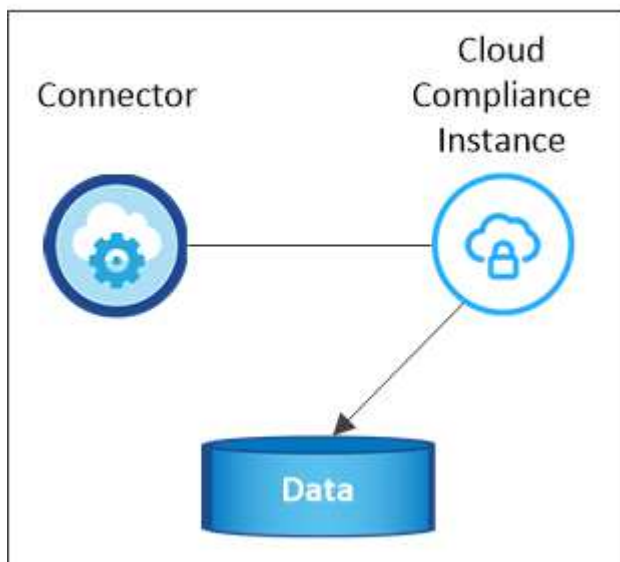
### The Cloud Compliance instance

When you enable Cloud Compliance, Cloud Manager deploys a Cloud Compliance instance in the same subnet as the Connector. [Learn more about Connectors](#).



If the Connector is installed on-prem, it deploys the Cloud Compliance instance in same VPC or VNet as the first Cloud Volumes ONTAP system in the request.

#### VPC or VNet



Note the following about the instance:

- In Azure, Cloud Compliance runs on a Standard\_D16s\_v3 VM with a 512 GB disk.
- In AWS, Cloud Compliance runs on an m5.4xlarge instance with a 500 GB GP2 disk.

In regions where m5.4xlarge isn't available, Cloud Compliance runs on an m4.4xlarge instance instead.



Changing or resizing the instance/VM type isn't supported. You need to use the size that's provided.

- The instance is named *CloudCompliance* with a generated hash (UUID) concatenated to it. For example: *CloudCompliance-16bb6564-38ad-4080-9a92-36f5fd2f71c7*
- Only one Cloud Compliance instance is deployed per Connector.
- Upgrades of Cloud Compliance software is automated—you don't need to worry about it.



The instance should remain running at all times because Cloud Compliance continuously scans the data.

## How scans work

After you enable Cloud Compliance and select the volumes, buckets, or database schemas you want to scan, it immediately starts scanning the data to identify personal and sensitive data. It maps your organizational data, categorizes each file, and identifies and extracts entities and predefined patterns in the data. The result of the scan is an index of personal information, sensitive personal information, and data categories.

Cloud Compliance connects to the data like any other client by mounting NFS and CIFS volumes. NFS volumes are automatically accessed as read-only, while you need to provide Active Directory credentials to scan CIFS volumes.

After the initial scan, Cloud Compliance continuously scans each volume to detect incremental changes (this is why it's important to keep the instance running).

You can enable and disable scans at the [volume level](#), at the [bucket level](#), and at the [database schema level](#).

## Information that Cloud Compliance indexes

Cloud Compliance collects, indexes, and assigns categories to unstructured data (files). The data that Cloud Compliance indexes includes the following:

### Standard metadata

Cloud Compliance collects standard metadata about files: the file type, its size, creation and modification dates, and so on.

### Personal data

Personally identifiable information such as email addresses, identification numbers, or credit card numbers. [Learn more about personal data.](#)

### Sensitive personal data

Special types of sensitive information, such as health data, ethnic origin, or political opinions, as defined by

GDPR and other privacy regulations. [Learn more about sensitive personal data.](#)

## Categories

Cloud Compliance takes the data that it scanned and divides it into different types of categories. Categories are topics based on AI analysis of the content and metadata of each file. [Learn more about categories.](#)

## Name entity recognition

Cloud Compliance uses AI to extract natural persons' names from documents. [Learn about responding to Data Subject Access Requests.](#)

## Networking overview

Cloud Manager deploys the Cloud Compliance instance with a security group that enables inbound HTTP connections from the Connector instance.

When using Cloud Manager in SaaS mode, the connection to Cloud Manager is served over HTTPS, and the private data sent between your browser and the Cloud Compliance instance are secured with end-to-end encryption, which means NetApp and third parties can't read it.

If you need to use the local user interface instead of the SaaS user interface for any reason, you can still [access the local UI.](#)

Outbound rules are completely open. Internet access is needed to install and upgrade the Cloud Compliance software and to send usage metrics.

If you have strict networking requirements, [learn about the endpoints that Cloud Compliance contacts.](#)

## User access to compliance information

The role each user has been assigned provides different capabilities within Cloud Manager and within Cloud Compliance:

- **Account Admins** can manage compliance settings and view compliance information for all working environments.
- **Workspace Admins** can manage compliance settings and view compliance information only for systems that they have permissions to access. If a Workspace Admin can't access a working environment in Cloud Manager, then they can't see any compliance information for the working environment in the Compliance tab.
- Users with the **Cloud Compliance Viewer** role can only view compliance information and generate reports for systems that they have permission to access. These users cannot enable/disable scanning of volumes, buckets, or database schemas.

[Learn more about Cloud Manager roles](#) and how to [add users with specific roles.](#)

## Get started

### Deploy Cloud Compliance

Complete a few steps to deploy the Cloud Compliance instance in your Cloud Manager workspace.

## Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



### Create a Connector

If you don't already have a Connector, create a Connector in Azure or AWS. See [creating a Connector in AWS](#) or [creating a Connector in Azure](#).



### Review prerequisites

Ensure that your cloud environment can meet the prerequisites, which includes 16 vCPUs for the Cloud Compliance instance, outbound internet access for the instance, connectivity between the Connector and Cloud Compliance over port 80, and more. [See the complete list](#).



### Deploy Cloud Compliance

Launch the installation wizard to deploy the Cloud Compliance instance in Cloud Manager.



### Subscribe to the Cloud Compliance service

The first 1 TB of data that Cloud Compliance scans in Cloud Manager is free. A subscription to the AWS or Azure Marketplace is required to continue scanning data after that point.

## Creating a Connector

If you don't already have a Connector, create a Connector in Azure or AWS. See [creating a Connector in AWS](#) or [creating a Connector in Azure](#). In most cases you will probably have a Connector set up before you attempt to activate Cloud Compliance because most [Cloud Manager features require a Connector](#), but there are cases when you need to set one up now.

There are some scenarios where you have to use a Connector in AWS or Azure for Cloud Compliance.

- When scanning data in Cloud Volumes ONTAP in AWS or in AWS S3 buckets, you use a connector in AWS.
- When scanning data in Cloud Volumes ONTAP in Azure or in Azure NetApp Files, you use a connector in Azure.
- Databases can be scanned using either Connector.

As you can see, there may be some situations where you need to use [multiple Connectors](#).



If you are planning on scanning Azure NetApp Files, you need to make sure you're deploying in the same region as the volumes you wish to scan.



## Reviewing prerequisites

Review the following prerequisites to make sure that you have a supported configuration before you deploy Cloud Compliance.

### Enable outbound internet access

Cloud Compliance requires outbound internet access. If your virtual network uses a proxy server for internet access, ensure that the Cloud Compliance instance has outbound internet access to contact the following endpoints. Note that Cloud Manager deploys the Cloud Compliance instance in the same subnet as the Connector.

Endpoints	Purpose
<a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a>	Communication with the Cloud Manager service, which includes Cloud Central accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://auth0.com">https://auth0.com</a>	Communication with NetApp Cloud Central for centralized user authentication.
<a href="https://cloud-compliance-support-netapp.s3.us-west-2.amazonaws.com">https://cloud-compliance-support-netapp.s3.us-west-2.amazonaws.com</a> <a href="https://hub.docker.com">https://hub.docker.com</a> <a href="https://auth.docker.io">https://auth.docker.io</a> <a href="https://registry-1.docker.io">https://registry-1.docker.io</a> <a href="https://index.docker.io/">https://index.docker.io/</a> <a href="https://dseasb33srrrn.cloudfront.net/">https://dseasb33srrrn.cloudfront.net/</a> <a href="https://production.cloudflare.docker.com/">https://production.cloudflare.docker.com/</a>	Provides access to software images, manifests, and templates.
<a href="https://kinesis.us-east-1.amazonaws.com">https://kinesis.us-east-1.amazonaws.com</a>	Enables NetApp to stream data from audit records.
<a href="https://cognito-idp.us-east-1.amazonaws.com">https://cognito-idp.us-east-1.amazonaws.com</a> <a href="https://cognito-identity.us-east-1.amazonaws.com">https://cognito-identity.us-east-1.amazonaws.com</a>	Enables Cloud Compliance to access and download manifests and templates, and to send logs and metrics.

### Ensure that Cloud Manager has the required permissions

Ensure that Cloud Manager has permissions to deploy resources and create security groups for the Cloud Compliance instance. You can find the latest Cloud Manager permissions in [the policies provided by NetApp](#).

### Check your vCPU limits

Ensure that your cloud provider's vCPU limit allows for the deployment of an instance with 16 cores. You'll need to verify the vCPU limit for the relevant instance family in the region where Cloud Manager is running.

In AWS, the instance family is *On-Demand Standard instances*. In Azure, the instance family is *Standard DSv3 Family*.

For more details on vCPU limits, see the following:

- [AWS documentation: Amazon EC2 Service Limits](#)
- [Azure documentation: Virtual machine vCPU quotas](#)

### Ensure that Cloud Manager can access Cloud Compliance

Ensure connectivity between the Connector and the Cloud Compliance instance. The security group for the Connector must allow inbound and outbound traffic over port 80 to and from the Cloud Compliance

instance.

This connection enables deployment of the Cloud Compliance instance and enables you to view information in the Compliance tab.

### Set up discovery of Azure NetApp Files

Before you can scan volumes for Azure NetApp Files, [Cloud Manager must be set up to discover the configuration](#).

### Ensure that you can keep Cloud Compliance running

The Cloud Compliance instance needs to stay on to continuously scan your data.

### Ensure web browser connectivity to Cloud Compliance

After Cloud Compliance is enabled, ensure that users access the Cloud Manager interface from a host that has a connection to the Cloud Compliance instance.


The Cloud Compliance instance uses a private IP address to ensure that the indexed data isn't accessible to the internet. As a result, the web browser that you use to access Cloud Manager must have a connection to that private IP address. That connection can come from a direct connection to AWS or Azure (for example, a VPN), or from a host that's inside the same network as the Cloud Compliance instance.

## Deploying the Cloud Compliance instance


You deploy an instance of Cloud Compliance for each Cloud Manager instance.

### Steps

1. In Cloud Manager, click **Cloud Compliance**.
2. Click **Activate Cloud Compliance** to start the deployment wizard.



3. The wizard displays progress as it goes through the deployment steps. It will stop and ask for input if it runs into any issues.



4. When the instance is deployed, click **Continue to configuration** to go to the *Scan Configuration* page.

### Result

Cloud Manager deploys the Cloud Compliance instance in your cloud provider.

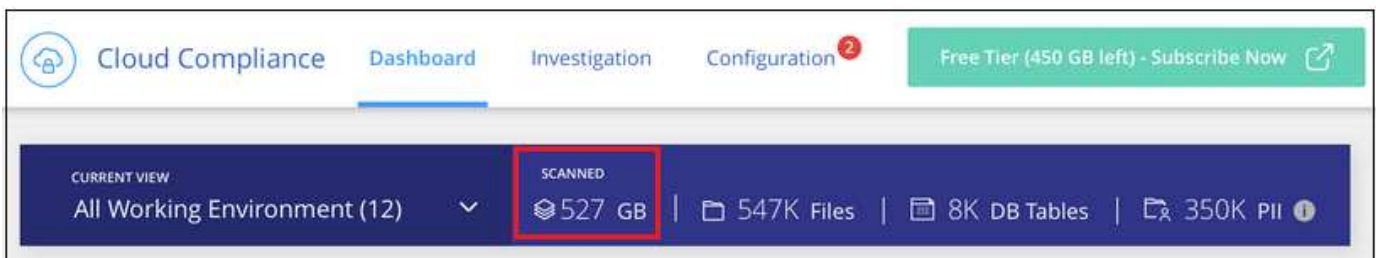
### What's Next

From the Scan Configuration page you can select the working environments, volumes, and buckets that you want to scan for compliance. You can also connect to a database server in order to scan specific database schemas. Activate Cloud Compliance on any of these data sources.

### Subscribing to the Cloud Compliance service

The first 1 TB of data that Cloud Compliance scans in a Cloud Manager workspace is free. A subscription to the AWS or Azure Marketplace is required to continue scanning data after that point.

You can subscribe at any time and you will not be charged until the amount of data exceeds 1 TB. You can always see the total amount of data that is being scanned from the Cloud Compliance Dashboard. And the *Subscribe Now* button makes it easy to subscribe when you are ready.




**Note:** If you are prompted by Cloud Compliance to subscribe, but you already have an Azure subscription, you're probably using the old **Cloud Manager** subscription and you need to change to the new **NetApp Cloud Manager** subscription. See [Changing to the new NetApp Cloud Manager plan in Azure](#) for details.

### Steps

These steps must be completed by a user who has the *Account Admin* role.

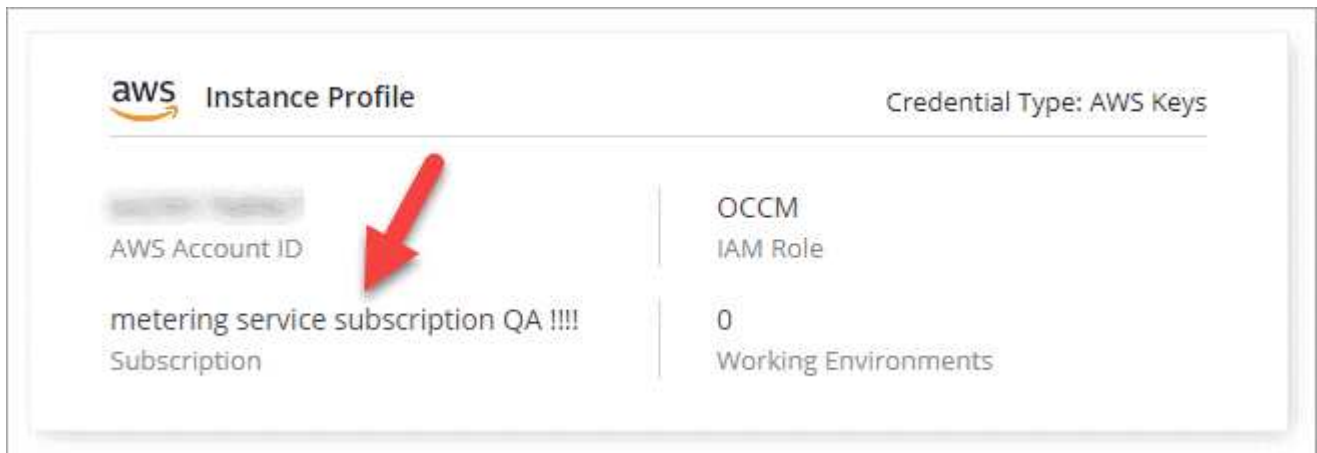
1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.



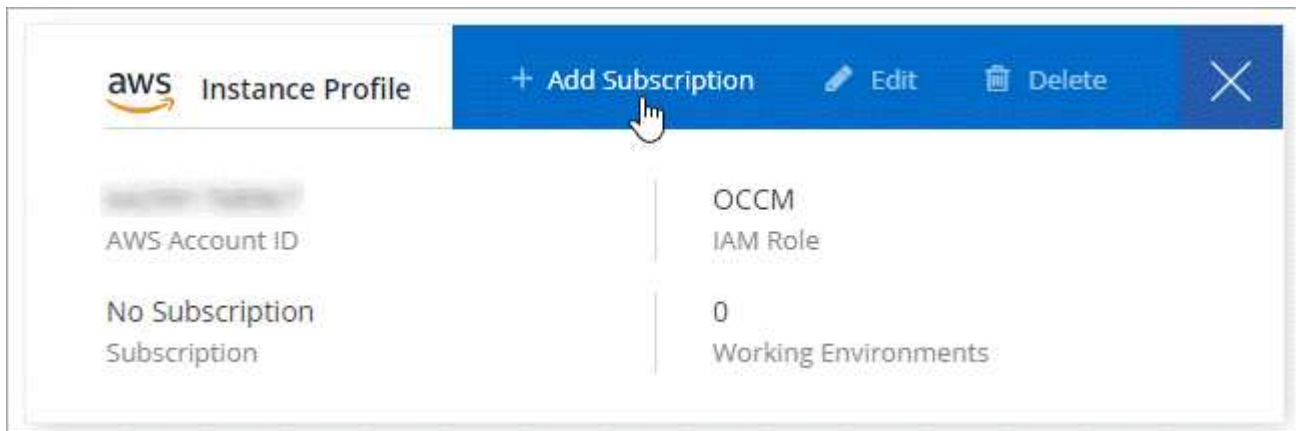
2. Find the credentials for the AWS Instance Profile or Azure Managed Service Identity.

The subscription must be added to the Instance Profile or Managed Service Identity. Charging won't work otherwise.

If you already have a subscription, then you're all set—there's nothing else that you need to do.



3. If you don't have a subscription yet, hover over the credentials and click the action menu.
4. Click **Add Subscription**.



5. Click **Add Subscription**, click **Continue**, and follow the steps.

The following video shows how to associate a Marketplace subscription to an AWS subscription:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_aws.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_aws.mp4) (video)

The following video shows how to associate a Marketplace subscription to an Azure subscription:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_azure.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_azure.mp4) (video)

## Changing to the new Cloud Manager plan in Azure

Cloud Compliance was added to the Azure Marketplace subscription named **NetApp Cloud Manager** as of October 7, 2020. If you already have the original Azure **Cloud Manager** subscription it will not allow you to use Cloud Compliance.

You need to follow these steps and select the new **NetApp Cloud Manager** subscription and then remove the old **Cloud Manager** subscription.



If your existing Subscription was issued with a special private offer, you need to contact NetApp so that we can issue a new special private offer with Compliance included.

### Steps

These steps are similar to adding a new subscription as described above, but vary in a few places.

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.
2. Find the credentials for the Azure Managed Service Identity that you want to change the subscription for and hover over the credentials and click **Associate Subscription**.

The details for your current Marketplace Subscription are displayed.

3. Click **Add Subscription**, click **Continue**, and follow the steps. You are redirected to Azure portal in order to create the new subscription.
4. Make sure you select the plan **NetApp Cloud Manager** that provides access to Cloud Compliance and not **Cloud Manager**.
5. Go through the steps in the video to associate a Marketplace subscription to an Azure subscription:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_azure.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_azure.mp4) (video)

6. Return to Cloud Manager, select the new subscription, and click **Associate**.
7. To verify your subscription has changed, hover over the “i” above subscription in the Credentials card.

Now you can unsubscribe your old subscription from the Azure portal.

8. In the Azure portal, go to Software as a Service (SaaS), select the subscription, and click **Unsubscribe**.

## Activate scanning on your data sources

### Getting started with Cloud Compliance for Cloud Volumes ONTAP and Azure NetApp Files

Complete a few steps to get started with Cloud Compliance for Cloud Volumes ONTAP or Azure NetApp Files.

#### Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



## Deploy the Cloud Compliance instance

Deploy [Cloud Compliance in Cloud Manager](#) if there isn't already an instance deployed.



## Enable Cloud Compliance in your working environments

Click **Cloud Compliance**, select the **Configuration** tab, and activate compliance scans for specific working environments.



## Ensure access to volumes

Now that Cloud Compliance is enabled, ensure that it can access volumes.

- The Cloud Compliance instance needs a network connection to each Cloud Volumes ONTAP subnet or Azure NetApp Files subnet.
- Security groups for Cloud Volumes ONTAP must allow inbound connections from the Cloud Compliance instance.
- NFS volume export policies must allow access from the Cloud Compliance instance.
- Cloud Compliance needs Active Directory credentials to scan CIFS volumes.

Click **Cloud Compliance > Scan Configuration > Edit CIFS Credentials** and provide the credentials. The credentials can be read-only, but providing admin credentials ensures that Cloud Compliance can read data that requires elevated permissions.



## Configure volumes to scan

Select the volumes that you'd like to scan and Cloud Compliance will start scanning them.

### Deploying the Cloud Compliance instance


Deploy [Cloud Compliance in Cloud Manager](#) if there isn't already an instance deployed.

### Enabling Cloud Compliance in your working environments

1. At the top of Cloud Manager, click **Cloud Compliance** and then select the **Configuration** tab.


View Dashboard >

Scan Configuration [How to add AWS accounts to scan S3](#)

 **AWS Account Number 1**  
Amazon S3

---


To enable Compliance for Amazon S3 on this AWS account or other, go to [Working Environment tab](#), select the Amazon S3 cloud and activate Compliance from the right hand panel.

 **Azure Netapp Files**  
Azure NetApp Files ⓘ

---

[Activate Compliance for All Volumes](#)

or select Volumes

 **Working Environment Name 1**  
Cloud Volumes ONTAP

---

[Activate Compliance for All Volumes](#)

or select Volumes

2. To scan all volumes in a working environment, click **Activate Compliance for All Volumes**.

To scan only certain volumes in a working environment, click **or select Volumes** and then choose the volumes you want to scan.

See [Enabling and disabling compliance scans on volumes](#) for details.

## Result

Cloud Compliance starts scanning the data on each working environment. Results will be available in the Compliance dashboard as soon as Cloud Compliance finishes the initial scans. The time that it takes depends on the amount of data—it could be a few minutes or hours.

## Verifying that Cloud Compliance has access to volumes

Make sure that Cloud Compliance can access volumes by checking your networking, security groups, and export policies. You'll need to provide Cloud Compliance with CIFS credentials so it can access CIFS volumes.

## Steps

1. Make sure that there's a network connection between the Cloud Compliance instance and each network that includes volumes for Cloud Volumes ONTAP or Azure NetApp Files.



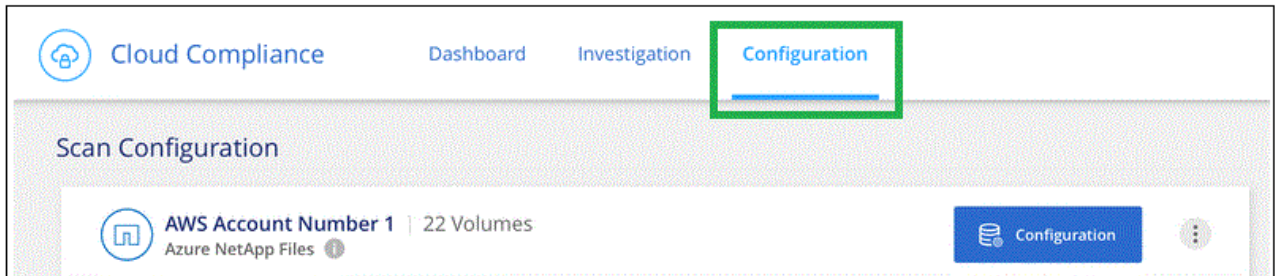
For Azure NetApp Files, Cloud Compliance can only scan volumes that are in the same region as Cloud Manager.

2. Ensure that the security group for Cloud Volumes ONTAP allows inbound traffic from the Cloud Compliance instance.



You can either open the security group for traffic from the IP address of the Cloud Compliance instance, or you can open the security group for all traffic from inside the virtual network.

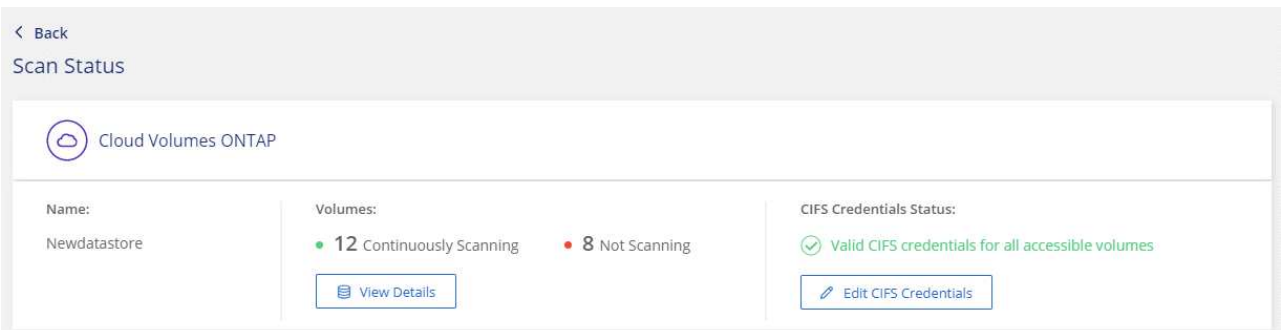
3. Ensure that NFS volume export policies include the IP address of the Cloud Compliance instance so it can access the data on each volume.
4. If you use CIFS, provide Cloud Compliance with Active Directory credentials so it can scan CIFS volumes.
  - a. At the top of Cloud Manager, click **Cloud Compliance**.
  - b. Click the **Configuration** tab.



- c. For each working environment, click **Edit CIFS Credentials** and enter the user name and password that Cloud Compliance needs to access CIFS volumes on the system.

The credentials can be read-only, but providing admin credentials ensures that Cloud Compliance can read any data that requires elevated permissions. The credentials are stored on the Cloud Compliance instance.

After you enter the credentials, you should see a message that all CIFS volumes were authenticated successfully.



5. On the *Scan Configuration* page, click **View Details** to review the status for each CIFS and NFS volume and correct any errors.

For example, the following image shows three volumes; one of which Cloud Compliance can't scan due to network connectivity issues between the Cloud Compliance instance and the volume.



### Enabling and disabling compliance scans on volumes

You can stop or start scanning volumes in a working environment at any time from the Scan Configuration page. We recommend that you scan all volumes.

To:	Do this:
Disable scanning for a volume	Move the volume slider to the left
Disable scanning for all volumes	Move the <b>Activate Compliance for all Volumes</b> slider to the left
Enable scanning for a volume	Move the volume slider to the right
Enable scanning for all volumes	Move the <b>Activate Compliance for all Volumes</b> slider to the right



New volumes added to the working environment are automatically scanned only when the **Activate Compliance for all Volumes** setting is enabled. When this setting is disabled, you'll need to activate scanning on each new volume you create in the working environment.

### Scanning data protection volumes

By default, data protection (DP) volumes are not scanned because they are not exposed externally and Cloud Compliance cannot access them. These volumes are typically the destination volumes for SnapMirror operations from an on-premises ONTAP cluster.

Initially, the Cloud Compliance volume list identifies these volumes as *Type DP* with the *Status Not Scanning* and the *Required Action Enable Access to DP volumes*.

Compliance	Volume Name	Type	Status	Required Action
<input type="checkbox"/>	VolumeName1	DP	Not Scanning	Enable access to DP Volumes
<input checked="" type="checkbox"/>	VolumeName2	NFS	Continuously Scanning	
<input type="checkbox"/>	VolumeName3	CIFS	Not Scanning	

## Steps

If you want to scan these data protection volumes:

1. Click the **Enable Access to DP volumes** button at the top of the page.
2. Activate each DP volume that you want to scan, or use the **Activate Compliance for all Volumes** control to enable all volumes, including all DP volumes.

Once enabled, Cloud Compliance creates an NFS share from each DP volume that was activated for Compliance so that it can be scanned. The share export policies only allow access from the Cloud Compliance instance.



Only volumes that were initially created as NFS volumes in the source ONTAP system are shown in the volume list. Source volumes that were created initially as CIFS do not currently appear in Cloud Compliance.

## Getting started with Cloud Compliance for Amazon S3

Cloud Compliance can scan your Amazon S3 buckets to identify the personal and sensitive data that resides in S3 object storage. Cloud Compliance can scan any bucket in the account, regardless if it was created for a NetApp solution.

### Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.



### 1 Set up the S3 requirements in your cloud environment

Ensure that your cloud environment can meet the requirements for Cloud Compliance, including preparing an IAM role and setting up connectivity from Cloud Compliance to S3. [See the complete list.](#)



### 2 Deploy the Cloud Compliance instance

[Deploy Cloud Compliance in Cloud Manager](#) if there isn't already an instance deployed.

### 3

#### Activate Compliance on your S3 working environment

Select the Amazon S3 working environment, click **Enable Compliance**, and select an IAM role that includes the required permissions.

### 4

#### Select the buckets to scan

Select the buckets that you'd like to scan and Cloud Compliance will start scanning them.

#### Reviewing S3 prerequisites

The following requirements are specific to scanning S3 buckets.

#### Set up an IAM role for the Cloud Compliance instance

Cloud Compliance needs permissions to connect to the S3 buckets in your account and to scan them. Set up an IAM role that includes the permissions listed below. Cloud Manager prompts you to select an IAM role when you enable Cloud Compliance on the Amazon S3 working environment.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:Get*",
        "s3:List*",
        "s3:HeadBucket"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedRolePolicies"
      ],
      "Resource": [
        "arn:aws:iam::*:policy/*",
        "arn:aws:iam::*:role*"
      ]
    }
  ]
}
```

## Provide connectivity from Cloud Compliance to Amazon S3

Cloud Compliance needs a connection to Amazon S3. The best way to provide that connection is through a VPC Endpoint to the S3 service. For instructions, see [AWS Documentation: Creating a Gateway Endpoint](#).

When you create the VPC Endpoint, be sure to select the region, VPC, and route table that corresponds to the Cloud Compliance instance. You must also modify the security group to add an outbound HTTPS rule that enables traffic to the S3 endpoint. Otherwise, Cloud Compliance can't connect to the S3 service.

If you experience any issues, see [AWS Support Knowledge Center: Why can't I connect to an S3 bucket using a gateway VPC endpoint?](#)

An alternative is to provide the connection by using a NAT Gateway.



You can't use a proxy to get to S3 over the internet.

## Deploying the Cloud Compliance instance

[Deploy Cloud Compliance in Cloud Manager](#) if there isn't already an instance deployed.

You need to deploy the instance in an AWS Connector so that Cloud Manager automatically discovers the S3 buckets in this AWS account and displays them in an Amazon S3 working environment.

## Activating Compliance on your S3 working environment

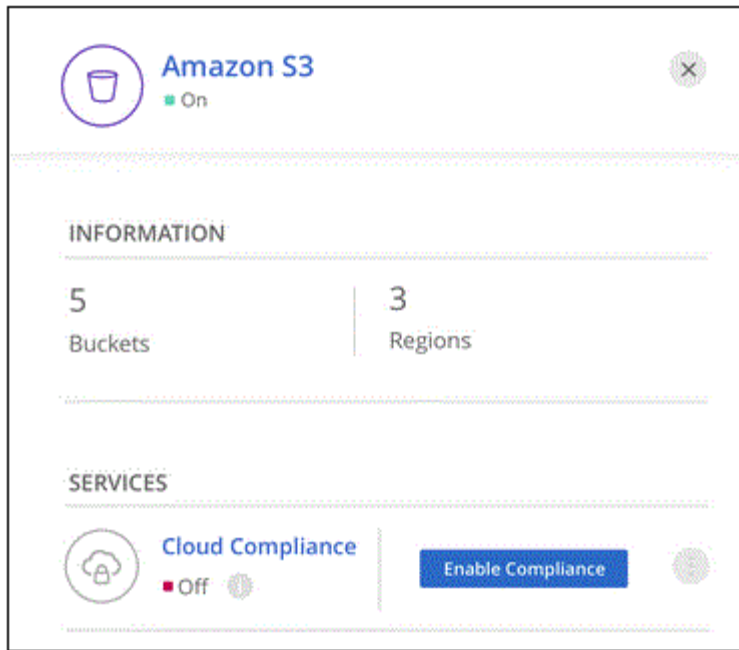
Enable Cloud Compliance on Amazon S3 after you verify the prerequisites.

### Steps

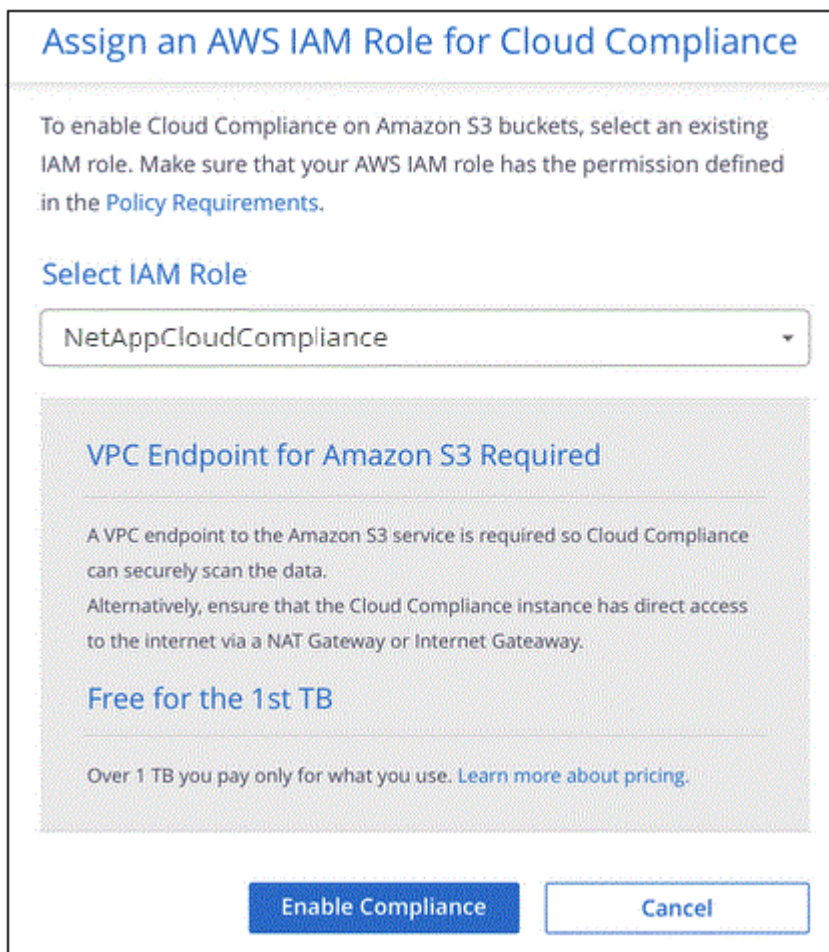
1. At the top of Cloud Manager, click **Working Environments**.
2. Select the Amazon S3 working environment.



3. In the pane on the right, click **Enable Compliance**.




4. When prompted, assign an IAM role to the Cloud Compliance instance that has [the required permissions](#).



5. Click **Enable Compliance**.



You can also enable compliance scans for a working environment from the Scan Configuration page by clicking the  button and selecting **Activate Compliance**.

## Result

Cloud Manager assigns the IAM role to the instance.

## Enabling and disabling compliance scans on S3 buckets

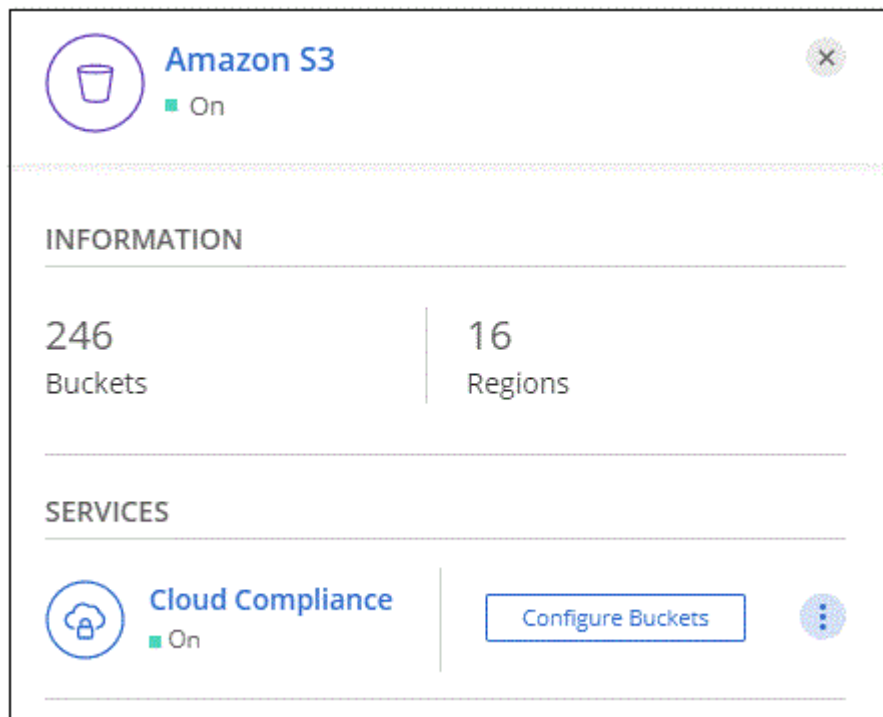
After Cloud Manager enables Cloud Compliance on Amazon S3, the next step is to configure the buckets that you want to scan.

When Cloud Manager is running in the AWS account that has the S3 buckets you want to scan, it discovers those buckets and displays them in an Amazon S3 working environment.

Cloud Compliance can also [scan S3 buckets that are in different AWS accounts](#).

## Steps

1. Select the Amazon S3 working environment.
2. In the pane on the right, click **Configure Buckets**.



3. Enable compliance on the buckets that you want to scan.

Compliance	Bucket Name	Status	Required Action
<input checked="" type="checkbox"/>	BucketName1	Not Scanning	Add Credentials
<input checked="" type="checkbox"/>	BucketName2	Continuously Scanning	
<input type="checkbox"/>	BucketName3	Not Scanning	

## Result

Cloud Compliance starts scanning the S3 buckets that you enabled. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

## Scanning buckets from additional AWS accounts

You can scan S3 buckets that are under a different AWS account by assigning a role from that account to access the existing Cloud Compliance instance.

## Steps

1. Go to the target AWS account where you want to scan S3 buckets and create an IAM role by selecting **Another AWS account**.

## Create role



### Select type of trusted entity

**AWS service**  
EC2, Lambda and others

**Another AWS account**  
Belonging to you or 3rd party

**Web identity**  
Cognito or any OpenID provider

**SAML 2.0 federation**  
Your corporate directory

Allows entities in other accounts to perform actions in this account. [Learn more](#)

### Specify accounts that can use this role

Account ID\*  ⓘ

- Options**
- Require external ID (Best practice when a third party will assume this role)
  - Require MFA ⓘ

Be sure to do the following:

- Enter the ID of the account where the Cloud Compliance instance resides.
- Change the **Maximum CLI/API session duration** from 1 hour to 12 hours and save that change.
- Attach the Cloud Compliance IAM policy. Make sure it has the required permissions.



```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:Get*",
        "s3:List*",
        "s3:HeadBucket"
      ],
      "Resource": "*"
    },
  ],
}
```

2. Go to the source AWS account where the Cloud Compliance instance resides and select the IAM role that is attached to the instance.
  - a. Change the **Maximum CLI/API session duration** from 1 hour to 12 hours and save that change.
  - b. Click **Attach policies** and then click **Create policy**.
  - c. Create a policy that includes the "sts:AssumeRole" action and the ARN of the role that you created in the target account.




```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::<ADDITIONAL-ACCOUNT-
ID>:role/<ADDITIONAL_ROLE_NAME>"
    },
    {
      "Effect": "Allow",
      "Action": [
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedRolePolicies"
      ],
      "Resource": [
        "arn:aws:iam::*:policy/*",
        "arn:aws:iam::*:role/*"
      ]
    }
  ]
}

```

The Cloud Compliance instance profile account now has access to the additional AWS account.

3. Go to the **Amazon S3 Scan Configuration** page and the new AWS account is displayed. Note that it can take a few minutes for Cloud Compliance to sync the new account's working environment and show this information.



4. Click **Activate Compliance & Select Buckets** and select the buckets you want to scan.

### Result

Cloud Compliance starts scanning the new S3 buckets that you enabled.

### Scanning database schemas

Complete a few steps to start scanning your database schemas with Cloud Compliance.

## Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



### Review database prerequisites

Ensure that your database is supported and that you have the information necessary to connect to the database.



### Deploy the Cloud Compliance instance

[Deploy Cloud Compliance in Cloud Manager](#) if there isn't already an instance deployed.



### Add the database server

Add the database server that you want to access.



### Select the schemas

Select the schemas that you want to scan.

## Reviewing prerequisites

Review the following prerequisites to make sure that you have a supported configuration before you enable Cloud Compliance.

## Supported databases

Cloud Compliance can scan schemas from the following databases:

- MongoDB
- Oracle
- PostgreSQL
- SAP HANA
- SQL Server (MSSQL)



The statistics gathering feature **must be enabled** in the database.

## Database requirements

Any database with connectivity to the Cloud Compliance instance can be scanned, regardless of where it is hosted. You just need the following information to connect to the database:

- IP Address or host name
- Port

- Service name (only for accessing Oracle databases)
- Credentials that allow read access to the schemas

When choosing a user name and password, it's important to choose one that has full read permissions to all the schemas and tables you want to scan. We recommend that you create a dedicated user for the Cloud Compliance system with all the required permissions.

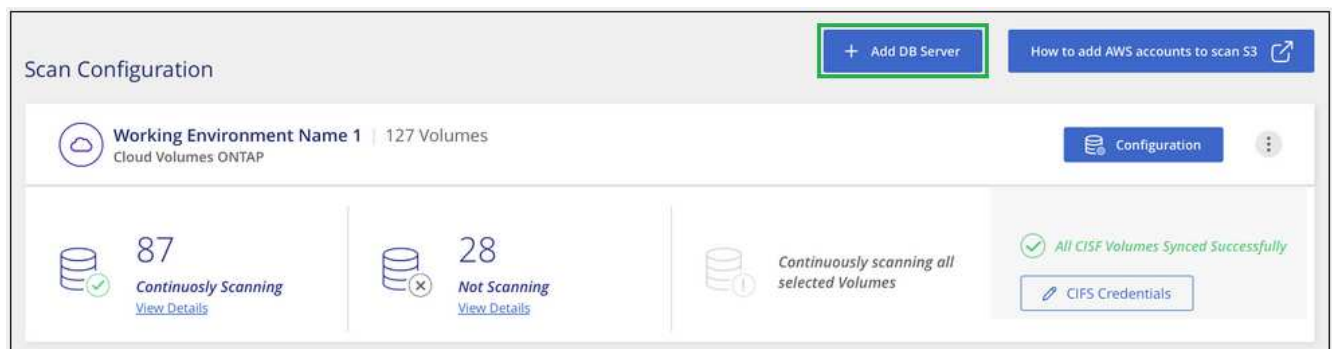
**Note:** For MongoDB, a read-only Admin role is required.

### Adding the database server

You must have [deployed an instance of Cloud Compliance in Cloud Manager already](#).

Add the database server where the schemas reside.

1. From the *Scan Configuration* page, click the **Add DB Server** button.



2. Enter the required information to identify the database server.
  - a. Select the database type.
  - b. Enter the port and the host name or IP address to connect to the database.
  - c. For Oracle databases, enter the Service name.
  - d. Enter the credentials so that Cloud Compliance can access the server.
  - e. Click **Add DB Server**.

## Add DB Server

To activate Compliance on Databases, first add a Database Server. After this step, you'll be able to select which Database Schemas you would like to activate Compliance for.

**Database**

Database Type  Host Name or IP Address

Port  Service Name

**Credentials**

Username  Password

The database is added to the list of working directories.

### Enabling and disabling compliance scans on database schemas

You can stop or start scanning schemas at any time.

1. From the *Scan Configuration* page, click the **Configuration** button for the database you want to configure.

Scan Configuration

Oracle DB 1 | 41 Schemas  
Oracle

No Schemas selected for Compliance

7 Not Scanning  
[View Details](#)

2. Select the schemas that you want to scan by moving the slider to the right.

'Working Environment Name' Scan Configuration			
Compliance	Schema Name	Status	Required Action
<input type="checkbox"/>	DB1 - SchemaName1	Not Scanning	Add Credentials
<input checked="" type="checkbox"/>	DB1 - SchemaName2	Continuously Scanning	
<input checked="" type="checkbox"/>	DB1 - SchemaName3	Continuously Scanning	
<input checked="" type="checkbox"/>	DB1 - SchemaName4	Continuously Scanning	

## Result

Cloud Compliance starts scanning the database schemas that you enabled. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

## Removing a database from Cloud Manager

If you no longer want to scan a certain database, you can delete it from the Cloud Manager interface and stop all scans.

From the *Scan Configuration* page, click the  button in the row for the database, and then click **Remove DB Server**.



## Scanning on-premises ONTAP data with Cloud Compliance by using SnapMirror

You can scan your on-premises ONTAP data with Cloud Compliance by replicating the on-prem NFS or CIFS data to a Cloud Volumes ONTAP working environment and then enabling compliance. Scanning the data directly from an on-premises ONTAP working environment isn't supported.

You must have [deployed an instance of Cloud Compliance in Cloud Manager already](#).

## Steps

1. From Cloud Manager, create a SnapMirror relationship between the on-premises ONTAP cluster and Cloud Volumes ONTAP.
  - a. [Discover the on-premises cluster in Cloud Manager](#).
  - b. [Create a SnapMirror replication between the on-premises ONTAP cluster and Cloud Volumes ONTAP](#)

from Cloud Manager.

2. For DP volumes that were created from SMB source volumes, from the ONTAP CLI, configure the SMB destination volumes for data access. (This is not required for NFS volumes because data access is enabled automatically through Cloud Compliance.)
  - a. [Create a SMB share on the destination volume.](#)
  - b. [Apply the appropriate ACLs to the SMB share at the destination volume.](#)
3. From Cloud Manager, activate Cloud Compliance on the Cloud Volumes ONTAP working environment that contains the SnapMirror data:
  - a. Click **Working Environments**.
  - b. Select the working environment that contains the SnapMirror data and click **Enable Compliance**.

[Click here if you need help with enabling Cloud Compliance on a Cloud Volumes ONTAP system.](#)

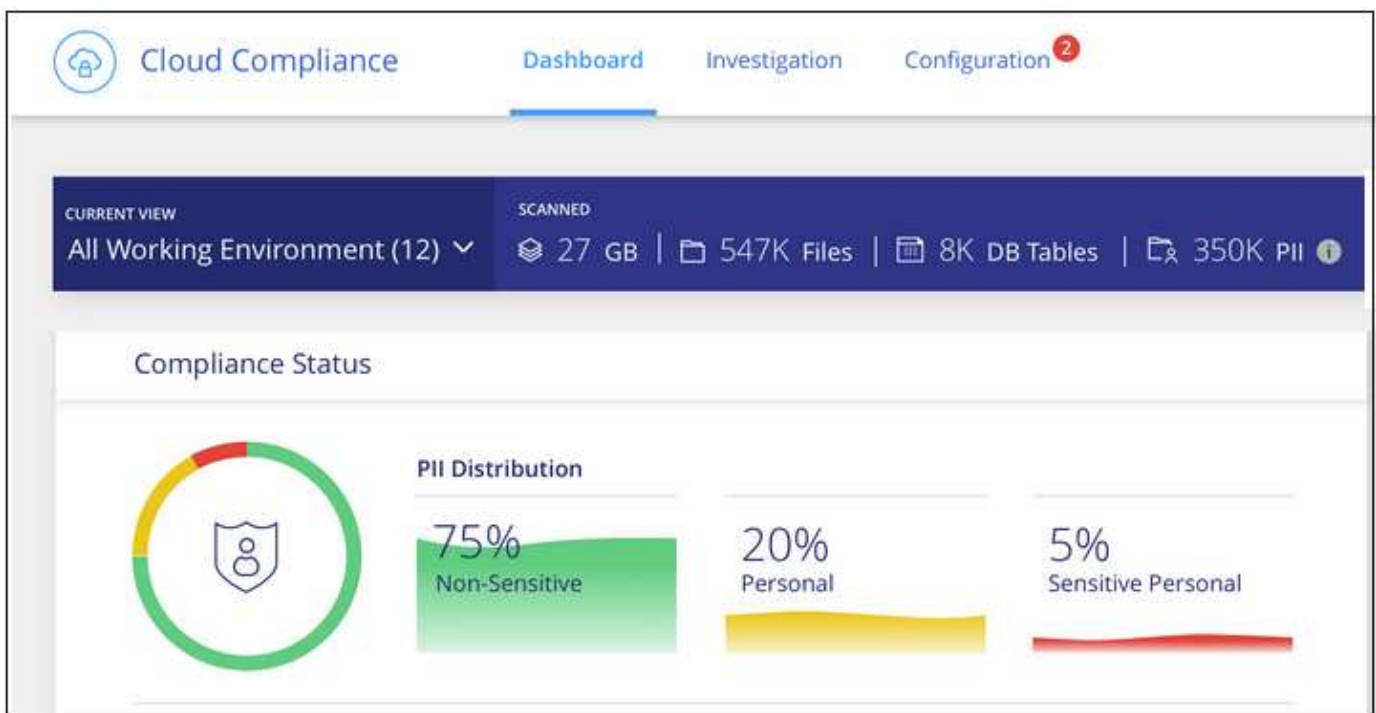
- c. Click the **Enable Access to DP volumes** button at the top of the *Scan Configuration* page.
- d. Activate each DP volume that you want to scan, or use the **Activate Compliance for all Volumes** control to enable all volumes, including all DP volumes.

See [Scanning data protection volumes](#) for more information about scanning DP volumes.

## Gaining visibility and control of private data

Gain control of your private data by viewing details about the personal data and sensitive personal data in your organization. You can also gain visibility by reviewing the categories and file types that Cloud Compliance found in your data.

By default, the Cloud Compliance dashboard displays compliance data for all working environments and databases.



If you want to see data for only some of the working environments, [select those working environments](#).

## Personal data

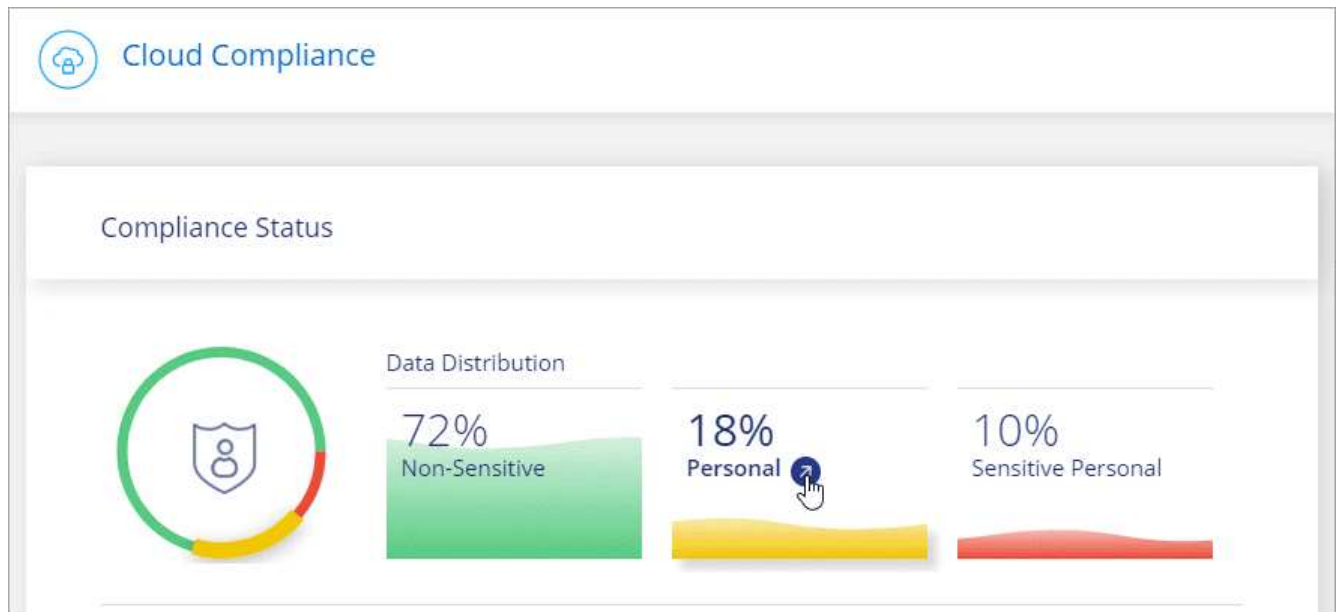
Cloud Compliance automatically identifies specific words, strings, and patterns (Regex) inside the data. For example, Personal Identification Information (PII), credit card numbers, social security numbers, bank account numbers, and more. [See the full list](#).

For some types of personal data, Cloud Compliance uses *proximity validation* to validate its findings. The validation occurs by looking for one or more predefined keywords in proximity to the personal data that was found. For example, Cloud Compliance identifies a U.S. social security number (SSN) as a SSN if it sees a proximity word next to it—for example, *SSN* or *social security*. [The list below](#) shows when Cloud Compliance uses proximity validation.

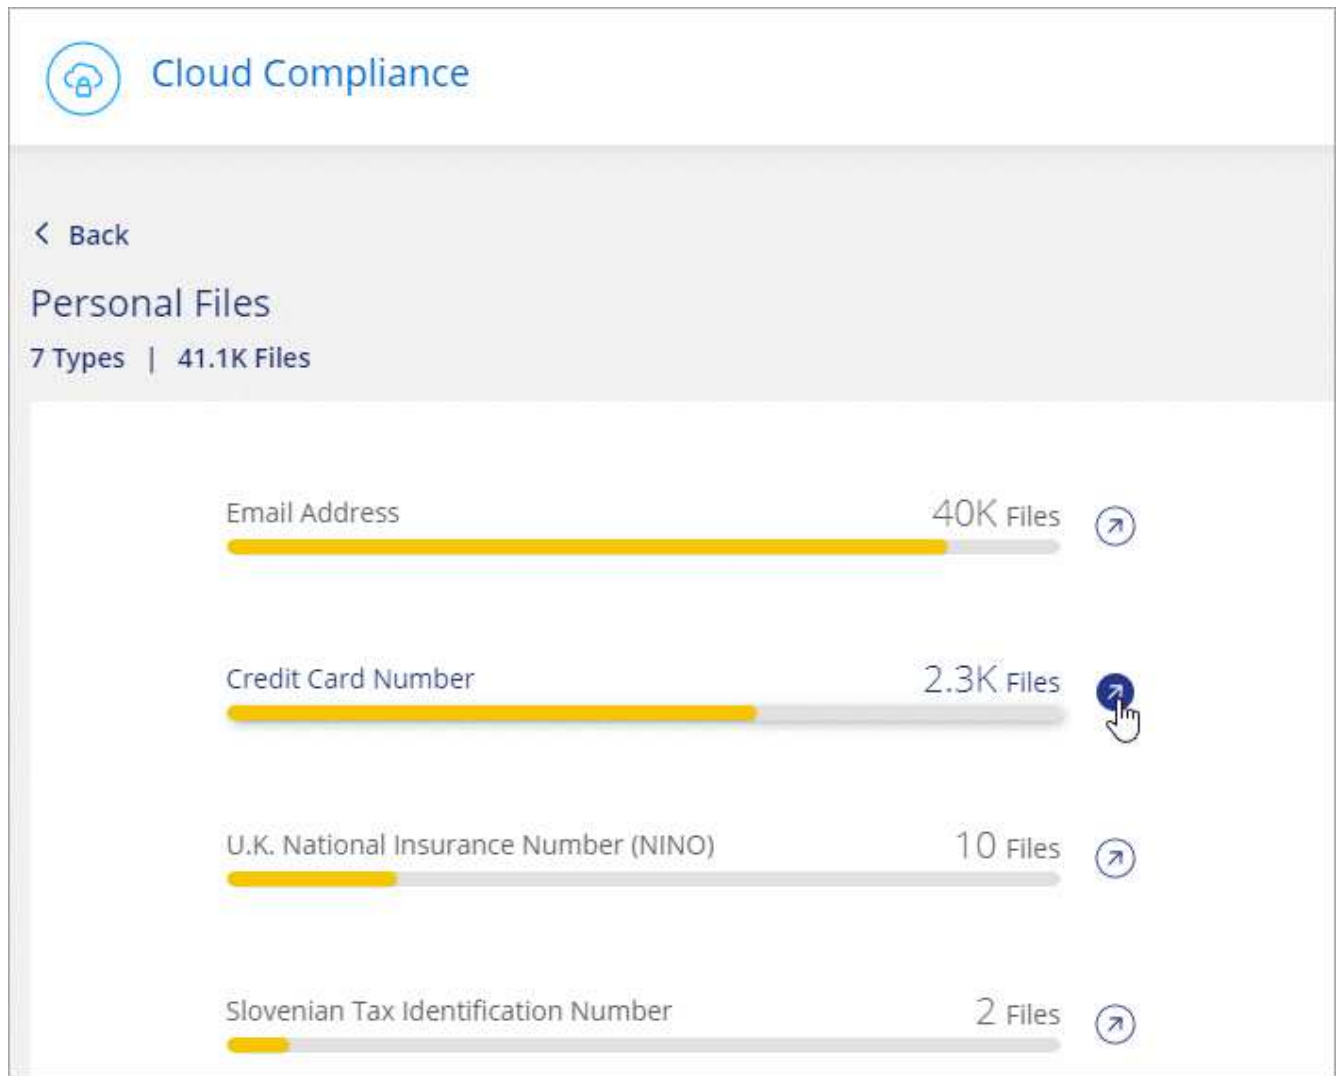
## Viewing files that contain personal data

### Steps

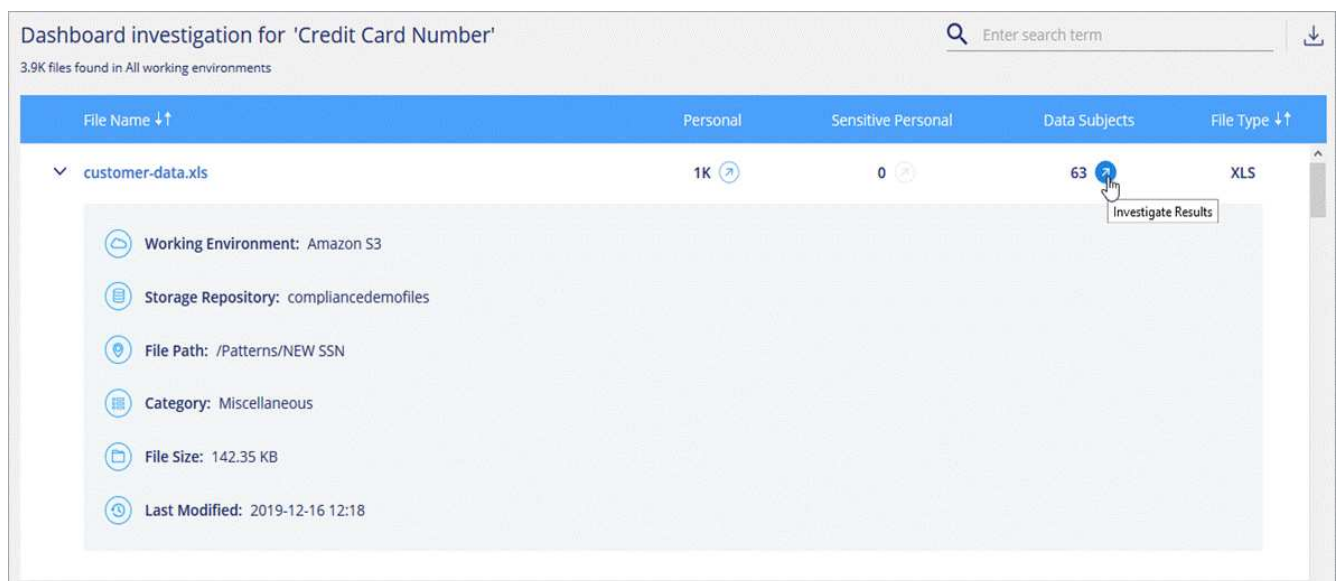
1. At the top of Cloud Manager, click **Cloud Compliance** and click the **Dashboard** tab.
2. To investigate the details for all personal data, click the icon next to the personal data percentage.



3. To investigate the details for a specific type of personal data, click **View All** and then click the **Investigate Results** icon for a specific type of personal data.



- Investigate the data by searching, sorting, expanding details for a specific file, clicking **Investigate Results** to see masked information, or by downloading the file list.



- You can also filter the contents of the investigation page to display only the results you want to see. The



top-level tabs allow you to view data from files (unstructured data) or from databases (structured data).

Then you have filters for working environment, storage repository, category, private data, file type, last modified date, and whether the S3 object's permissions are open to public access.

The screenshot shows a web interface titled "Dashboard Investigation". At the top, there are two tabs: "Unstructured (32K Files)" (selected) and "Structured (323 DB Tables)". On the left, there is a "FILTERS:" section with a "Clear All" button and several filter categories: "Working Environment" (4 items), "Storage Repository", "Category", "Private Data" (6 items), and "File Type". The main table displays a list of files. Each row represents a file with the following columns: "File Name" (with a dropdown arrow), "Personal" (with a dropdown arrow), "Sensitive Personal" (with a dropdown arrow), "Data Subjects" (with a dropdown arrow), and "File Type". The first row shows an "Expense Report EXP-TPO-10603888765435" with a "cvo" tag, 6 items in the "Personal" column, 3 items in the "Sensitive Personal" column, 16 items in the "Data Subjects" column, and a "PDF" file type. This pattern repeats for several rows.

### Types of personal data

The personal data found in files can be general personal data or national identifiers. The third column identifies whether Cloud Compliance uses [proximity validation](#) to validate its findings for the identifier.

Type	Identifier	Proximity validation?
General	Email address	No
	Credit card number	No
	IBAN number (International Bank Account Number)	No

Type	Identifier	Proximity validation?
National Identifiers	Belgian ID (Numero National)	Yes
	Brazilian ID (CPF)	Yes
	Bulgarian ID (UCN)	Yes
	California Driver's License	Yes
	Croatian ID (OIB)	Yes
	Cyprus Tax Identification Number (TIC)	Yes
	Czech/Slovak ID	Yes
	Danish ID (CPR)	Yes
	Dutch ID (BSN)	Yes
	Estonian ID	Yes
	Finnish ID (HETU)	Yes
	French Tax Identification Number (SPI)	Yes
	German Tax Identification Number (Steuerliche Identifikationsnummer)	Yes
	Greek ID	Yes
	Hungarian Tax Identification Number	Yes
	Irish ID (PPS)	Yes
	Israeli ID	Yes
	Italian Tax Identification Number	Yes
	Latvian ID	Yes
	Lithuanian ID	Yes
	Luxembourg ID	Yes
	Maltese ID	Yes
	Polish ID (PESEL)	Yes
	Portuguese Tax Identification Number (NIF)	Yes
	Romanian ID (CNP)	Yes
	Slovenian ID (EMSO)	Yes
	South African ID	Yes
	Spanish Tax Identification Number	Yes
	Swedish ID	Yes
	U.K. ID (NINO)	Yes
USA Social Security Number (SSN)	Yes	

## Sensitive personal data

Cloud Compliance automatically identifies special types of sensitive personal information, as defined by privacy regulations such as [articles 9 and 10 of the GDPR](#). For example, information regarding a person's health, ethnic origin, or sexual orientation. [See the full list](#).

Cloud Compliance uses artificial intelligence (AI), natural language processing (NLP), machine learning (ML), and cognitive computing (CC) to understand the meaning of the content that it scans in order to extract entities and categorize it accordingly.

For example, one sensitive GDPR data category is ethnic origin. Because of its NLP abilities, Cloud Compliance can distinguish the difference between a sentence that reads "George is Mexican" (indicating sensitive data as specified in article 9 of the GDPR), versus "George is eating Mexican food."

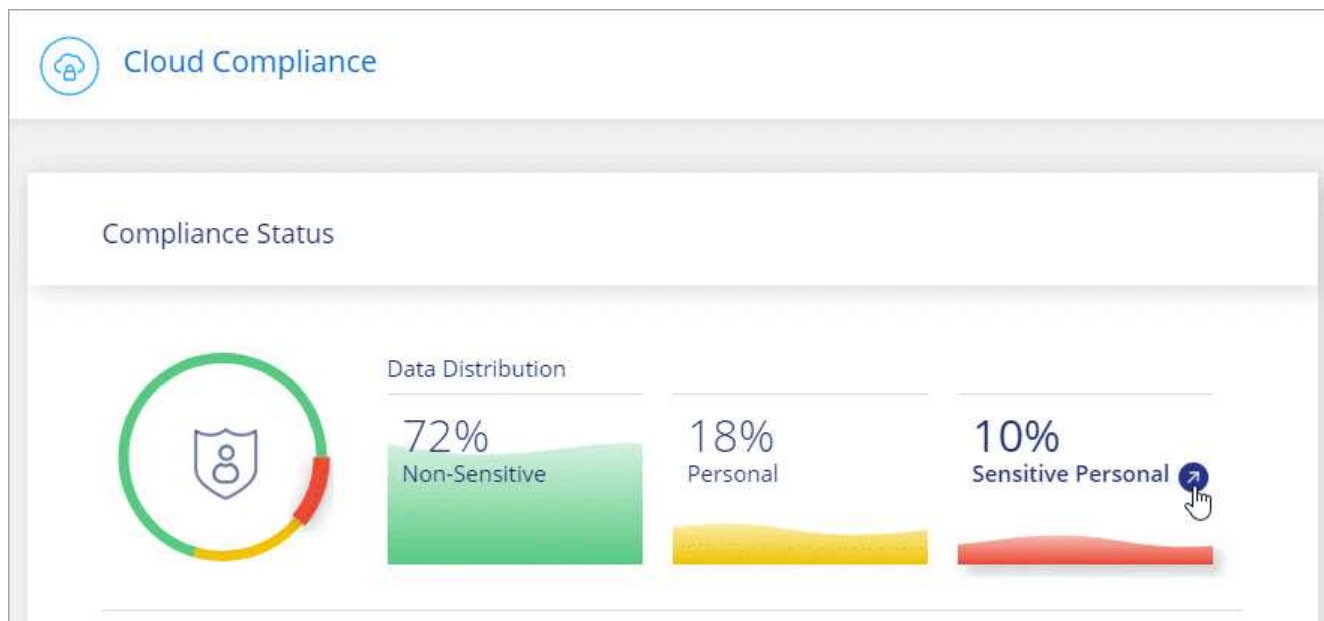


Only English is supported when scanning for sensitive personal data. Support for more languages will be added later.


### Viewing files that contain sensitive personal data

#### Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. To investigate the details for all sensitive personal data, click the icon next to the sensitive personal data percentage.



3. To investigate the details for a specific type of sensitive personal data, click **View All** and then click the **Investigate Results** icon for a specific type of sensitive personal data.



4. Investigate the data by searching, sorting, expanding details for a specific file, clicking **Investigate Results** to see masked information, or by downloading the file list.

### Types of sensitive personal data

The sensitive personal data that Cloud Compliance can find in files includes the following:

#### **Criminal Procedures Reference**

Data concerning a natural person's criminal convictions and offenses.

#### **Ethnicity Reference**

Data concerning a natural person's racial or ethnic origin.

#### **Health Reference**

Data concerning a natural person's health.

#### **ICD-9-CM Medical Codes**

Codes used in the medical and health industry.

#### **ICD-10-CM Medical Codes**

Codes used in the medical and health industry.

### Philosophical Beliefs Reference

Data concerning a natural person's philosophical beliefs.

### Religious Beliefs Reference

Data concerning a natural person's religious beliefs.

### Sex Life or Orientation Reference

Data concerning a natural person's sex life or sexual orientation.

## Categories

Cloud Compliance takes the data that it scanned and divides it into different types of categories. Categories are topics based on AI analysis of the content and metadata of each file. [See the list of categories.](#)

Categories can help you understand what's happening with your data by showing you the types of information that you have. For example, a category like resumes or employee contracts can include sensitive data. When you investigate the results, you might find that employee contracts are stored in an insecure location. You can then correct that issue.

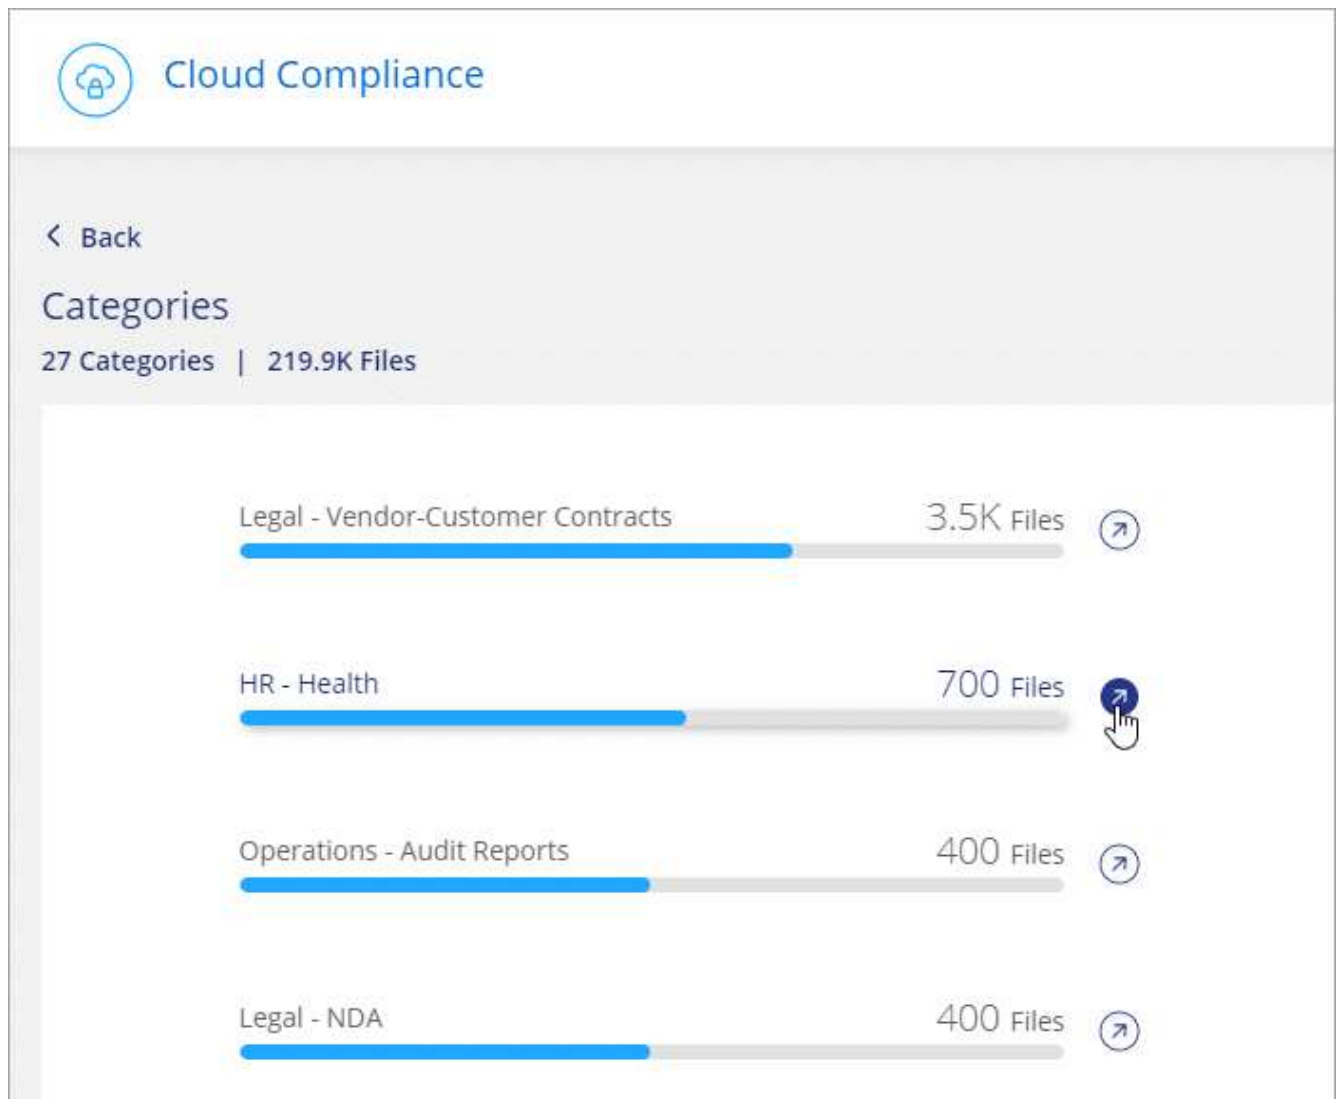


Only English is supported for categories. Support for more languages will be added later.

### Viewing files by categories

#### Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Click the **Investigate Results** icon for one of the top 4 categories directly from the main screen, or click **View All** and then click the icon for any of the categories.



3. Investigate the data by searching, sorting, expanding details for a specific file, clicking **Investigate Results** to see masked information, or by downloading the file list.

### Types of categories

Cloud Compliance categorizes your data as follows:

#### Finance

- Balance Sheets
- Purchase Orders
- Invoices
- Quarterly Reports

#### HR

- Background Checks
- Compensation Plans
- Employee Contracts
- Employee Reviews

- Health
- Resumes

### **Legal**

- NDAs
- Vendor-Customer contracts

### **Marketing**

- Campaigns
- Conferences

### **Operations**

- Audit Reports

### **Sales**

- Sales Orders

### **Services**

- RFI
- RFP
- SOW
- Training

### **Support**

- Complaints and Tickets

### **Metadata categories**

- Application Data
- Archive Files
- Audio
- Business Application Data
- CAD Files
- Code
- Database and index files
- Design Files
- Email Application Data
- Executables
- Financial Application Data
- Health Application Data
- Images
- Logs
- Miscellaneous Documents
- Miscellaneous Presentations

- Miscellaneous Spreadsheets
- Videos

## File types


Cloud Compliance takes the data that it scanned and breaks it down by file type. Reviewing your file types can help you control your sensitive data because you might find that certain file types are not stored correctly. [See the list of file types.](#)

For example, you might be storing CAD files that include very sensitive information about your organization. If they are unsecured, you can take control of the sensitive data by restricting permissions or moving the files to another location.

### Viewing file types

#### Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Click the **Investigate Results** icon for one of the top 4 file types directly from the main screen, or click **View All** and then click the icon for any of the file types.



3. Investigate the data by searching, sorting, expanding details for a specific file, clicking **Investigate Results** to see masked information, or by downloading the file list.



## Types of files

Cloud Compliance scans all files for category and metadata insights and displays all file types in the file types section of the dashboard.

But when Cloud Compliance detects Personal Identifiable Information (PII), or when it performs a DSAR search, only the following file formats are supported:

.PDF, .DOCX, .DOC, .PPTX, .XLS, .XLSX, .CSV, .TXT, .RTF, and .JSON.

## Viewing data from specific working environments

You can filter the contents of the Cloud Compliance dashboard to see compliance data for all working environments and databases, or for just specific working environments.

When you filter the dashboard, Cloud Compliance scopes the compliance data and reports to just those working environments that you selected.

### Steps

1. Click the filter drop-down, select the working environments that you'd like to view data for, and click **View**.

The screenshot shows the Cloud Compliance dashboard interface. At the top left, there is a 'Cloud Compliance' header with a house icon. Below this, a filter menu is open, displaying 'All Working Environments (12)' with an upward arrow. The menu includes a 'Select all' option and five individual working environments, each with a checkbox and a 'View' button: 'ANF - Azure NetApp Files' (ANF), 'Working Environment Name 1' (CVO), 'Working Environment Name 2' (CVS), 'Working Environment Name 3' (CVS), and 'Working Environment Name 4' (CVO). At the bottom of the menu are 'View' and 'Cancel' buttons. To the right of the menu, the dashboard displays compliance statistics: '20% Personal' (represented by a yellow bar) and '5% Sensitive Personal' (represented by a red bar). Below these, there are two sections: 'Personal Files' (7,000) and 'Sensitive Personal Files' (7,000), each with a 'View All' button. Under 'Personal Files', there are two categories: 'Email Address' and 'Credit Card', both showing 2,700 files with yellow progress bars. Under 'Sensitive Personal Files', there are two categories: 'Health' and 'Ethnicity', both showing 2,700 files with red progress bars.

## Accuracy of information found

NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that Cloud Compliance identifies. You should always validate the information by reviewing the data.

Based on our testing, the table below shows the accuracy of the information that Cloud Compliance finds. We break it down by *precision* and *recall*:

### Precision

The probability that what Cloud Compliance finds has been identified correctly. For example, a precision rate of 90% for personal data means that 9 out of 10 files identified as containing personal information, actually contain personal information. 1 out of 10 files would be a false positive.

### Recall

The probability for Cloud Compliance to find what it should. For example, a recall rate of 70% for personal data means that Cloud Compliance can identify 7 out of 10 files that actually contain personal information in your organization. Cloud Compliance would miss 30% of the data and it won't appear in the dashboard.

Cloud Compliance is in a Controlled Availability release and we are constantly improving the accuracy of our results. Those improvements will be automatically available in future Cloud Compliance releases.

Type	Precision	Recall
Personal data - General	90%-95%	60%-80%
Personal data - Country identifiers	30%-60%	40%-60%
Sensitive personal data	80%-95%	20%-30%
Categories	90%-97%	60%-80%

## What's included in each file list report (CSV file)

From each Investigation page you can download file lists (in CSV format) that include details about the identified files. If there are more than 10,000 results, only the top 10,000 appear in the list.

Each file list includes the following information:

- File name
- Location type
- Working environment
- Storage repository
- Protocol
- File path
- File type
- Category
- Personal information
- Sensitive personal information
- Deletion detection date

A deletion detection date identifies the date that the file was deleted or moved. This enables you to identify when sensitive files have been moved. Deleted files aren't part of the file number count that appears in the dashboard or on the Investigation page. The files only appear in the CSV reports.

## Viewing compliance reports

Cloud Compliance provides reports that you can use to better understand the status of your organization's data privacy program.

By default, the Cloud Compliance dashboard displays compliance data for all working environments and databases. If you want to view reports that contain data for only some of the working environments, [select those working environments](#).



NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that Cloud Compliance identifies. You should always validate the information by reviewing the data.

## Privacy Risk Assessment Report

The Privacy Risk Assessment Report provides an overview of your organization's privacy risk status, as required by privacy regulations such as GDPR and CCPA. The report includes the following information:

### Compliance status

A [severity score](#) and the distribution of data, whether it's non-sensitive, personal, or sensitive personal.

### Assessment overview

A breakdown of the types of personal data found, as well as the categories of data.

### Data subjects in this assessment


The number of people, by location, for which national identifiers were found.

## Generating the Privacy Risk Assessment Report

Go to the Compliance tab to generate the report.

### Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Under **Reports**, click the download icon next to **Privacy Risk Assessment**.



## Result

Cloud Compliance generates a PDF report that you can review and send to other groups as needed.

## Severity score

Cloud Compliance calculates the severity score for the Privacy Risk Assessment Report on the basis of three variables:

- The percentage of personal data out of all data.
- The percentage of sensitive personal data out of all data.
- The percentage of files that include data subjects, determined by national identifiers such as national IDs, Social Security numbers, and tax ID numbers.

The logic used to determine the score is as follows:

Severity score	Logic
0	All three variables are exactly 0%
1	One of the variables are larger than 0%
2	One of the variables are larger than 3%
3	Two of the variables are larger than 3%
4	Three of the variables are larger than 3%
5	One of the variables are larger than 6%
6	Two of the variables are larger than 6%
7	Three of the variables are larger than 6%
8	One of the variables are larger than 15%
9	Two of the variables are larger than 15%
10	Three of the variables are larger than 15%

## PCI DSS Report

The Payment Card Industry Data Security Standard (PCI DSS) Report can help you identify the distribution of credit card information across your files. The report includes the following information:

### Overview

How many files contain credit card information and in which working environments.

### Encryption

The percentage of files containing credit card information that are on encrypted or unencrypted working environments. This information is specific to Cloud Volumes ONTAP.

### Ransomware Protection

The percentage of files containing credit card information that are on working environments that do or don't have ransomware protection enabled. This information is specific to Cloud Volumes ONTAP.

### Retention

The timeframe in which the files were last modified. This is helpful because you shouldn't keep credit card information for longer than you need to process it.

### Distribution of Credit Card Information

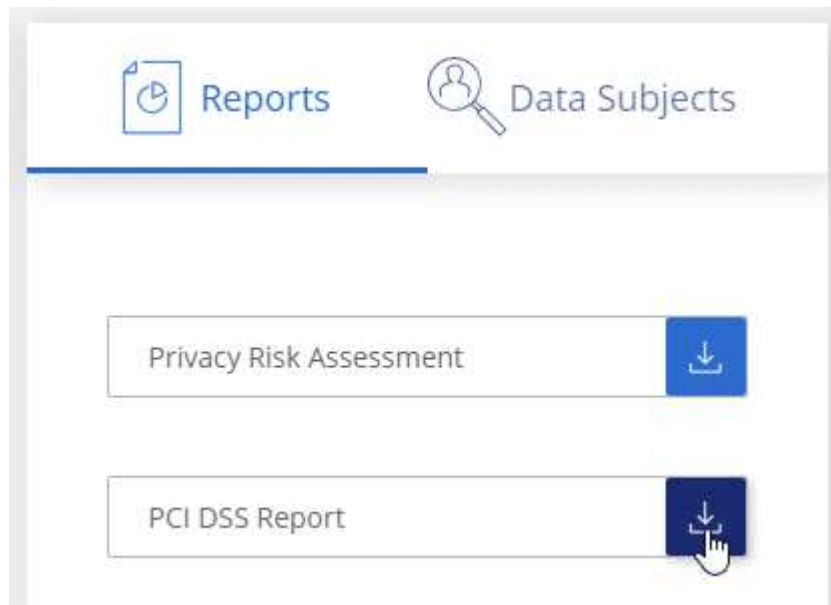
The working environments where the credit card information was found and whether encryption and ransomware protection are enabled.

## Generating the PCI DSS Report

Go to the Compliance tab to generate the report.

### Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Under **Reports**, click the download icon next to **PCI DSS Report**.



### Result

Cloud Compliance generates a PDF report that you can review and send to other groups as needed.

## HIPAA Report

The Health Insurance Portability and Accountability Act (HIPAA) Report can help you identify files containing health information. It is designed to aid in your organization's requirement to comply with HIPAA data privacy laws. The information Cloud Compliance looks for includes:

- Health reference pattern
- ICD-10-CM Medical code
- ICD-9-CM Medical code
- HR – Health category
- Health Application Data category

The report includes the following information:

### Overview

How many files contain health information and in which working environments.

### Encryption

The percentage of files containing health information that are on encrypted or unencrypted working environments. This information is specific to Cloud Volumes ONTAP.

### Ransomware Protection

The percentage of files containing health information that are on working environments that do or don't have ransomware protection enabled. This information is specific to Cloud Volumes ONTAP.

### Retention

The timeframe in which the files were last modified. This is helpful because you shouldn't keep health information for longer than you need to process it.

### Distribution of Health Information

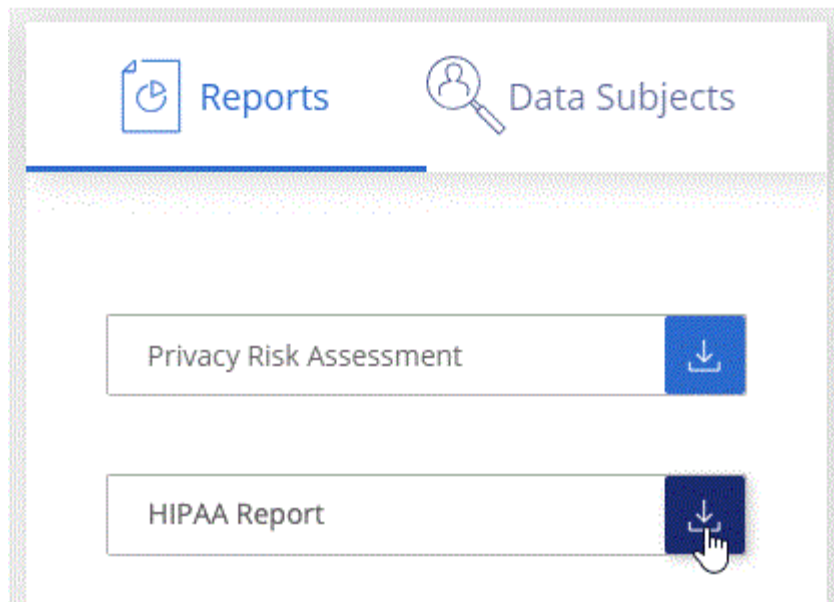
The working environments where the health information was found and whether encryption and ransomware protection are enabled.

## Generating the HIPAA Report

Go to the Compliance tab to generate the report.

### Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Under **Reports**, click the download icon next to **HIPAA Report**.



### Result

Cloud Compliance generates a PDF report that you can review and send to other groups as needed.


### Selecting the working environments for reports

You can filter the contents of the Cloud Compliance dashboard to see compliance data for all working environments and databases, or for just specific working environments.

When you filter the dashboard, Cloud Compliance scopes the compliance data and reports to just those working environments that you selected.

### Steps

1. Click the filter drop-down, select the working environments that you'd like to view data for, and click **View**.



## Responding to a Data Subject Access Request

Respond to a Data Subject Access Request (DSAR) by searching for a subject's full name or known identifier (such as an email address) and then downloading a report. The report is designed to aid in your organization's requirement to comply with GDPR or similar data privacy laws.



NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that Cloud Compliance identifies. You should always validate the information by reviewing the data.

### What is a Data Subject Access Request?

Privacy regulations such as the European GDPR grant data subjects (such as customers or employees) the right to access their personal data. When a data subject requests this information, this is known as a DSAR (data subject access request). Organizations are required to respond to these requests "without undue delay," and at the latest within one month of receipt.



## How can Cloud Compliance help you respond to a DSAR?

When you perform a data subject search, Cloud Compliance finds all of the files that has that person's name or identifier in it. Cloud Compliance checks the latest pre-indexed data for the name or identifier. It doesn't initiate a new scan.

After the search is complete, you can then download the list of files for a Data Subject Access Request report. The report aggregates insights from the data and puts it into legal terms that you can send back to the person.

### Searching for data subjects and downloading reports

Search for the data subject's full name or known identifier and then download a file list report or DSAR report. You can search by [any personal information type](#).

Only English is supported when searching for the names of data subjects. Support for more languages will be added later.

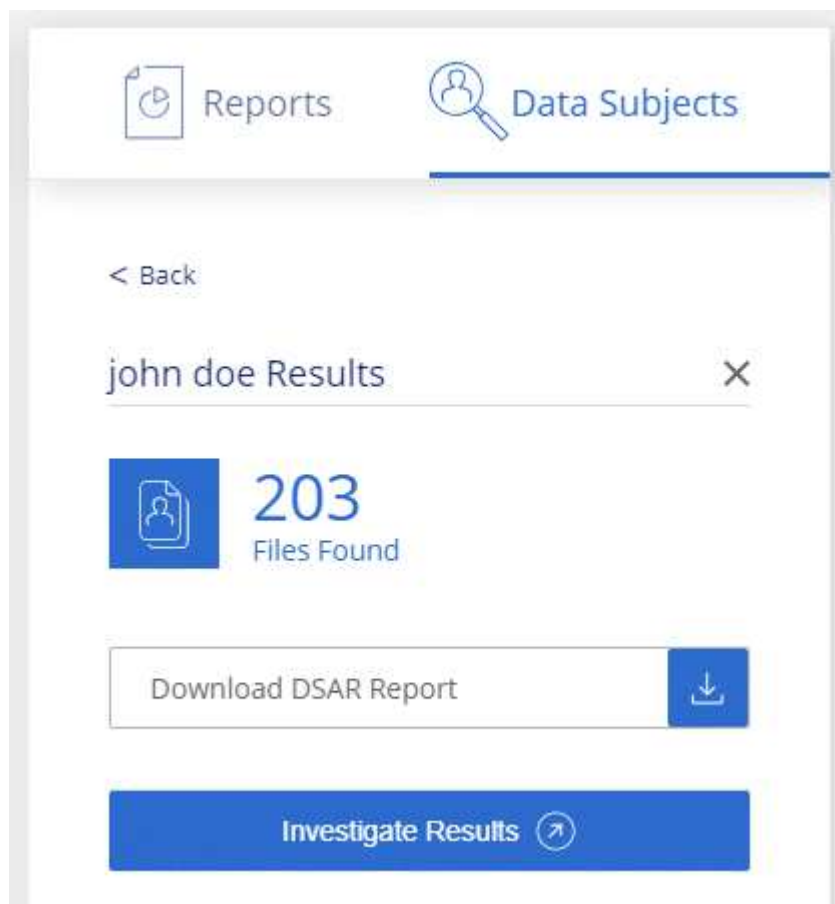


Data subject search is not supported within databases at this time.

#### Steps

1. At the top of Cloud Manager, click **Cloud Compliance**.
2. Click **Data Subjects**.
3. Search for the data subject's full name or known identifier.

Here's an example that shows a search for the name *john doe*:



4. Choose one of the available options:

- **Download DSAR Report:** A formal response to the access request that you can send to the data subject. This report contains automatically-generated information based on data that Cloud Compliance found on the data subject and is designed to be used as a template. You should complete the form and review it internally before sending it to the data subject.
- **Investigate Results:** A page that enables you to investigate the data by searching, sorting, expanding details for a specific file, and by downloading the file list.



If there are more than 10,000 results, only the top 10,000 appear in the file list.


## Disabling Cloud Compliance

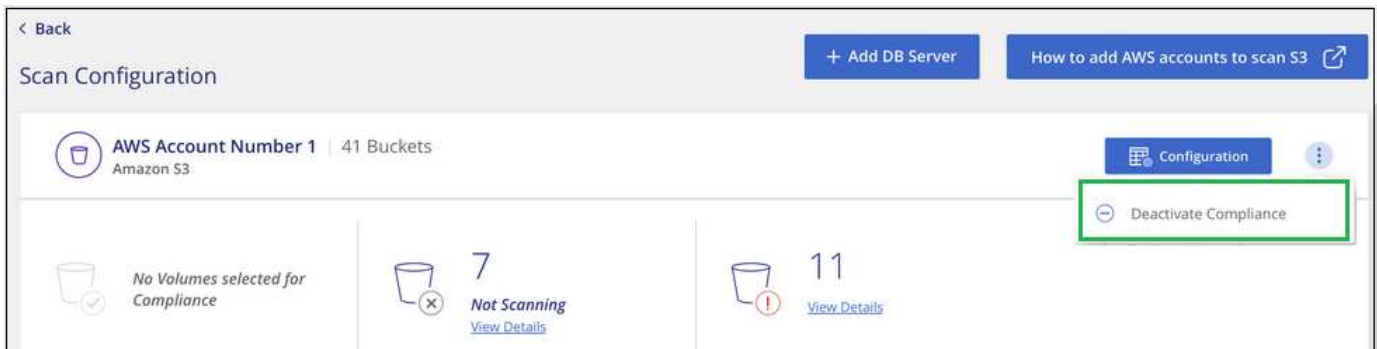
If you need to, you can stop Cloud Compliance from scanning one or more working environments or databases. You can also delete the Cloud Compliance instance if you no longer want to use Cloud Compliance with your working environments.

### Deactivating compliance scans for a working environment

When you deactivate scans, Cloud Compliance no longer scans the data on the system and it removes the indexed compliance insights from the Cloud Compliance instance (the data from the working environment or database itself isn't deleted).

#### Steps

From the *Scan Configuration* page, click the  button in the row for the working environment, and then click **Deactivate Compliance**.



You can also disable compliance scans for a working environment from the Services panel when you select the working environment.

### Deleting the Cloud Compliance instance

You can delete the Cloud Compliance instance if you no longer want to use Cloud Compliance. Deleting the instance also deletes the associated disks where the indexed data resides.

#### Step

1. Go to your cloud provider's console and delete the Cloud Compliance instance.

The instance is named *CloudCompliance* with a generated hash (UUID) concatenated to it. For example:

## Frequently asked questions about Cloud Compliance

This FAQ can help if you're just looking for a quick answer to a question.

### What is Cloud Compliance?

Cloud Compliance is a cloud offering that uses Artificial Intelligence (AI) driven technology to help organizations understand data context and identify sensitive data across your Azure NetApp Files configurations, Cloud Volumes ONTAP systems hosted in AWS or Azure, Amazon S3 buckets, and databases.

Cloud Compliance provides pre-defined parameters (such as sensitive information types and categories) to address new data compliance regulations for data privacy and sensitivity, such as GDPR, CCPA, HIPAA, and more.

### Why should I use Cloud Compliance?

Cloud Compliance can empower you with data to help you:

- Comply with data compliance and privacy regulations.
- Comply with data retention policies.
- Easily locate and report on specific data in response to data subjects, as required by GDPR, CCPA, HIPAA, and other data privacy regulations.

### What are the common use cases for Cloud Compliance?

- Identify Personal Identifiable Information (PII).
- Identify a wide scope of sensitive information as required by GDPR and CCPA privacy regulations.
- Comply with new and upcoming data privacy regulations.

[Learn more about the use cases for Cloud Compliance.](#)

### What types of data can be scanned with Cloud Compliance?

Cloud Compliance supports scanning of unstructured data over NFS and CIFS protocols that are managed by Cloud Volumes ONTAP and Azure NetApp Files. Cloud Compliance can also scan data stored on Amazon S3 buckets.

Additionally, Cloud Compliance can scan databases that are located anywhere - they are not required to be managed by Cloud Manager.

[Learn how scans work.](#)

### Which cloud providers are supported?

Cloud Compliance operates as part of Cloud Manager and currently supports AWS and Azure. This provides your organization with unified privacy visibility across different cloud providers. Support for Google Cloud Platform (GCP) will be added soon.

## How do I access Cloud Compliance?

Cloud Compliance is operated and managed through Cloud Manager. You can access Cloud Compliance features from the **Compliance** tab in Cloud Manager.

## How does Cloud Compliance work?

Cloud Compliance deploys another layer of Artificial Intelligence alongside your Cloud Manager system and storage systems. It then scans the data on volumes, buckets, and databases and indexes the data insights that are found.

[Learn more about how Cloud Compliance works.](#)

## How much does Cloud Compliance cost?

The cost to use Cloud Compliance depends on the amount of data that you're scanning. The first 1 TB of data that Cloud Compliance scans in a Cloud Manager workspace is free. A subscription to the AWS or Azure Marketplace is required to continue scanning data after that point. See [pricing](#) for details.

## How often does Cloud Compliance scan my data?

Data changes frequently, so Cloud Compliance scans your data continuously with no impact to your data. While the initial scan of your data might take longer, subsequent scans only scan the incremental changes, which reduces system scan times.

[Learn how scans work.](#)

## Does Cloud Compliance offer reports?

Yes. The information offered by Cloud Compliance can be relevant to other stakeholders in your organizations, so we enable you to generate reports to share the insights.

The following reports are available for Cloud Compliance:

### Privacy Risk Assessment report

Provides privacy insights from your data and a privacy risk score. [Learn more.](#)

### Data Subject Access Request report

Enables you to extract a report of all files that contain information regarding a data subject's specific name or personal identifier. [Learn more.](#)

### PCI DSS report

Helps you identify the distribution of credit card information across your files. [Learn more.](#)

### HIPAA report

Helps you identify the distribution of health information across your files. [Learn more.](#)

### Reports on a specific information type

Reports are available that include details about the identified files that contain personal data and sensitive personal data. You can also see files broken down by category and file type. [Learn more.](#)

## What type of instance or VM is required for Cloud Compliance?

- In Azure, Cloud Compliance runs on a Standard\_D16s\_v3 VM with a 512 GB disk.
- In AWS, Cloud Compliance runs on an m5.4xlarge instance with a 500 GB GP2 disk.

In regions where m5.4xlarge isn't available, Cloud Compliance runs on an m4.4xlarge instance instead.



Changing or resizing the instance/VM type isn't supported. You need to use the default size that's provided.

[Learn more about how Cloud Compliance works.](#)

## Does scan performance vary?

Scan performance can vary based on the network bandwidth and the average file size in your cloud environment.

## Which file types are supported?

Cloud Compliance scans all files for category and metadata insights and displays all file types in the file types section of the dashboard.

When Cloud Compliance detects Personal Identifiable Information (PII), or when it performs a DSAR search, only the following file formats are supported:

.PDF, .DOCX, .DOC, .PPTX, .XLS, .XLSX, .CSV, .TXT, .RTF, and .JSON.

## How do I enable Cloud Compliance?

First you need to deploy an instance of Cloud Compliance in Cloud Manager. Once the instance is running, you can enable it on existing working environments and databases from the **Compliance** tab or by selecting a specific working environment.

[Learn how to get started.](#)



Activating Cloud Compliance results in an immediate initial scan. Compliance results display shortly after.

## How do I disable Cloud Compliance?

You can disable Cloud Compliance from the Working Environments page after you select an individual working environment.

[Learn more.](#)



To completely remove the Cloud Compliance instance, you can manually remove the Cloud Compliance instance from your cloud provider's portal.

## What happens if data tiering is enabled on Cloud Volumes ONTAP?

You might want to enable Cloud Compliance on a Cloud Volumes ONTAP system that tiers cold data to object storage. If data tiering is enabled, Cloud Compliance scans all of the data—data that's on disks and cold data

tiered to object storage.

The compliance scan doesn't heat up the cold data—it stays cold and tiered to object storage.

### **Can I use Cloud Compliance to scan on-premise ONTAP storage?**

Scanning the data directly from an on-premises ONTAP working environment isn't supported. But you can scan your on-premises ONTAP data by replicating the on-prem NFS or CIFS data to a Cloud Volumes ONTAP working environment and then activating compliance on those volumes. We're planning to support Cloud Compliance with additional cloud offerings such as Cloud Volumes Service.

[Learn more.](#)

### **Can Cloud Compliance send notifications to my organization?**

No, but you can download status reports that you can share internally in your organization.

### **Can I customize the service to my organization's need?**

Cloud Compliance provides out-of-the-box insights to your data. These insights can be extracted and used for your organization's needs.

### **Can I limit Cloud Compliance information to specific users?**

Yes, Cloud Compliance is fully integrated with Cloud Manager. Cloud Manager users can only see information for the working environments they are eligible to view according to their workspace privileges.

Additionally, if you want to allow certain users to just view Cloud Compliance scan results without having the ability to manage Cloud Compliance settings, you can assign those users the *Cloud Compliance Viewer* role.

[Learn more.](#)


# Enable real-time global file sharing

## Learn about Global File Cache

NetApp Global File Cache enables you to consolidate silos of distributed file servers into one cohesive global storage footprint in the public cloud. This creates a globally accessible file system in the cloud that all remote locations can use as if they were local.

### Overview

Implementing Global File Cache results in a single, centralized storage footprint, versus a distributed storage architecture that requires local data management, backup, security management, storage, and infrastructure footprint in each location.



### Features

Global File Cache enables the following features:

- Consolidate and centralize your data into the public cloud and leverage the scalability and performance from enterprise-grade storage solutions
- Create a single set of data for users globally and leverage intelligent file caching to improve global data access, collaboration, and performance
- Rely on a self-sustaining, self-managing cache, and eliminate full data copies and backups. Utilize local file caching for active data and cut storage costs
- Transparent access from branch locations through a global namespace with real time central file locking


See more about Global File Cache features and use cases [here](#).

### Global File Cache components

Global File Cache consists of the following components:

- Global File Cache Management Server
- Global File Cache Core
- Global File Cache Edge (deployed in your remote locations)

The Global File Cache Core instance mounts to your corporate file shares hosted on your backend storage platform of choice (such as Cloud Volumes ONTAP, Cloud Volumes Service, and Azure NetApp Files) and creates the Global File Cache Fabric that provides the ability to centralize and consolidate unstructured data into a single set of data, whether it resides on one or multiple storage platforms in the public cloud.



## Supported storage platforms

The supported storage platforms for Global File Cache differ depending on the deployment option you select.

### Automated deployment options

Global File Cache is supported with the following types of working environments when deployed using Cloud Manager:

- Cloud Volumes ONTAP in Azure
- Cloud Volumes ONTAP in AWS

This configuration lets you deploy and manage the entire Global File Cache server-side deployment, including Global File Cache Management Server and Global File Cache Core, from within Cloud Manager.

### Manual deployment options


Global File Cache configurations are also supported with Cloud Volumes ONTAP, Cloud Volumes Service, or Azure NetApp Files installed on Microsoft Azure, Google Cloud Platform, or Amazon Web Services public cloud storage infrastructure. On-premises solutions are also available on NetApp AFF and FAS platforms. In these installations the Global File Cache server-side components must be configured and deployed manually, not using Cloud Manager.

See the [NetApp Global File Cache User Guide](#) for details.

## How Global File Cache works

Global File Cache creates a software fabric that caches active data sets in remote offices globally. As a result, business users are guaranteed transparent data access and optimal performance on a global scale.





The topology referenced in this example is a hub and spoke model, whereby the network of remote offices/locations is accessing one common set of data in the cloud. The key points of this example are:

- Centralized data store:
  - Enterprise public cloud storage platform, such as Cloud Volumes ONTAP
- Global File Cache Fabric:
  - Extension of the central data store to the remote locations
  - Global File Cache Core instance, mounting to corporate file shares (SMB).
  - Global File Cache Edge instances running in each remote location.
  - Presents a virtual file share in each remote location that provides access to central data.
  - Hosts the Intelligent File Cache on a custom-sized NTFS volume (D:\).
- Network configuration:
  - Multiprotocol Label Switching (MPLS), ExpressRoute, or VPN connectivity
- Integration with customer's Active Directory domain services.
- DFS namespace for the use of a global namespace (recommended).

## Cost

The cost to use Global File Cache depends on the type of installation you have chosen.

- All installations require that you deploy one or more volumes in the cloud (Cloud Volumes ONTAP, Cloud Volumes Service, or Azure NetApp Files). This results in charges from the selected cloud provider.
- All installations also require that you deploy two or more virtual machines (VMs) in the cloud. This results in charges from the selected cloud provider.
  - Global File Cache Management Server:

In Azure, this runs on a D2s\_V3 or equivalent (2 vCPU/8GB RAM) VM with 127GB premium SSD

In AWS, this runs on a m4.large or equivalent (2 vCPU/8GB RAM) instance with 127GB general purpose SSD

- Global File Cache Core:

In Azure, this runs on a D4s\_V3 or equivalent (4 vCPU/16GB RAM) VM with 127GB premium SSD

In AWS, this runs on a m4.xlarge or equivalent (4 vCPU/16GB RAM) instance with 127GB general purpose SSD

- When installed with Cloud Volumes ONTAP in Azure or AWS (the supported configurations deployed completely through Cloud Manager), there is a charge of \$3,000 per site (for each Global File Cache Edge instance), per year.
- When installed using the manual deployment options the pricing is different. To see a high-level estimate of costs, see [Calculate Your Savings Potential](#) or consult your Global File Cache Solutions Engineer to discuss the best options for your enterprise deployment.

## Licensing

Global File Cache includes a software-based License Management Server (LMS), which allows you to consolidate your license management and deploy licenses to all Core and Edge instances using an automated mechanism.

When you deploy your first Core instance in the datacenter or cloud, you can choose to designate that instance as the LMS for your organization. This LMS instance is configured once, connects to the subscription service (over HTTPS) and validates your subscription using the customer ID provided by our support/operations department upon enablement of the subscription. After you have made this designation, you associate your Edge instances with the LMS by providing your customer ID and the IP address of the LMS instance.

When you purchase additional Edge licenses or renew your subscription, our support/operations department updates the license details, for example, the number of sites or subscription end date. After the LMS queries the subscription service, the license details are automatically updated on the LMS instance and will apply to your GFC Core and Edge instances.

See the [NetApp Global File Cache User Guide](#) for additional details about licensing.

## Limitations

- The version of Global File Cache supported within Cloud Manager requires that the backend storage platform used as your central storage must be a working environment where you have deployed a Cloud Volumes ONTAP single node or HA pair in Azure or AWS.

Other storage platforms and other cloud providers are not supported at this time using Cloud Manager, but can be deployed using legacy deployment procedures.

These other configurations, for example, Global File Cache using Cloud Volumes ONTAP, Cloud Volumes Service, and Azure NetApp Files on Microsoft Azure, Google Cloud, and AWS continue to be supported using the legacy procedures. See [Global File Cache overview and onboarding](#) for details.

## Before you begin to deploy Global File Cache

There are many requirements you need to be aware of before you begin to deploy Global File Cache in the cloud and in your remote offices.

## Global File Cache Core design considerations

Depending on your requirements, you may need to deploy one or multiple Global File Cache Core instances to create the Global File Cache Fabric. The Core instance is designed to act as a traffic cop between your distributed Global File Cache Edge instances and the data center file server resources, for example, file shares, folders, and files.

When you are designing your Global File Cache deployment you need to determine what's right for your environment in terms of scale, availability of resources, and in terms of redundancy. Global File Cache Core can be deployed in the following ways:

- GFC Core stand-alone instance
- GFC Core Load Distributed design (Cold Standby)

See [Sizing guidelines](#) to understand the maximum number of Edge instances and total users that each configuration can support:

Consult your Global File Cache Solutions Engineer to discuss the best options for your enterprise deployment.

### Sizing guidelines

There are a few sizing guideline ratios that you need to keep in mind when configuring the initial system. You should revisit these ratios after some usage history has accumulated to make sure you are using the system optimally. These include:

- Global File Cache Edges/Core ratio
- Distributed users/Global File Cache Edge ratio
- Distributed users/Global File Cache Core ratio

### Number of Edge Instances per Core Instance

Our guidelines recommend up to 10 Edge instances per Global File Cache Core instance, with a maximum of 20 Edges per Global File Cache Core instance. This is dependent to a significant degree upon the type and mean file size of the most common workload. In some cases, with more common workloads you can add more Edge instances per Core, but in these cases you should contact NetApp Support to correctly size the number of Edge and Core instances depending on the types and sizes of the file sets.



You can leverage multiple Global File Cache Edge and Core instances simultaneously to scale out your infrastructure depending on the requirements.

### Number of concurrent users per Edge instance

Global File Cache Edge handles the heavy lifting in terms of caching algorithms and file-level differencing. A single Global File Cache Edge instance can serve up to 400 users per dedicated physical Edge instance, and up to 200 users for dedicated virtual deployments. This is dependent to a significant degree upon the type and mean file size of the most common workload. For larger collaborative file types, guide towards 50% of the maximum users per Global File Cache Edge lower boundary (depending on physical or virtual deployment). For more common Office items with a mean file size <1MB, guide towards the 100% users per Global File Cache Edge upper boundary (depending on physical or virtual deployment).



Global File Cache Edge detects whether it is running on a virtual or physical instance and it will limit the number of SMB connections to the local virtual file share to the maximum of 200 or 400 concurrent connections.

## Number of concurrent users per Core instance

The Global File Cache Core instance is extremely scalable, with a recommended concurrent user count of 3,000 users per Core. This is dependent to a significant degree upon the type and mean file size of the most common workload.

Consult your Global File Cache Solutions Engineer to discuss the best options for your enterprise deployment.

## Prerequisites

The prerequisites described in this section are for the components installed in the cloud: the Global File Cache Management Server and the Global File Cache Core.

Global File Cache Edge prerequisites are described [here](#).

## Cloud Manager instance

When using Cloud Volumes ONTAP for Azure as your storage platform, ensure that Cloud Manager has permissions as shown in the latest [Cloud Manager policy for Azure](#).

Newly created instances will have all the required permissions by default. If you deployed your instance prior to version 3.8.7 (August 3, 2020), then you will need to add these items.

```
"Microsoft.Resources/deployments/operationStatuses/read",  
"Microsoft.Insights/Metrics/Read",  
"Microsoft.Compute/virtualMachines/extensions/write",  
"Microsoft.Compute/virtualMachines/extensions/read",  
"Microsoft.Compute/virtualMachines/extensions/delete",  
"Microsoft.Compute/virtualMachines/delete",  
"Microsoft.Network/networkInterfaces/delete",  
"Microsoft.Network/networkSecurityGroups/delete",  
"Microsoft.Resources/deployments/delete",
```

## Storage platform (volumes)

The back-end storage platform – in this case, your deployed Cloud Volumes ONTAP instance - should present SMB file shares. Any shares that will be exposed through Global File Cache must allow the Everyone group Full Control at the share level, while restricting permissions through NTFS permissions.

If you have not set up at least one SMB file share on the Cloud Volumes ONTAP instance, then you need to have the following information ready so you can configure this information during installation:

- Active Directory domain name, name server IP address, Active Directory admin credentials.
- The name and size of the volume you want to create, the name of the aggregate on which the volume will be created, and the share name.

We recommend that the volume is large enough to accommodate the total data set for the application along with the ability to scale accordingly as the data set grows. If you have multiple aggregates in the working environment, see [Managing existing aggregates](#) to determine which aggregate has the most available space for the new volume.

## Global File Cache Management Server

This Global File Cache Management Server requires external access over HTTPS (TCP port 443) to connect to the cloud provider subscription service and to access these URLs:

- <https://talonazuremicroservices.azurewebsites.net>
- <https://talonlicensing.table.core.windows.net>

This port must be excluded from any WAN optimization devices or firewall restriction policies for the Global File Cache software to operate properly.

The Global File Cache Management Server also requires a unique (geographical) NetBIOS name for the instance (such as GFC-MS1).



One Management Server can support multiple Global File Cache Core instances deployed in different working environments. When deployed from Cloud Manager, each working environment has its own separate backend storage and would not contain the same data.

## Global File Cache Core

This Global File Cache Core listens on TCP port range 6618-6630. Depending on your firewall or Network Security Group (NSG) configuration you may need to explicitly allow access to these ports through Inbound Port Rules. Also, these ports must be excluded from any WAN optimization devices or firewall restriction policies for the Global File Cache software to operate properly.

The Global File Cache Core requirements are:

- A unique (geographical) NetBIOS name for the instance (such as GFC-CORE1)
- Active Directory domain name
  - Global File Cache instances should be joined to your Active Directory domain.
  - Global File Cache instances should be managed in a Global File Cache specific Organizational Unit (OU) and excluded from inherited company GPOs.
- Service account. The services on this Global File Cache Core run as a specific domain user account. This account, also known as the Service Account, must have the following privileges on each of the SMB servers that will be associated with the Global File Cache Core instance:
  - The provisioned Service Account must be a domain user.

Depending on the level of restrictions and GPOs in the network environment, this account might require domain admin privileges.

- It must have "Run as a Service" privileges.
- The password should be set to "Never Expire".
- The account option "User Must Change Password at Next Logon" should be DISABLED (unchecked).
- It must be a member of the back-end file server Built-in Backup Operators group (this is automatically enabled when deployed through Cloud Manager).

## License Management Server

- The Global File Cache License Management Server (LMS) should be configured on a Microsoft Windows Server 2016 Standard or Datacenter edition or Windows Server 2019 Standard or Datacenter edition, preferably on the Global File Cache Core instance in the datacenter or cloud.
- If you require a separate Global File Cache LMS instance, you need to install the latest Global File Cache software installation package on a pristine Microsoft Windows Server instance.
- The LMS instance needs to be able to connect to the subscription service (Azure Services / public internet) using HTTPS (TCP port 443).
- The Core and Edge instances need to connect to the LMS instance using HTTPS (TCP port 443).

## Networking

- Firewall: TCP ports should be allowed between Global File Cache Edge and Core instances.
- Global File Cache TCP Ports: 443 (HTTPS), 6618–6630.
- Network optimization devices (such as Riverbed Steelhead) must be configured to pass-thru Global File Cache specific ports (TCP 6618-6630).

# Getting started

You use Cloud Manager to deploy the Global File Cache Management Server and Global File Cache Core software in the working environment.

## Enable Global File Cache using Cloud Manager

In this configuration you will deploy the Global File Cache Management Server and Global File Cache Core in the same working environment where you created your Cloud Volumes ONTAP system using Cloud Manager.

Watch [this video](#) to see the steps from start to finish.

### Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details:



#### Deploy Cloud Volumes ONTAP

Deploy Cloud Volumes ONTAP in Azure or AWS and configure SMB file shares. For more information, see [Launching Cloud Volumes ONTAP in Azure](#) or [Launching Cloud Volumes ONTAP in AWS](#).



#### Deploy the Global File Cache Management Server

Deploy an instance of the Global File Cache Management Server in the same working environment as the instance of Cloud Volumes ONTAP.



#### Deploy the Global File Cache Core

Deploy an instance, or multiple instances, of the Global File Cache Core in the same working environment as the instance of Cloud Volumes ONTAP and join it to your Active Directory domain.



#### **License Global File Cache**

Configure the Global File Cache License Management Server (LMS) service on a Global File Cache Core instance. You will need your NSS Credentials or a customer ID provided by NetApp to activate your subscription.



#### **Deploy the Global File Cache Edge instances**

See [Deploying Global File Cache Edge instances](#) to deploy the Global File Cache Edge instances in each remote location. This step is not done using Cloud Manager.

### **Deploy Cloud Volumes ONTAP as your storage platform**

In the current release, Global File Cache supports Cloud Volumes ONTAP deployed in Azure or AWS. For detailed prerequisites, requirements, and deployment instructions, see [Launching Cloud Volumes ONTAP in Azure](#) or [Launching Cloud Volumes ONTAP in AWS](#).

Note the following additional Global File Cache requirement:

- You should configure SMB file shares on the instance of Cloud Volumes ONTAP.

If no SMB file shares are set up on the instance, then you are prompted to configure the SMB shares during the installation of the Global File Cache components.

### **Enable Global File Cache in your working environment**

The Global File Cache wizard walks you through the steps to deploy the Global File Cache Management Server instance and the Global File Cache Core instance, as highlighted below.



Cloud Manager Working Environments Compliance Replication K8s Backup Sync On-Prem Tiering Global File Cache

Enable GFC 1 Overview 2 Enable GFC Service 3 GFC Service (Setup) 4 Deploy GFC Core 5 GFC Core (Setup)

### Thank you for enabling NetApp Global File Cache

Global File Cache allows distributed enterprises to securely consolidate silos of file servers into one cohesive global storage footprint in the public cloud, which streamlines overall IT management, significantly cuts costs and boosts business productivity on a global scale.

From a high-level perspective, we will guide you through the process of deploying the GFC Core Instance in the public cloud and provide you with the instructions to start off your branch office deployment, the Edge Instance.

Continue

Cloud Manager 3.8.7 Build: 1 Jul 16, 2020 09:53:22 am UTC Help API API documentation

## Steps

1. Select the working environment where you deployed Cloud Volumes ONTAP.
2. In the Services panel, click **Enable GFC**.



3. Read the Overview page and click **Continue**.
4. If no SMB shares are available on the Cloud Volumes ONTAP instance, you are prompted to enter the SMB Server and SMB Share details to create the share now. For details about the SMB configuration, see [Storage platform](#).

When finished, click **Continue** to create the SMB share.



### SMB Setup

<p><b>SMB Server</b></p> <p>Active Directory Domain  <input type="text" value="gfc.netapp.com"/></p> <p>Name Server IP Address  <input type="text" value="10.0.2.4"/></p> <p>Active Directory Admin User  <input type="text" value="cvoadmin"/></p> <p>Active Directory Admin Password  <input type="password" value="*****"/></p>	<p><b>SMB Share</b></p> <p>Volume Name                      Volume Size(GB)  <input type="text" value="Enter Volume Name"/>    <input type="text"/></p> <p>Select Aggregate  <input type="text" value="Select Aggregate"/></p> <p>Share Name  <input type="text" value="Enter Share Name"/></p> <hr/> <p>Thin provisioning                      Enabled ⓘ</p> <p>Deduplication                              Enabled ⓘ</p>
--	---

- On the Global File Cache Service page, enter the number of Global File Cache Edge instances you plan to deploy, and then make sure your system meets the requirements for Network Configuration and Firewall Rules, Active Directory settings, and Antivirus exclusions. See [Prerequisites](#) for more details.

## Enable Global File Cache Service

### Licensing Global File Cache:

Once you've completed this deployment process, you will need your NSS Credentials to activate your subscription. If you haven't purchased or received your NetApp Global File Cache licenses, which are available as an Edge-based license, they can be purchased through your NetApp Partner or NetApp Sales Representative.

How many edge instances are you planning to deploy?

### Before you begin:

Here are the most important requirements for your environment before you can deploy the NetApp Global File Cache solution:

Configure the required Network Configuration and Firewall Rules for Global File Cache



Create a "Service Account" in your Active Directory domain: GFC.NETAPP.COM



Update Antivirus Exclusions for your Windows Server infrastructure by committing the required exclusions to your Antivirus services



For more information on all the solution requirements [Click Here](#)

Continue

6. After you have verified that the requirements have been met, or that you have the information to meet these requirements, click **Continue**.
7. Enter the admin credentials you will use to access to the Global File Cache Management Server VM and click **Enable GFC Service**. For Azure you enter the credentials as a user name and password; for AWS you select the appropriate key pair. You can change the VM/instance name if you want.

## Global File Cache Service (Setup)

### Information

Subscription Name	OCCM Dev
Azure Region	eastus
VNet	Vnet1
Subnet	Subnet2
Resource Group	occm_group_eastus

### Credentials & Virtual Machine

Local Admin Name

Local Admin Password

VM Name

8. After the Global File Cache Management Service is successfully deployed, click **Continue**.
9. For the Global File Cache Core, enter the admin user credentials to join the Active Directory domain, and the service account user credentials. Then click **Continue**.
  - The Global File Cache Core instance must be deployed in the same Active Directory domain as the Cloud Volumes ONTAP instance.
  - The service account is a domain user and it is part of the BUILTIN\Backup Operators group on the Cloud Volumes ONTAP instance.

## Deploy Global File Cache Core

### Active Directory and Admin Credentials

Provide administrative credentials to join the GFC Core instance to the Active Directory domain

Join Active Directory Domain ⓘ

Admin User ⓘ

Admin Password ⓘ

### Account User Credentials

Provide Service Account credentials

Service Account User ⓘ

Service Account Password ⓘ

10. Enter the admin credentials you will use to access to the Global File Cache Core VM and click **Deploy GFC Core**. For Azure you enter the credentials as a user name and password; for AWS you select the appropriate key pair. You can change the VM/instance name if you want.

## Global File Cache Core (Setup)

### Information

Subscription Name	Subscription_1234567891234...
Region	East US   Virginia
VNet	VNet_1234567
Subnet	10.0.0.0/24
Resource Group	Resource Group 1

### Credentials & Virtual Machine

Local Admin Name

Local Admin Password

VM Name

Local Admin Name & Password are inherited from the Global File Cache Management Service. The Virtual Machine Name is associated to your Cloud Manager Account

11. After the Global File Cache Core is successfully deployed, click **Go to Dashboard**.

Global File Cache

**Global File Cache Management Instance**

	www.working-environment-1.com <small>Hostname</small>	ON <small>Status</small>
141.226.210.219 <small>IP Address</small>	East US <small>Region</small>	VNet1 <small>VNet</small>
10.10.10.10/24 <small>Subnet</small>	RGName <small>Resource Group</small>	26% <small>CPU Utilization</small>

**1 Working Environment**

	Working Environment_1 <small>Name</small>	High Availability <small>Type</small>	ON <small>Status</small>	2 <small>Core Instances</small>	<input style="background-color: #0070C0; color: white; padding: 5px 10px; border: none; cursor: pointer;" type="button" value="Add Core Instance"/>						
<p><b>Instance Core 1</b>   ON</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 15%; text-align: center;">www.working-environment-1.com <small>Hostname</small></td> <td style="width: 15%; text-align: center;">141.226.210.219 <small>IP Address</small></td> <td style="width: 10%; text-align: center;">26% <small>CPU Utilization</small></td> <td style="width: 15%; text-align: center;">2.5 TB <small>Network Inbound</small></td> <td style="width: 15%; text-align: center;">2.5 TB <small>Network Outbound</small></td> <td style="width: 20%; text-align: right;"> <input style="border: 1px solid #ccc; padding: 5px 10px; border-radius: 5px; cursor: pointer;" type="button" value="Deploy GFC Edge"/> </td> </tr> </table>						www.working-environment-1.com <small>Hostname</small>	141.226.210.219 <small>IP Address</small>	26% <small>CPU Utilization</small>	2.5 TB <small>Network Inbound</small>	2.5 TB <small>Network Outbound</small>	<input style="border: 1px solid #ccc; padding: 5px 10px; border-radius: 5px; cursor: pointer;" type="button" value="Deploy GFC Edge"/>
www.working-environment-1.com <small>Hostname</small>	141.226.210.219 <small>IP Address</small>	26% <small>CPU Utilization</small>	2.5 TB <small>Network Inbound</small>	2.5 TB <small>Network Outbound</small>	<input style="border: 1px solid #ccc; padding: 5px 10px; border-radius: 5px; cursor: pointer;" type="button" value="Deploy GFC Edge"/>						

The Dashboard shows that the Management Server instance and the Core instance are both **On** and working.

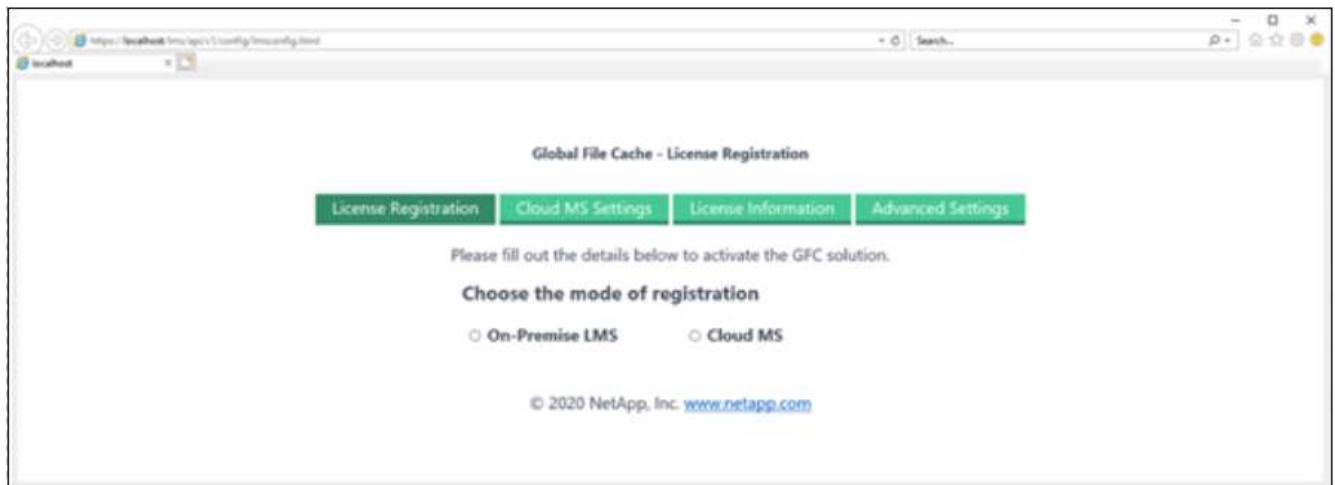
## License your Global File Cache installation

Before you can use Global File Cache, you need to configure the Global File Cache License Management Server (LMS) service on a Global File Cache Core instance. You will need your NSS Credentials or a customer ID provided NetApp to activate your subscription.

In this example, we will configure the LMS service on a Core instance that you just deployed in the public cloud. This is a one-time process that sets up your LMS service.

### Steps

1. Open the Global File Cache License Registration page on the Global File Cache Core (the Core you are designating as your LMS service) using the following URL. Replace `<ip_address>` with the IP address of the Global File Cache Core:  
[https://<ip\\_address>/lms/api/v1/config/lmsconfig.html](https://<ip_address>/lms/api/v1/config/lmsconfig.html)
2. Click “Continue to this website (not recommended)” to continue. A page that allows you to configure the LMS, or check existing license information, is displayed.



3. Choose the mode of registration by selecting “On-Premise LMS” or “Cloud MS”.
  - “On-Premise LMS” is used for existing or trial customers who have received a Customer ID through NetApp Support.
  - “Cloud MS” is used for customers who have purchased NetApp Global File Cache Edge licenses from NetApp or its certified partners and have their NetApp credentials.
4. For Cloud MS, click **Cloud MS**, enter your NSS Credentials, and click **Submit**.

The screenshot shows the 'Global File Cache - License Registration' page. At the top, there are four tabs: 'License Registration', 'Cloud MS Settings', 'License Information', and 'Advanced Settings'. Below the tabs, there are two radio buttons: 'SPN Information' (unselected) and 'NSS Credentials' (selected). Under 'NSS Credentials', there are two input fields: 'NSS username:' and 'NSS password:'. Below these fields is a checkbox labeled 'Update'. At the bottom right of the form is a blue 'SUBMIT' button.

5. For On-Premise LMS, click **On-Premise LMS**, enter your Customer ID, and click **Register LMS**.

The screenshot shows the 'Global File Cache - License Registration' page. At the top, there are four tabs: 'License Registration', 'Cloud MS Settings', 'License Information', and 'Advanced Settings'. Below the tabs, there is a message: 'Please fill out the details below to activate the GFC solution.' Underneath, there is a section titled 'Choose the mode of registration' with two radio buttons: 'On-Premise LMS' (selected) and 'Cloud MS' (unselected). Below the radio buttons is a text input field labeled 'Customer ID:' containing the text 'XYZINC|'. At the bottom right of the form is a blue 'REGISTER LMS' button.

### What's Next?

If you have determined that you need to deploy multiple Global File Cache Cores to support your configuration, click **Add Core Instance** from the Dashboard and follow the deployment wizard.


After you have completed your Core deployment, you need to [deploy the Global File Cache Edge instances](#) in each of your remote offices.

### Deploy additional Core instances

If your configuration requires more than one Global File Cache Core to be installed because of a large number of Edge instances, you can add another Core to the working environment.

When deploying Edge instances, you will configure some to connect to the first Core and others to the second Core. Both Core instances access the same backend storage (your Cloud Volumes ONTAP instance) in the working environment.

1. From the Global File Cache Dashboard, click **Add Core Instance**.



2. Enter the admin user credentials to join the Active Directory domain, and the service account user credentials. Then click **Continue**.

- The Global File Cache Core instance must be in the same Active Directory domain as the Cloud Volumes ONTAP instance.
- The service account is a domain user and it is part of the BUILTIN\Backup Operators group on the Cloud Volumes ONTAP instance.

The screenshot shows a form titled 'Deploy Global File Cache Core'. It is divided into two columns: 'Active Directory and Admin Credentials' and 'Account User Credentials'. Under 'Active Directory and Admin Credentials', there are three input fields: 'Join Active Directory Domain' (with a sub-field 'Active Directory Domain'), 'Admin User' (with 'Enter Admin User'), and 'Admin Password' (with 'Enter Admin Password'). Under 'Account User Credentials', there are two input fields: 'Service Account User' (with 'Enter Service Account User') and 'Service Account Password' (with 'Enter Service Account Password'). A blue 'Continue' button is at the bottom center.

3. Enter the admin credentials you will use to access to the Global File Cache Core VM and click **Deploy GFC Core**. For Azure you enter the credentials as a user name and password; for AWS you select the appropriate key pair. You can change the VM name if you want.



## Global File Cache Core (Setup)

### Information

Subscription Name	Subscription_1234567891234...
Region	East US   Virginia
VNet	VNet_1234567
Subnet	10.0.0.0/24
Resource Group	Resource Group 1

### Credentials & Virtual Machine

Local Admin Name

Local Admin Password

VM Name

Local Admin Name & Password are inherited from the Global File Cache Management Service. The Virtual Machine Name is associated to your Cloud Manager Account

4. After the Global File Cache Core is successfully deployed, click **Go to Dashboard**.

**1 Working Environment**

	Working Environment_1 <small>Name</small>	High Availability <small>Type</small>	ON <small>Status</small>	2 <small>Core Instances</small>	<input style="background-color: #0070C0; color: white; padding: 5px 10px; border: none;" type="button" value="Add Core Instance"/>
<b>Instance Core 1</b>   ON					
www.working-environment-1.com <small>Hostname</small>	141.226.210.219 <small>IP Address</small>	26% <small>CPU Utilization</small>	2.5 TB <small>Network Inbound</small>	2.5 TB <small>Network Outbound</small>	<input style="border: 1px solid #ccc; padding: 5px 10px;" type="button" value="Deploy GFC Edge"/>
<b>Instance Core 1</b>   ON					
www.working-environment-1.com <small>Hostname</small>	141.226.210.219 <small>IP Address</small>	26% <small>CPU Utilization</small>	2.5 TB <small>Network Inbound</small>	2.5 TB <small>Network Outbound</small>	<input style="border: 1px solid #ccc; padding: 5px 10px;" type="button" value="Deploy GFC Edge"/>

The Dashboard reflects the second Core instance for the working environment.

## Before you begin to deploy Global File Cache Edge instances

There are many requirements you need to be aware of before you begin to install Global File Cache Edge software in your remote offices.

## Download required resources

Download the Global File Cache Virtual Templates you are planning to use in your branch offices, the software installation package, and additional reference documentation:

- Windows Server 2016 Virtual Template:

[Windows Server 2016 .OVA including NetApp GFC \(VMware VSphere 6.5+\)](#)  
[Windows Server 2016 .VHDX including NetApp GFC \(Microsoft Hyper-v\)](#)

- Windows Server 2019 Virtual Template:

[Windows Server 2019 .OVA including NetApp GFC \(VMware VSphere 6.5+\)](#)  
[Windows Server 2019 .VHDX including NetApp GFC \(Microsoft Hyper-v\)](#)

- Global File Cache Edge Software:

[NetApp GFC Software \(.EXE\)](#)

- Global File Cache Documentation:

[NetApp Global File Cache User Guide](#)

## Designing and deploying Global File Cache Edge

Depending on your requirements, you might need to deploy one or multiple Global File Cache Edge instances based on the concurrent user sessions in a branch office. The Edge instance presents the virtual file share to the end users within the branch office, which has been transparently extended from the associated Global File Cache Core instance. The Global File Cache Edge should contain a D:\ NTFS volume, which contains the cached files within the branch office.



For the Global File Cache Edge, it is important to understand the [sizing guidelines](#). This will assist you in making the correct design for your Global File Cache deployment. You would also need to determine what's right for your environment in terms of scale, availability of resources, and in terms of redundancy.

### Global File Cache Edge instance

When deploying a Global File Cache Edge instance, you need to provision a single VM, either by deploying Windows Server 2016 Standard or Datacenter Edition, or Windows Server 2019 Standard or Datacenter Edition, or using the Global File Cache .OVA or .VHD template, which includes the Windows Server operating system of choice and Global File Cache software.

#### Quick steps

1. Deploy the Global File Cache Virtual Template, or Windows Server 2016 VM, or Windows Server 2019 Standard or Datacenter edition.
2. Ensure the VM is connected to the network, joined to the domain, and accessible through RDP.
3. Install the latest Global File Cache Edge software.
4. Identify the Global File Cache Management Server and Core instance.
5. Configure the Global File Cache Edge instance.

## Global File Cache Edge requirements

Global File Cache Edge is designed to function across all platforms supporting Windows Server 2016 and 2019, bringing simplified IT to corporate remote offices and beyond. Critically, Global File Cache can be deployed on your existing hardware infrastructure, virtualization, or hybrid/public cloud environments in almost every case if they meet a few base-level requirements.

Global File Cache Edge requires the following hardware and software resources to function optimally. For more information about overall sizing guidelines, see [Sizing guidelines](#).

### Hardened server appliance

The Global File Cache installation package creates a hardened software appliance on any Microsoft Windows Server instance. *Do Not Uninstall* the Global File Cache Package. Uninstalling Global File Cache will impact the functionality of the server instance and might require a full rebuild of the server instance.

### Physical hardware requirements

- Minimum 4 CPU cores
- Minimum 16 GB RAM
- Dedicated single or redundant 1 Gbps NIC
- 10k RPM SAS HDD or SSD (preferred)
- RAID controller with write-back caching functionality enabled

### Virtual deployment requirements

Hypervisor platforms are known to be subject to performance degradation from a storage subsystem perspective (for example, latency). For optimal performance using Global File Cache, a physical server instance with SSD is recommended.

For best performance in virtual environments, in addition to the physical host requirements, the following requirements and resource reservations must be met:

Microsoft Hyper-V 2012 R2 and later:

- Processor (CPU): CPUs must be set as **Static**: Minimum: 4 vCPU cores.
- Memory (RAM): Minimum: 16 GB set as **Static**.
- Hard-disk provisioning: Hard Disks must be configured as **Fixed Disk**.

VMware vSphere 6.x and later:

- Processor (CPU): Reservation of CPU Cycles must be set. Minimum: 4 vCPU cores @ 10000 MHz.
- Memory (RAM): Minimum: Reservation of 16 GB.
- Hard-disk provisioning:
  - Disk Provisioning must be set as **Thick Provisioned Eager Zeroed**.
  - Hard Disk Shares must be set to **High**.
  - Devices.hotplug must be set to **False** using the vSphere Client to prevent Microsoft Windows from presenting Global File Cache drives as removable.
- Networking: Network Interface must be set to **VMXNET3** (requires VM Tools).

Global File Cache runs on Windows Server 2016 and 2019, hence the virtualization platform needs to support the operating system, as well as integration with utilities enhancing the performance of the VM's guest operating system and management of the VM, such as VM Tools.

### Partition sizing requirements

- C:\ - minimum 250GB (system/boot volume)
- D:\ - minimum 1TB (separate data volume for Global File Cache Intelligent File Cache\*)

\*Minimum size is 2x the active data set. The cache volume (D:\) can be extended and is only restricted by the limitations of the Microsoft Windows NTFS file system.

### Global File Cache Intelligent File Cache disk requirements

Disk Latency on the Global File Cache Intelligent File Cache disk (D:\) should deliver < 0.5ms average I/O disk latency and 1MiBps throughput per concurrent user.

For more information, see the [NetApp Global File Cache User Guide](#).

### Networking

- Firewall: TCP ports should be allowed between the Global File Cache Edge and Management Server and Core instances.

Global File Cache TCP Ports: 443 (HTTPS - LMS), 6618 – 6630.

- Network optimization devices (such as Riverbed Steelhead) must be configured to pass-thru Global File Cache specific ports (TCP 6618-6630).

### Client workstation and application best practices

Global File Cache transparently integrates into customer's environments, allowing users to access centralized data using their client workstations, running enterprise applications. Using Global File Cache, data is accessed through a direct drive mapping or through a DFS namespace. For more information about the Global File Cache Fabric, Intelligent File Caching, and key aspects of the software, consult the [Before you begin to Deploy Global File Cache](#) section.

To ensure an optimal experience and performance, it is important to comply with the Microsoft Windows Client requirements and best practices as outlined in the Global File Cache User Guide. This applies to all versions of Microsoft Windows.

For more information, see the [NetApp Global File Cache User Guide](#).

### Firewall and Antivirus best practices

While Global File Cache makes a reasonable effort to validate that the most common antivirus application suites are compatible with Global File Cache, NetApp cannot guarantee and is not responsible for any incompatibilities or performance issues caused by these programs, or their associated updates, service packs, or modifications.

Global File Cache does not recommend the installation nor application of monitoring or antivirus solutions on any Global File Cache enabled instance (Core or Edge). Should a solution be installed, by choice or by policy, the following best practices and recommendations must be applied. For common antivirus suites, see Appendix A in the [NetApp Global File Cache User Guide](#).

## Firewall settings

- Microsoft firewall:
  - Retain firewall settings as default.
  - Recommendation: Leave Microsoft firewall settings and services at the default setting of OFF, and not started for standard Global File Cache Edge instances.
  - Recommendation: Leave Microsoft firewall settings and services at the default setting of ON, and started for Edge instances that also run the Domain Controller role.
- Corporate firewall:
  - Global File Cache Core instance listens on TCP ports 6618-6630, ensure that Global File Cache Edge instances can connect to these TCP ports.
  - Global File Cache instances require communications to the Global File Cache Management Server on TCP port 443 (HTTPS).
- Network optimization solutions/devices must be configured to pass-thru Global File Cache specific ports.

## Antivirus best practices

This section helps you to understand the requirements when running antivirus software on a Windows Server instance running Global File Cache. Global File Cache has tested most commonly used antivirus products including Cylance, McAfee, Symantec, Sophos, Trend Micro, Kaspersky and Windows Defender for use in conjunction with Global File Cache.



Adding antivirus to an Edge appliance can introduce a 10–20% impact on user performance.

For more information, see the [NetApp Global File Cache User Guide](#).

## Configure exclusions

Antivirus software or other third-party indexing or scanning utilities should never scan drive D:\ on the Edge instance. These scans of Edge server drive D:\ will result in numerous file open requests for the entire cache namespace. This will result in file fetches over the WAN to all file servers being optimized at the data center. WAN connection flooding and unnecessary load on the Edge instance will occur resulting in performance degradation.

In addition to the D:\ drive, the following Global File Cache directory and processes should generally be excluded from all antivirus applications:

- C:\Program Files\TalonFAST\
  - C:\Program Files\TalonFAST\Bin\LMClientService.exe
  - C:\Program Files\TalonFAST\Bin\LMServerService.exe
  - C:\Program Files\TalonFAST\Bin\Optimus.exe
  - C:\Program Files\TalonFAST\Bin\tafsexport.exe
  - C:\Program Files\TalonFAST\Bin\tafsutils.exe
  - C:\Program Files\TalonFAST\Bin\tapp.exe
  - C:\Program Files\TalonFAST\Bin\tfs.exe
  - C:\Program Files\TalonFAST\Bin\TService.exe

- C:\Program Files\TalonFAST\Bin\tum.exe
- C:\Program Files\TalonFAST\FastDebugLogs\
- C:\Windows\System32\drivers\tfast.sys
- \\?\TafsMtPt:\ or \\?\TafsMtPt\*
- \Device\TalonCacheFS\
- \\?\GLOBALROOT\Device\TalonCacheFS\
- \\?\GLOBALROOT\Device\TalonCacheFS\\*

## NetApp Support policy

Global File Cache instances are designed specifically for Global File Cache as the primary application running on a Windows Server 2016 and 2019 platform. Global File Cache requires priority access to platform resources, for example, disk, memory, network interfaces, and can place high demands on these resources. Virtual deployments require memory/CPU reservations and high-performance disks.

- For branch office deployments of Global File Cache, supported services and applications on the server running Global File Cache are limited to:
  - DNS/DHCP
  - Active Directory domain controller (Global File Cache must be on a separate volume)
  - Print services
  - Microsoft System Center Configuration Manager (SCCM)
  - Global File Cache approved client-side system agents and anti-virus applications
- NetApp Support and maintenance applies only to Global File Cache.
- Line of business productivity software, which are typically resource intensive, for example, database servers, mail servers, and so on, are not supported.
- The customer is responsible for any non-Global File Cache software which might be installed on the server running Global File Cache:
  - If any third-party software package causes software or resource conflicts with Global File Cache or performance is compromised, Global File Cache's support organization might require the customer to disable or remove the software from the server running Global File Cache.
  - It is the customer's responsibility for all installation, integration, support, and upgrade of any software added to the server running the Global File Cache application.
- Systems management utilities/agents such as antivirus tools and licensing agents might be able to coexist. However, except for the supported services and applications listed above, these applications are not supported by Global File Cache and the same guidelines as above must still be followed:
  - It is the customer's responsibility for all installation, integration, support, and upgrade of any software added.
  - If a customer does install any third-party software package that causes, or is suspected to be causing, software or resource conflicts with Global File Cache or performance is compromised, there might be a requirement by Global File Cache's support organization to disable/remove the software.

# Deploy Global File Cache Edge instances

After you have verified that your environment meets all the requirements, you install Global File Cache Edge software in each remote office.

## Before you begin

To complete Global File Cache Edge configuration tasks, you need the following information:

- Static IP addresses for each Global File Cache instance
- Subnet mask
- Gateway IP address
- The FQDN you wish to assign to each Global File Cache server
- The DNS suffix (optional)
- The user name and password of an administrative user in the domain
- The FQDN and/or IP address of the associated Core servers
- A volume to be used as the Intelligent File Cache. It is recommended this be at least 2x the size of the active dataset. This should be formatted as NTFS and assigned as D:\.

## Commonly used TCP ports

There are several TCP ports used by Global File Cache services. It is mandatory that the devices can communicate on these ports and they be excluded from any WAN optimization devices or firewall restriction policies:

- Global File Cache Licensing TCP Port: 443
- Global File Cache TCP Ports: 6618-6630

## Deploy the Global File Cache Virtual Template

The virtual template (.OVA and .VHD) images contain the latest release of the Global File Cache software. If you are deploying Global File Cache using the .OVA or .VHD virtual machine (VM) template, follow the steps as outlined in this section. It is assumed that you understand how to deploy the .OVA or .VHD template on the designated hypervisor platform.

Ensure that VM preferences, including resource reservations, are in line with the requirements as outlined in [Virtual deployment requirements](#).

### Steps

1. Extract the package from the template you downloaded.
2. Deploy the virtual template. Refer to the following videos before you start the deployment:
  - [Deploy the Virtual Template on VMware](#)
  - [Deploy the Virtual Template on Hyper-V](#)
3. After the Virtual Template has been deployed, and you have configured the VM settings, start the VM.
4. During initial boot, when the Windows Server 2016 or 2019 operating system is preparing for first use, complete the out-of-the-box experience by installing the correct drivers and installing the necessary components for the respective hardware.

5. When the base install of the Global File Cache Edge instance has been completed, the Windows Server 2016 or 2019 operating system will guide you through an initial configuration wizard to configure operating system specifics such as localization and product key.
6. After the initial configuration wizard has completed, log in locally to the Windows Server 2016 or 2019 operating system with the following credentials:
  - User name: **FASTAdmin**
  - Password: **Tal0nFAST!**
7. Configure your Windows Server VM, join to the organization's Active Directory domain, and proceed to the Global File Cache Edge configuration section.

## Configure the Global File Cache Edge instance

The Global File Cache Edge instance connects to a Global File Cache Core to provide users at the branch office access to data center file server resources.



The Edge instance must be licensed as part of your Cloud Volumes ONTAP deployment prior to beginning the configuration. See [Licensing](#) for more information about licensing.

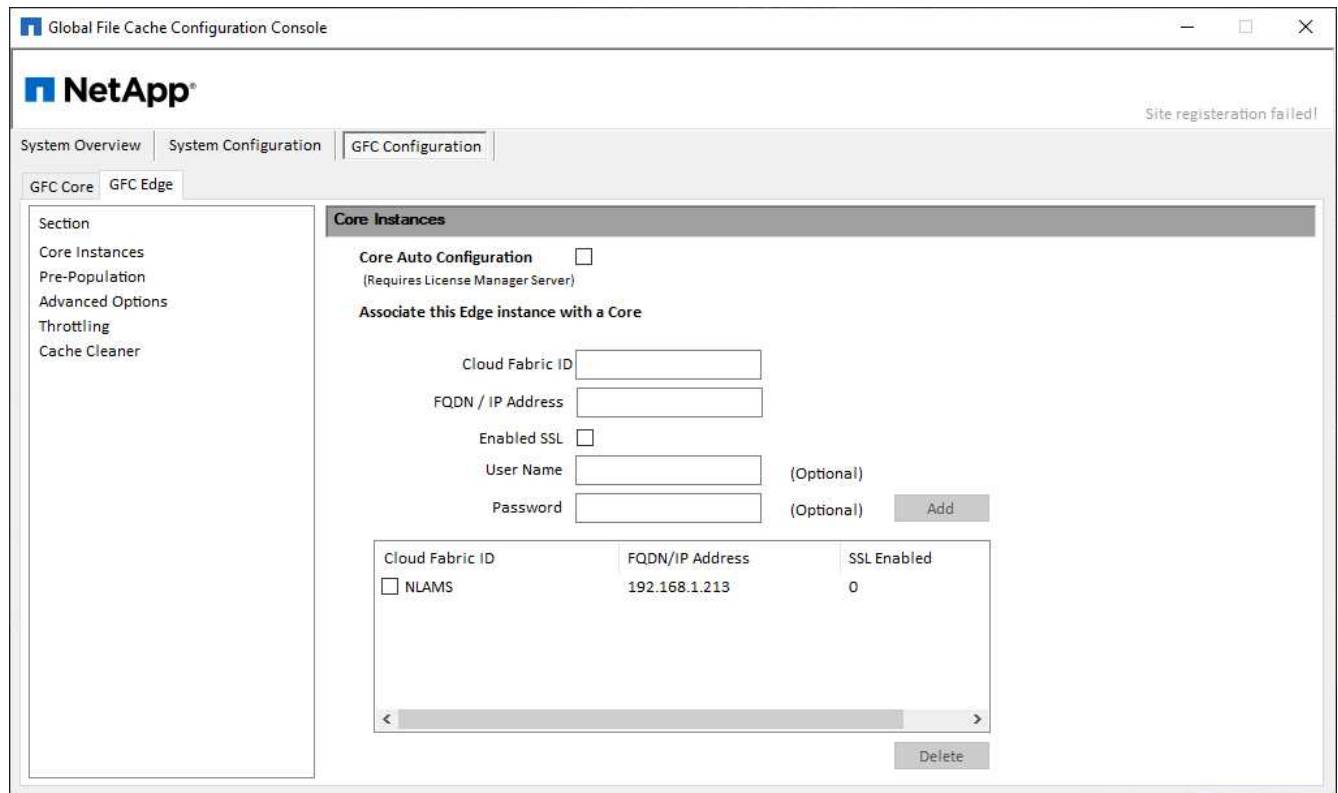
If your configuration requires more than one Global File Cache Core to be installed because of a large number of Edge instances, you will configure some Edge instances to connect to the first Core and others to connect to the second Core. Make sure you have the FQDN or IP address, and other required information, for the correct Core instance.

To configure the Edge instance, complete the following steps:

### Steps

1. Click **Perform** next to the unchecked Core Configuration step listed in the "Edge Configuration Steps" section of the Initial Configuration assistant. This opens a new tab, GFC Edge, and shows the section *Core Instances*.
2. Provide the **Cloud Fabric ID** of the Global File Cache Core server. The Cloud Fabric ID is typically the NetBIOS name or the geographical location of the backend file server.
3. Provide the **FQDN/IP Address** of the Global File Cache Core server:
  - a. (Optional) Check the **SSL** box to enable SSL support for enhanced encryption from the Edge to the Core.
  - b. Enter the User Name and Password, which are the credentials of the Service Account used on the Core.
4. Click **Add** to confirm the addition of the Global File Cache Core appliance. A confirmation box will appear. Click **OK** to dismiss it.





## Update Global File Cache Edge software

Global File Cache frequently releases updates to the software, either patches, enhancements, or new features/functionality. Although the virtual template (.OVA and .VHD) images contain the latest release of the Global File Cache software, it is possible that a newer version is available on the NetApp Support Download portal.

Ensure that your Global File Cache instances are up to date with the latest version.



This software package can also be used for pristine installations on Microsoft Windows Server 2016 Standard or Datacenter edition, or Windows Server 2019 Standard or Datacenter edition, or used as part of your upgrade strategy.

Below you can find the steps required to update the Global File Cache installation package:

### Steps

1. After saving the latest installation package to the desired Windows Server instance, double-click it to run the installation executable.
2. Click **Next** to continue the process.
3. Click **Next** to continue.
4. Accept the Licensing Agreement and click **Next**.
5. Select the desired Installation Destination Location.

NetApp recommends that the default installation location be used.

6. Click **Next** to continue.
7. Select the Start Menu Folder.

8. Click **Next** to continue.
9. Verify the desired installation parameters and click **Install** to begin the installation.

The installation process will execute.

10. After the installation has completed, reboot the server when prompted.

### What's Next?

For details about Global File Cache Edge advanced configuration, see the [NetApp Global File Cache User Guide](#).

## End-user training

You will want to train your users on the best practices for accessing the shared files through Global File Cache.

This is the final phase of the Global File Cache deployment, the end-user implementation phase.

In order to prepare and streamline the end user on-boarding process, use the email template below that will help you to educate end users on what it means to work in a "central data" environment. This will help your users leverage all of the benefits of the Global File Cache solution. We have also published a video that can be shared to "train" users where needed.

Customize and forward the following resources to end users to prepare them for roll-out:

- User Training video  
[End user training video](#)
- Email Template  
[Mac Email Template \(.emltpl\)](#)  
[Windows Email Template \(.msg\)](#)
- Onboarding Communications  
[Word Document \(.docx\)](#)

See Chapter 12 in the [NetApp Global File Cache User Guide](#) for additional material.

## Additional information

Use the following links to learn more about Global File Cache and other NetApp products:

- Global File Cache FAQ
  - See a list of frequently asked questions and answers [here](#)
- [NetApp Global File Cache User Guide](#)
- NetApp Product Documentation
  - See additional documentation for NetApp cloud products [here](#)
  - See additional documentation for all NetApp products [here](#)
- Customer support for Global File Cache users with Cloud Volumes ONTAP is available through these channels:

- Guided Problem Solving, Case Management, Knowledgebase, Downloads, Tools, and more go [here](#)
- Login to the NetApp Support at <https://mysupport.netapp.com> with your NSS credentials
- For immediate assistance for a P1 issue call: +1 856.481.3990 (Option 2)
- Customer support for Global File Cache users utilizing Cloud Volumes Services and Azure NetApp Files is available through standard support from your provider. Please contact Google Customer Support or Microsoft Customer Support respectively.

# Optimize cloud compute costs

## Learn about the Compute service

By leveraging [Spot's Cloud Analyzer service](#), Cloud Manager can provide a high-level cost analysis of your cloud compute spending and identify potential savings.

Cloud Analyzer is a cloud infrastructure management solution that uses advanced analytics to provide visibility and insights into your cloud costs. It shows you where you can optimize those costs and lets you implement that optimization using Spot's portfolio of continuous optimization products in just a few clicks.

### Features

- A cost analysis that shows current cost for the month, projected monthly costs, and missed savings
- A view of spend efficiency by account, including the estimated additional savings
- A link to Spot's Cloud Analyzer for more in-depth details about the spending for all accounts

### Supported cloud providers

This service is supported with AWS.

### Cost

There's no cost to use this service through Cloud Manager.

### How Cloud Analyzer works with Cloud Manager

At a high-level, Cloud Analyzer integration with Cloud Manager works like this:

1. You click **Compute** and connect your AWS master payer account.
2. NetApp configures your environment as follows:
  - a. Creates an organization in the Spot platform.
  - b. Sends an email welcoming you to Spot.

You can log in to the Spot service using the same single-sign on credentials that you use with Cloud Central and Cloud Manager.

- c. Cloud Analyzer starts processing your AWS account data.
3. In Cloud Manager, the Compute page refreshes and you use the information to gain insights on past, current, and future cloud costs.
  4. You click **Get Full Analysis** at any time to go to Spot's Cloud Analyzer, which provides a full analysis of your cloud spend and savings opportunities.

### Data security

Cloud Analyzer data is encrypted at rest and no credentials are stored for any account.

# Start optimizing your cloud compute costs

Connect your AWS account and then view the analysis to start optimizing your cloud compute costs.

## Connect Cloud Analyzer to your AWS account

Click **Compute** and connect your AWS payer account.

### Steps

1. Click **Compute**.
2. Click **Add AWS Credentials to Start**.
3. Follow the steps on the page to connect your AWS account:
  - a. Log in to your AWS master payer account.
  - b. Set up cost and usage reports on the AWS account.
  - c. Run the CloudFormation template.
  - d. Paste the Spot RoleARN.

[View more details about these steps.](#)

### Connect your AWS Account to Optimize Costs

Connecting your billing data will allow Cloud Analyzer to access your Cost and Usage data.

**Step 1**  
Log in to your AWS Master Payer account.

**Step 2**  
Set up your Cost and Usage Reports on your AWS account.  
([Learn How](#) or skip this if the report is already enabled.)  
Enter the bucket name where the report is located:

**Step 3**  
Open CloudFormation with Spot template.  
Under capabilities, mark "I acknowledge that AWS CloudFormation might create IAM resources" and click 'Create'.


**Step 4**  
Copy the Spot RoleARN from the Output tab and paste below.

## Result

Cloud Analyzer starts processing your AWS account data. If you have multiple accounts, Cloud Analyzer starts with read-only capabilities for all linked accounts under the master payer account. If you want to get more details about the potential savings for those accounts, then you'll need to connect them, as well. You can find more details about that process in the section below.

## Analyze your compute costs

After Cloud Analyzer processes your account data, the Compute tab shows you insights on past, current, and future cloud costs.



### Month to date cost

The total cost of your workloads from the beginning of the current month to present.

### Projected Costs

The forecasted cost at the end of the month based on analysis of your usage pattern.

### Last Week Missed Savings

Savings that could have been achieved in the previous seven days using optimization of spot instances and reservations.

### Top AWS Account Spend Efficiency

The top 10 accounts according to the greatest amount of estimated additional savings.

Each account is assigned an efficiency score based on current and additional potential savings. The estimated additional savings indicates how much can be further saved by leveraging the use of spot and reserved instances.

You can take the following actions to further optimize your accounts:

- **View details:** View your cost optimization opportunities by going to Spot's Cloud Analyzer.
- **Connect:** Connect an account that is not yet managed. You will be directed to the wizard that connects the account.

## Top AWS Accounts


This is a bar graph showing your top ten accounts by cost. The graph is based on the last 30 days of spend activity.

[Learn more about the Cost Analysis page that's available in Spot's Cloud Analyzer.](#)

## Go to Cloud Analyzer for more analysis and recommendations

Click **Get Full Analysis** at any time to access more charts and analysis, in-depth recommendations, a use case optimization breakdown (containers, ElasticApps, and reservations), and more.

Here's an example of what you'll see in Cloud Analyzer:



- [View the product page for Cloud Analyzer to learn more about its capabilities.](#)
- [View the documentation for Spot to get help using Cloud Analyzer.](#)

# Tier data to the cloud

## Learn about Cloud Tiering

NetApp's Cloud Tiering service extends your data center to the cloud by automatically tiering inactive data from on-premises ONTAP clusters to object storage. This frees valuable space on the cluster for more workloads, without making changes to the application layer. Cloud Tiering can reduce costs in your data center and enables a switch from a CAPEX model to an OPEX model.

The Cloud Tiering service leverages the capabilities of *FabricPool*. FabricPool is a NetApp Data Fabric technology that enables automated tiering of data to low-cost object storage. Active data remains on high-performance SSDs, while inactive data is tiered to low-cost object storage while preserving ONTAP data efficiencies.

### Features

Cloud Tiering offers automation, monitoring, reports, and a common management interface:

- Automation makes it easier to set up and manage data tiering from on-prem ONTAP clusters to the cloud
- A single pane of glass removes the need to independently manage FabricPool across several clusters
- Reports show the amount of active and inactive data on each cluster
- A tiering health status helps you identify and correct issues as they occur
- If you have Cloud Volumes ONTAP systems, you'll find them in the Cluster Dashboard so you get a full view of data tiering in your hybrid cloud infrastructure



Cloud Volumes ONTAP systems are read-only from Cloud Tiering. [You set up tiering for Cloud Volumes ONTAP from the working environment in Cloud Manager.](#)

For more details about the value that Cloud Tiering provides, [check out the Cloud Tiering page on NetApp Cloud Central.](#)



While Cloud Tiering can significantly reduce storage footprints, it is not a backup solution.

### Supported object storage providers

You can tier inactive data from an ONTAP cluster to Amazon S3, Microsoft Azure Blob storage, Google Cloud Storage, or StorageGRID (private cloud).

### Pricing and licenses

Pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP tiering license called *FabricPool*, or a combination of both. A 30-day free trial is available for your first cluster if you don't have a license.

There are no charges when tiering data to StorageGRID. Neither a BYOL license or PAYGO registration is required.

[View pricing details.](#)



## 30-day free trial

If you don't have a FabricPool license, a 30-day free trial of Cloud Tiering starts when you set up tiering to your first cluster. After that 30-day free trial ends, you'll need to pay for Cloud Tiering through a pay-as-you-go subscription, a FabricPool license, or a combination of both.

If your free trial ends and you haven't subscribed or added a license, then ONTAP no longer tiers cold data to object storage, but existing data is still available for access.

## Pay-as-you-go subscription

Cloud Tiering offers consumption-based licensing in a pay-as-you-go model. After subscribing through your cloud provider's marketplace, you pay per GB for data that's tiered—there's no up-front payment. You are billed by your cloud provider through your monthly bill.

You should subscribe even if you have a free trial or if you bring your own license (BYOL):

- Subscribing ensures that there's no disruption of service after your free trial ends.

When the trial ends, you'll be charged hourly according to the amount of data that you tier.

- If you tier more data than allowed by your FabricPool license, then data tiering continues through your pay-as-you-go subscription.

For example, if you have a 10 TB license, all capacity beyond the 10 TB is charged through the pay-as-you-go subscription.

You won't be charged from your pay-as-you-go subscription during your free trial or if you haven't exceeded your FabricPool license.

[Learn how to set up a pay-as-you-go subscription.](#)

## Bring your own license

Bring your own license by purchasing an ONTAP FabricPool license from NetApp. You can purchase term-based or perpetual licenses.

After you purchase a FabricPool license, you'll need to add it to the cluster, [which you can do directly from Cloud Tiering](#).

After you activate the license through Cloud Tiering, if you purchase additional add-on capacity at a later time, the license on the cluster is automatically updated with the new capacity. There's no need to apply a new NetApp License File (NLF) to the cluster.


As noted above, we recommend that you set up a pay-as-you-go subscription, even if your cluster has a BYOL license.

[Contact us to purchase a license.](#)

## How Cloud Tiering works

Cloud Tiering is a NetApp-managed service that uses FabricPool technology to automatically tier inactive (cold) data from your on-premises ONTAP clusters to object storage in your public cloud or private cloud. Connections to ONTAP take place from a Connector.

The following image shows the relationship between each component:



At a high level, Cloud Tiering works like this:

1. You discover your on-prem cluster from Cloud Manager.
2. You set up tiering by providing details about your object storage, including the bucket/container and a storage class or access tier.
3. Cloud Manager configures ONTAP to use the object storage provider and discovers the amount of active and inactive data on the cluster.
4. You choose the volumes to tier and the tiering policy to apply to those volumes.
5. ONTAP starts tiering inactive data to the object store, as soon as the data has reached the thresholds to be considered inactive (see [Volume tiering policies](#)).

## Object storage

Each ONTAP cluster tiers inactive data to a single object store. When you set up data tiering, you have the choice to add a new bucket/container or to select an existing bucket/container, along with a storage class or access tier.

- [Learn about supported S3 storage classes](#)
- [Learn about supported Azure Blob access tiers](#)
- [Learn about supported Google Cloud storage classes](#)

## Volume tiering policies

When you select the volumes that you want to tier, you choose a *volume tiering policy* to apply to each volume. A tiering policy determines when or whether the user data blocks of a volume are moved to the cloud.

## No tiering policy

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

## Cold snapshots (Snapshot only)

ONTAP tiers cold Snapshot blocks in the volume that are not shared with the active file system to object storage. If read, cold data blocks on the cloud tier become hot and are moved to the performance tier.

Data is tiered only after an aggregate has reached 50% capacity and when the data has reached the cooling period. The default number of cooling days is 2, but you can adjust the number of days.



Writes from the cloud tier to the performance tier are disabled if performance tier capacity is greater than 70%. If this occurs, blocks are accessed directly from the cloud tier.

## Cold user data (Auto)

ONTAP tiers all cold blocks in the volume (not including metadata) to object storage. The cold data includes not just Snapshot copies but also cold user data from the active file system.

If read by random reads, cold data blocks on the cloud tier become hot and are moved to the performance tier. If read by sequential reads, such as those associated with index and antivirus scans, cold data blocks on the cloud tier stay cold and are not written to the performance tier.

Data is tiered only after an aggregate has reached 50% capacity and when the data has reached the cooling period. The cooling period is the time that user data in a volume must remain inactive for the data to be considered "cold" and moved to the object store. The default number of cooling days is 31, but you can adjust the number of days.



Writes from the cloud tier to the performance tier are disabled if performance tier capacity is greater than 70%. If this occurs, blocks are accessed directly from the cloud tier.

## All user data (All)

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

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

Take the following into consideration before you choose this tiering policy:

- Tiering data immediately reduces storage efficiencies (inline only).
- You should use this policy only if you are confident that cold data on the volume will not change.
- Object storage is not transactional and will result in significant fragmentation if subjected to change.
- Consider the impact of SnapMirror transfers before assigning the All tiering policy to source volumes in data protection relationships.

Because data is tiered immediately, SnapMirror will read data from the cloud tier rather than the performance tier. This will result in slower SnapMirror operations—possibly slowing other SnapMirror operations later in queue—even if they are using different tiering policies.

## All DP user data (Backup)

All data on a data protection volume (not including metadata) is immediately moved to the cloud tier. If read, cold data blocks on the cloud tier stay cold and are not written back to the performance tier (starting with ONTAP 9.4).



This policy is available for ONTAP 9.5 or earlier. It was replaced with the **All** tiering policy starting with ONTAP 9.6.

# Get started

## Tiering data from on-premises ONTAP clusters to Amazon S3

Free space on your on-prem ONTAP clusters by tiering data to Amazon S3. Data tiering is powered by NetApp's Cloud Tiering service.

### Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



### Prepare to tier data to Amazon S3

You need the following:

- An AFF or FAS system with all-SSD aggregates that's running ONTAP 9.2 or later and has an HTTPS connection to Amazon S3.
- An AWS account that has an access key and [the required permissions](#) so the ONTAP cluster can tier inactive data in and out of S3.
- A Connector installed in an AWS VPC or on your premises.
- Networking for the Connector that enables an outbound HTTPS connection to the ONTAP cluster, to S3 storage, and to the Cloud Tiering service.



### Set up tiering

In Cloud Manager, select an on-prem working environment, click **Setup Tiering** and follow the prompts to tier data to Amazon S3.



### Set up licensing

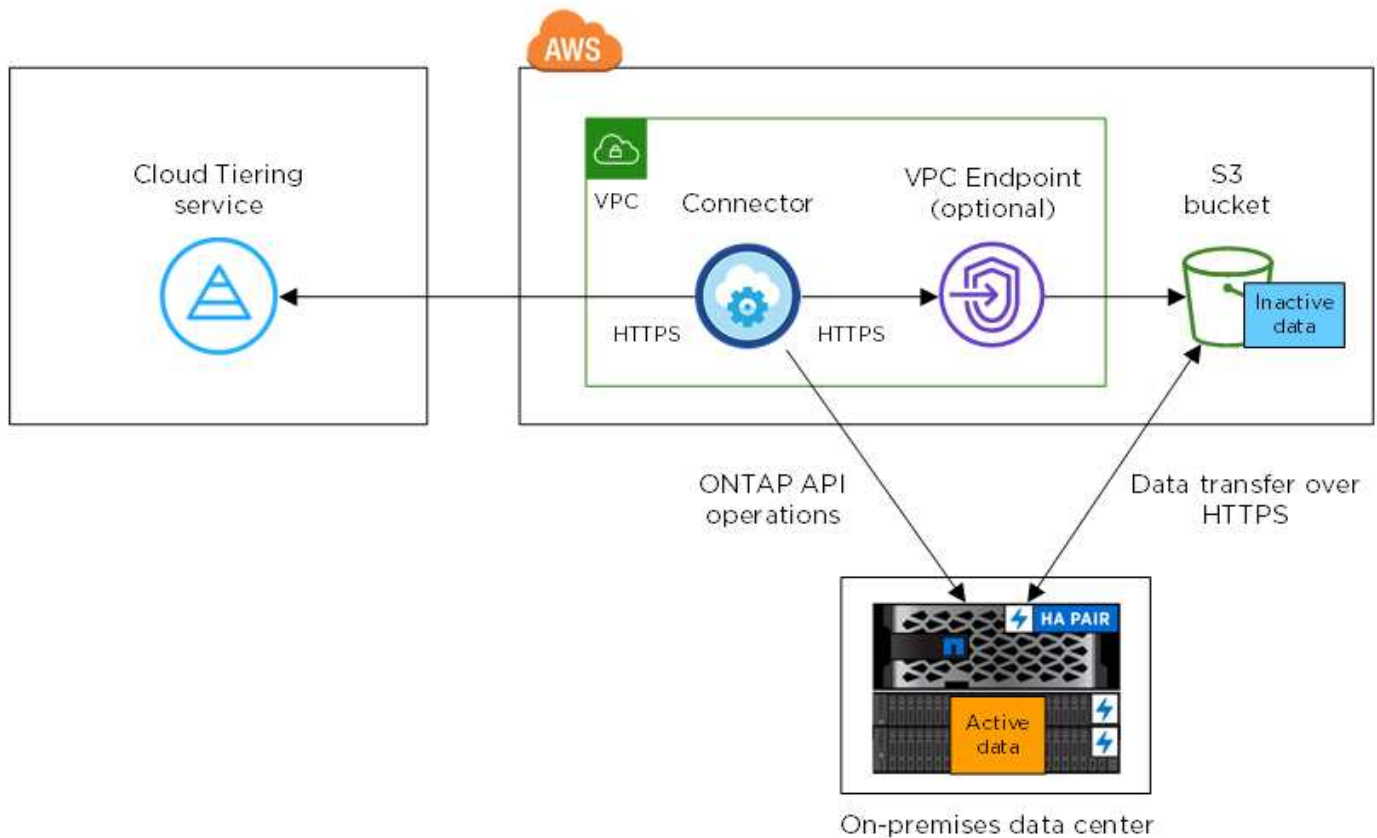
After your free trial ends, pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP tiering license, or a combination of both:

- To subscribe from the AWS Marketplace, click **Tiering > Licensing**, click **Subscribe**, and then follow the prompts.
- To pay using a tiering license, [contact us if you need to purchase one](#), and then [add it to your cluster from Cloud Tiering](#).

## Requirements

Verify support for your ONTAP cluster, set up your networking, and prepare your object storage.

The following image shows each component and the connections that you need to prepare between them:



Communication between a Connector and S3 is for object storage setup only. The Connector can reside on your premises, instead of in the cloud.

## Preparing your ONTAP clusters

Your ONTAP clusters must meet the following requirements when tiering data to Amazon S3.

### Supported ONTAP platforms

Cloud Tiering supports AFF systems and all-SSD aggregates on FAS systems.

### Supported ONTAP version

ONTAP 9.2 or later

### Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 to Amazon S3.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

Although AWS Direct Connect provides better performance and lower data transfer charges, it's not required between the ONTAP cluster and S3. Because performance is significantly better when using AWS Direct Connect, doing so is the recommended best practice.

- An inbound connection is required from the Connector, which can reside in an AWS VPC or on your premises.

A connection between the cluster and the Cloud Tiering service is not required.

- An intercluster LIF is required on each ONTAP node that hosts tiered volumes. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

IPspaces enable network traffic segregation, allowing for separation of client traffic for privacy and security. [Learn more about IPspaces.](#)

When you set up data tiering, Cloud Tiering prompts you for the IPspace to use. You should choose the IPspace that each LIF is associated with. That might be the "Default" IPspace or a custom IPspace that you created.

## Supported volumes and aggregates

The total number of volumes that Cloud Tiering can tier might be less than the number of volumes on your ONTAP system. That's because volumes can't be tiered from some aggregates. For example, you can't tier data from SnapLock volumes or from MetroCluster configurations. Refer to ONTAP documentation for [functionality or features not supported by FabricPool.](#)



Cloud Tiering supports FlexGroup volumes, starting with ONTAP 9.5. Setup works the same as any other volume.

## Creating or switching Connectors

A Connector is required to tier data to the cloud. When tiering data to AWS S3, you can use a Connector that's in an AWS VPC or on your premises. You'll either need to create a new Connector or make sure that the currently selected Connector resides in AWS or on-prem.

- [Learn about Connectors](#)
- [Creating a Connector in AWS](#)
- [Connector host requirements](#)
- [Installing the Connector on an existing Linux host](#)
- [Switching between Connectors](#)

## Preparing networking for the Connector

Ensure that the Connector has the required networking connections. A Connector can be installed on-prem or in AWS.

### Steps

1. Ensure that the network where the Connector is installed enables the following connections:
  - An outbound internet connection to the Cloud Tiering service over port 443 (HTTPS)
  - An HTTPS connection over port 443 to S3
  - An HTTPS connection over port 443 to your ONTAP clusters
2. If needed, enable a VPC Endpoint to S3.

A VPC Endpoint to S3 is recommended if you have a Direct Connect or VPN connection from your ONTAP cluster to the VPC and you want communication between the Connector and S3 to stay in your AWS

internal network.

### Preparing Amazon S3

When you set up data tiering to a new cluster, you're prompted to create an S3 bucket or to select an existing S3 bucket in the AWS account where the Connector is set up.

The AWS account must have permissions and an access key that you can enter in Cloud Tiering. The ONTAP cluster uses the access key to tier data in and out of S3.

### Steps

1. Provide the following permissions to the IAM user:

```
"s3:ListAllMyBuckets",  
"s3:ListBucket",  
"s3:GetBucketLocation",  
"s3:GetObject",  
"s3:PutObject",  
"s3:DeleteObject"
```

[AWS Documentation: Creating a Role to Delegate Permissions to an IAM User](#)

2. Create or locate an access key.

Cloud Tiering passes the access key on to the ONTAP cluster. The credentials are not stored in the Cloud Tiering service.

[AWS Documentation: Managing Access Keys for IAM Users](#)

### Tiering inactive data from your first cluster to Amazon S3


After you prepare your AWS environment, start tiering inactive data from your first cluster.

### What you'll need

- [An on-premises working environment.](#)
- An AWS access key for an IAM user who has the required S3 permissions.

### Steps

1. Select an on-prem cluster.
2. Click **Setup Tiering**.



You're now on the Tiering dashboard.

3. Click **Set up Tiering** next to the cluster.
4. Complete the steps on the **Tiering Setup** page:
  - a. **S3 Bucket:** Add a new S3 bucket or select an existing S3 bucket that starts with the prefix *fabric-pool* and click **Continue**.

The *fabric-pool* prefix is required because the IAM policy for the Connector enables the instance to perform S3 actions on buckets named with that exact prefix.

For example, you could name the S3 bucket *fabric-pool-AFF1*, where AFF1 is the name of the cluster.

- b. **Storage Class:** Select the S3 storage class that you want to transition the data to after 30 days and click **Continue**.

If you choose Standard, then the data remains in that storage class.

- c. **Credentials:** Enter the access key ID and secret key for an IAM user who has the required S3 permissions.


The IAM user must be in the same AWS account as the bucket that you selected or created on the **S3 Bucket** page.

- d. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and click **Continue**.

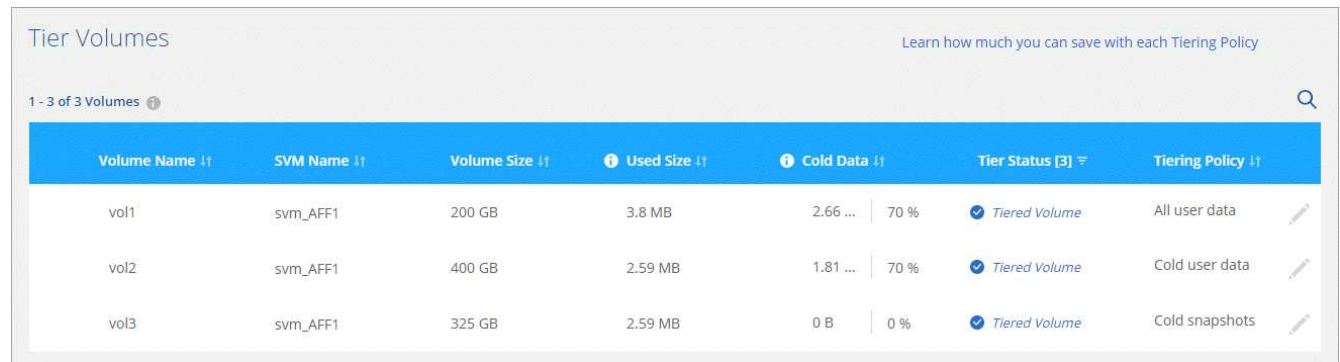
Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to your cloud provider's object storage.

5. Click **Continue** to select the volumes that you want to tier.



6. On the **Tier Volumes** page, set up tiering for each volume. Click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

[Learn more about volume tiering policies.](#)



Volume Name	SVM Name	Volume Size	Used Size	Cold Data	Tier Status	Tiering Policy
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ...	70 %	Tiered Volume All user data
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ...	70 %	Tiered Volume Cold user data
vol3	svm_AFF1	325 GB	2.59 MB	0 B	0 %	Tiered Volume Cold snapshots

## Result

You've successfully set up data tiering from volumes on the cluster to S3 object storage.

## What's next?

[Be sure to subscribe from the Cloud Tiering service.](#)

You can also add additional clusters or review information about the active and inactive data on the cluster. For details, see [Managing data tiering from your clusters.](#)

## Tiering data from on-premises ONTAP clusters to Azure Blob storage

Free space on your on-prem ONTAP clusters by tiering data to Azure Blob storage. Data tiering is powered by NetApp's Cloud Tiering service.

### Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



### Prepare to tier data to Azure Blob storage

You need the following:

- An AFF or FAS system with all-SSD aggregates that's running ONTAP 9.4 or later and has an HTTPS connection to Azure Blob storage.
- A Connector installed in an Azure VNet.
- Networking for a Connector that enables an outbound HTTPS connection to the ONTAP cluster in your data center, to Azure Blob storage, and to the Cloud Tiering service.



### Set up tiering

In Cloud Manager, select an on-prem working environment, click **Setup Tiering** and follow the prompts to tier

data to Azure Blob storage.

### 3 Set up licensing

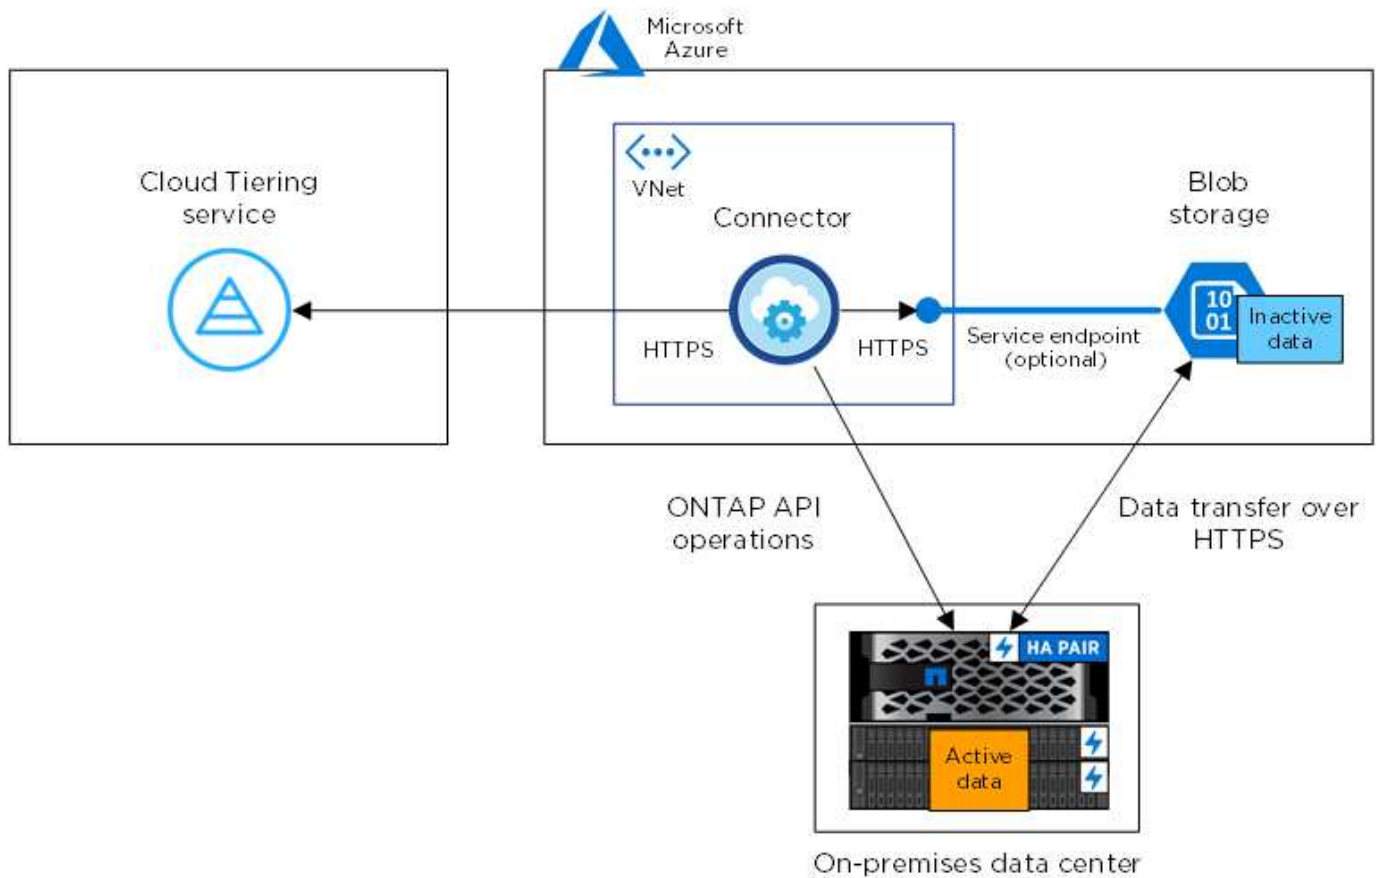
After your free trial ends, pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP tiering license, or a combination of both:

- To subscribe from the Azure Marketplace, click **Tiering > Licensing**, click **Subscribe**, and then follow the prompts.
- To add a tiering license, [contact us if you need to purchase one](#), and then [add it to your cluster from Cloud Tiering](#).

#### Requirements

Verify support for your ONTAP cluster, set up your networking, and prepare your object storage.

The following image shows each component and the connections that you need to prepare between them:



Communication between the Connector and Blob storage is for object storage setup only.

#### Preparing your ONTAP clusters

Your ONTAP clusters must meet the following requirements when tiering data to Azure Blob storage.

## Supported ONTAP platforms

Cloud Tiering supports AFF systems and all-SSD aggregates on FAS systems.

## Supported ONTAP version

ONTAP 9.4 or later

## Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 to Azure Blob storage.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

Although ExpressRoute provides better performance and lower data transfer charges, it's not required between the ONTAP cluster and Azure Blob storage. Because performance is significantly better when using ExpressRoute, doing so is the recommended best practice.

- An inbound connection is required from the NetApp Service Connector, which resides in an Azure VNet.

A connection between the cluster and the Cloud Tiering service is not required.

- An intercluster LIF is required on each ONTAP node that hosts tiered volumes. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

IPspaces enable network traffic segregation, allowing for separation of client traffic for privacy and security. [Learn more about IPspaces.](#)

When you set up data tiering, Cloud Tiering prompts you for the IPspace to use. You should choose the IPspace that each LIF is associated with. That might be the "Default" IPspace or a custom IPspace that you created.

## Supported volumes and aggregates

The total number of volumes that Cloud Tiering can tier might be less than the number of volumes on your ONTAP system. That's because volumes can't be tiered from some aggregates. For example, you can't tier data from SnapLock volumes or from MetroCluster configurations. Refer to ONTAP documentation for [functionality or features not supported by FabricPool.](#)



Cloud Tiering supports FlexGroup volumes, starting with ONTAP 9.5. Setup works the same as any other volume.

## Creating or switching Connectors

A Connector is required to tier data to the cloud. When tiering data to Azure Blob storage, a Connector must be available in an Azure VNet. You'll either need to create a new Connector or make sure that the currently selected Connector resides in Azure.

- [Learn about Connectors](#)
- [Creating a Connector in Azure](#)
- [Switching between Connectors](#)

## Preparing networking for the Connector

Ensure that the Connector has the required networking connections.

## Steps

1. Ensure that the VNet where the Connector is installed enables the following connections:
  - An outbound internet connection to the Cloud Tiering service over port 443 (HTTPS)
  - An HTTPS connection over port 443 to Azure Blob storage
  - An HTTPS connection over port 443 to your ONTAP clusters
2. If needed, enable a VNet service endpoint to Azure storage.

A VNet service endpoint to Azure storage is recommended if you have an ExpressRoute or VPN connection from your ONTAP cluster to the VNet and you want communication between the Connector and Blob storage to stay in your virtual private network.

## Tiering inactive data from your first cluster to Azure Blob storage


After you prepare your Azure environment, start tiering inactive data from your first cluster.

### What you'll need

[An on-premises working environment.](#)

## Steps

1. Select an on-prem cluster.
2. Click **Setup Tiering**.



You're now on the Tiering dashboard.

3. Click **Set up Tiering** next to the cluster.
4. Complete the steps on the **Tiering Setup** page:


- a. **Resource Group:** Select a resource group where an existing container is managed, or where you would like to create a new container for tiered data.
- b. **Azure Container:** Add a new Blob container to a storage account or select an existing container and click **Continue**.

The storage account and containers that appear in this step belong to the resource group that you selected in the previous step.

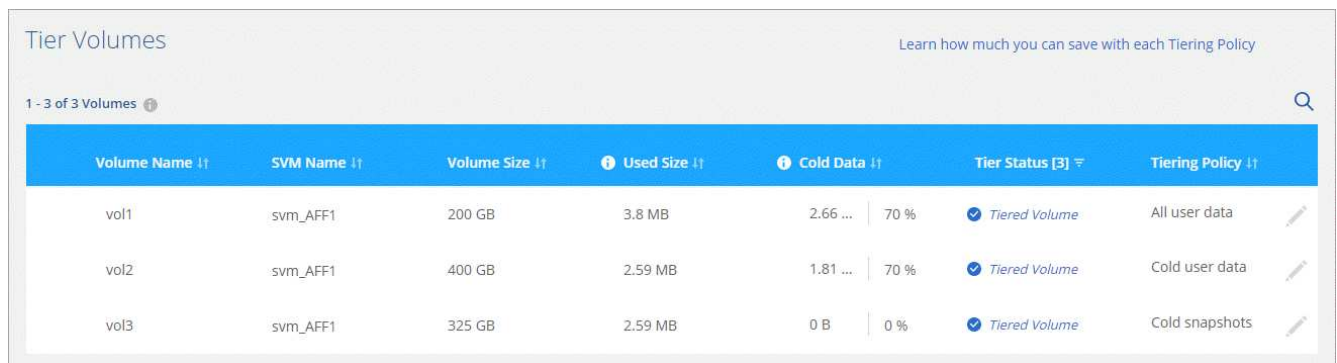
- c. **Access Tier:** Select the access tier that you want to use for the tiered data and click **Continue**.
- d. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and click **Continue**.

Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to your cloud provider's object storage.

5. Click **Continue** to select the volumes that you want to tier.

6. On the **Tier Volumes** page, set up tiering for each volume. Click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

[Learn more about volume tiering policies.](#)



Volume Name	SVM Name	Volume Size	Used Size	Cold Data	Tier Status [3]	Tiering Policy
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ... 70 %	Tiered Volume	All user data
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ... 70 %	Tiered Volume	Cold user data
vol3	svm_AFF1	325 GB	2.59 MB	0 B 0 %	Tiered Volume	Cold snapshots

## Result

You've successfully set up data tiering from volumes on the cluster to Azure Blob object storage.

## What's next?

[Be sure to subscribe from the Cloud Tiering service.](#)

You can also add additional clusters or review information about the active and inactive data on the cluster. For details, see [Managing data tiering from your clusters](#).

## Tiering data from on-premises ONTAP clusters to Google Cloud Storage

Free space on your on-prem ONTAP clusters by tiering data to Google Cloud Storage. Data tiering is powered by NetApp's Cloud Tiering service.

## Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.

# 1

## Prepare to tier data to Google Cloud Storage

You need the following:

- An AFF or FAS system with all-SSD aggregates that's running ONTAP 9.6 or later and has an HTTPS connection to Google Cloud Storage.
- A service account that has the predefined Storage Admin role and storage access keys.
- A Connector installed in a Google Cloud Platform VPC.
- Networking for the Connector that enables an outbound HTTPS connection to the ONTAP cluster in your data center, to Google Cloud Storage, and to the Cloud Tiering service.

# 2

## Set up tiering

In Cloud Manager, select an on-prem working environment, click **Setup Tiering** and follow the prompts to tier data to Google Cloud Storage.

# 3

## Set up licensing

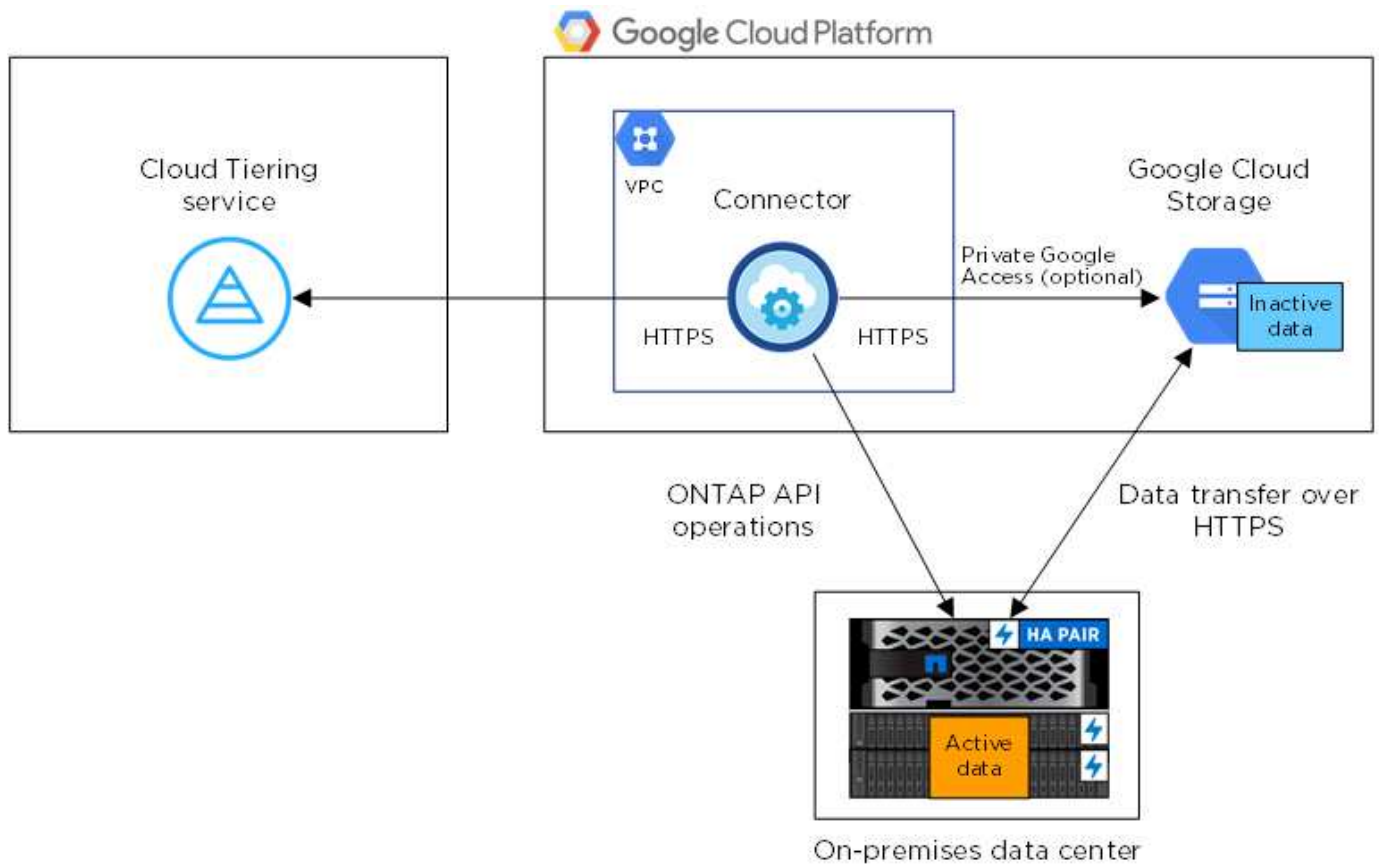
After your free trial ends, pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP tiering license, or a combination of both:

- To subscribe from the GCP Marketplace, click **Tiering > Licensing**, click **Subscribe**, and then follow the prompts.
- To add a tiering license, [contact us if you need to purchase one](#), and then [add it to your cluster from Cloud Tiering](#).

## Requirements

Verify support for your ONTAP cluster, set up your networking, and prepare your object storage.

The following image shows each component and the connections that you need to prepare between them:



Communication between the Connector and Google Cloud Storage is for object storage setup only.

### Preparing your ONTAP clusters

Your ONTAP clusters must meet the following requirements when tiering data to Google Cloud Storage.

### Supported ONTAP platforms

Cloud Tiering supports AFF systems and all-SSD aggregates on FAS systems.

### Supported ONTAP versions

ONTAP 9.6 or later

### Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 to Google Cloud Storage.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

Although a Google Cloud Interconnect provides better performance and lower data transfer charges, it's not required between the ONTAP cluster and Google Cloud Storage. Because performance is significantly better when using Google Cloud Interconnect, doing so is the recommended best practice.

- An inbound connection is required from the NetApp Service Connector, which resides in an Google Cloud Platform VPC.

A connection between the cluster and the Cloud Tiering service is not required.

- An intercluster LIF is required on each ONTAP node that hosts tiered volumes. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

IPspaces enable network traffic segregation, allowing for separation of client traffic for privacy and security. [Learn more about IPspaces.](#)

When you set up data tiering, Cloud Tiering prompts you for the IPspace to use. You should choose the IPspace that each LIF is associated with. That might be the "Default" IPspace or a custom IPspace that you created.

## Supported volumes and aggregates

The total number of volumes that Cloud Tiering can tier might be less than the number of volumes on your ONTAP system. That's because volumes can't be tiered from some aggregates. For example, you can't tier data from SnapLock volumes or from MetroCluster configurations. Refer to ONTAP documentation for [functionality or features not supported by FabricPool.](#)



Cloud Tiering supports FlexGroup volumes. Setup works the same as any other volume.

## Creating or switching Connectors

A Connector is required to tier data to the cloud. When tiering data to Google Cloud Storage, a Connector must be available in a Google Cloud Platform VPC. You'll either need to create a new Connector or make sure that the currently selected Connector resides in GCP.

- [Learn about Connectors](#)
- [Creating a Connector in GCP](#)
- [Switching between Connectors](#)

## Preparing networking for the Connector

Ensure that the Connector has the required networking connections.

### Steps

1. Ensure that the VPC where the Connector is installed enables the following connections:
  - An outbound internet connection to the Cloud Tiering service over port 443 (HTTPS)
  - An HTTPS connection over port 443 to Google Cloud Storage
  - An HTTPS connection over port 443 to your ONTAP clusters
2. Optional: Enable Private Google Access on the subnet where you plan to deploy the Service Connector.

[Private Google Access](#) is recommended if you have a direct connection from your ONTAP cluster to the VPC and you want communication between the Connector and Google Cloud Storage to stay in your virtual private network. Note that Private Google Access works with VM instances that have only internal (private) IP addresses (no external IP addresses).

## Preparing Google Cloud Storage for data tiering

When you set up tiering, you need to provide storage access keys for a service account that has Storage Admin permissions. A service account enables Cloud Tiering to authenticate and access Cloud Storage buckets used for data tiering. The keys are required so that Google Cloud Storage knows who is making the request.



## Steps

1. [Create a service account that has the predefined Storage Admin role.](#)
2. Go to [GCP Storage Settings](#) and create access keys for the service account:
  - a. Select a project, and click **Interoperability**. If you haven't already done so, click **Enable interoperability access**.
  - b. Under **Access keys for service accounts**, click **Create a key for a service account**, select the service account that you just created, and click **Create Key**.

You'll need to [enter the keys in Cloud Tiering](#) later when you set up tiering.

## Tiering inactive data from your first cluster to Google Cloud Storage

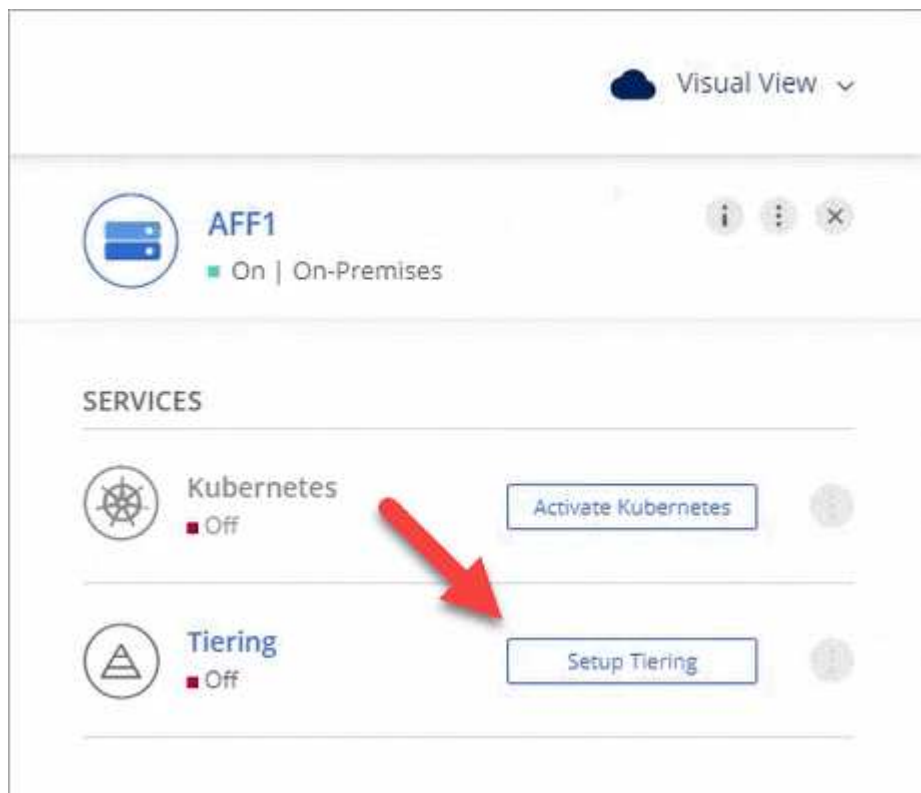
After you prepare your Google Cloud environment, start tiering inactive data from your first cluster.

### What you'll need

- [An on-premises working environment.](#)
- Storage access keys for a service account that has the Storage Admin role.

## Steps

1. Select an on-prem cluster.
2. Click **Setup Tiering**.




You're now on the Tiering dashboard.

3. Click **Set up Tiering** next to the cluster.
4. Complete the steps on the **Tiering Setup** page:

- a. **Bucket:** Add a new Google Cloud Storage bucket or select an existing bucket and click **Continue**.
- b. **Storage Class:** Select the storage class that you want to use for the tiered data and click **Continue**.
- c. **Credentials:** Enter the storage access key and secret key for a service account that has the Storage Admin role.
- d. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and click **Continue**.

Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to your cloud provider's object storage.

5. Click **Continue** to select the volumes that you want to tier.
6. On the **Tier Volumes** page, set up tiering for each volume. Click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

[Learn more about volume tiering policies.](#)

Tier Volumes							
Volume Name	SVM Name	Volume Size	Used Size	Cold Data		Tier Status [3]	Tiering Policy
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ...	70 %	<span>✓</span> Tiered Volume	All user data
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ...	70 %	<span>✓</span> Tiered Volume	Cold user data
vol3	svm_AFF1	325 GB	2.59 MB	0 B	0 %	<span>✓</span> Tiered Volume	Cold snapshots

## Result

You've successfully set up data tiering from volumes on the cluster to Google Cloud object storage.

## What's next?

[Be sure to subscribe from the Cloud Tiering service.](#)

You can also add additional clusters or review information about the active and inactive data on the cluster. For details, see [Managing data tiering from your clusters.](#)

## Tiering data from on-premises ONTAP clusters to StorageGRID

Free space on your on-prem ONTAP clusters by tiering data to StorageGRID. Data tiering is powered by NetApp's Cloud Tiering service.

### Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.



### 1 Prepare to tier data to StorageGRID

You need the following:

- An AFF or FAS system with all-SSD aggregates that's running ONTAP 9.4 or later, and a connection over a user-specified port to StorageGRID.
- StorageGRID 10.3 or later with AWS access keys that have S3 permissions.
- A Connector installed on your premises.
- Networking for the Connector that enables an outbound HTTPS connection to the ONTAP cluster, to StorageGRID, and to the Cloud Tiering service.

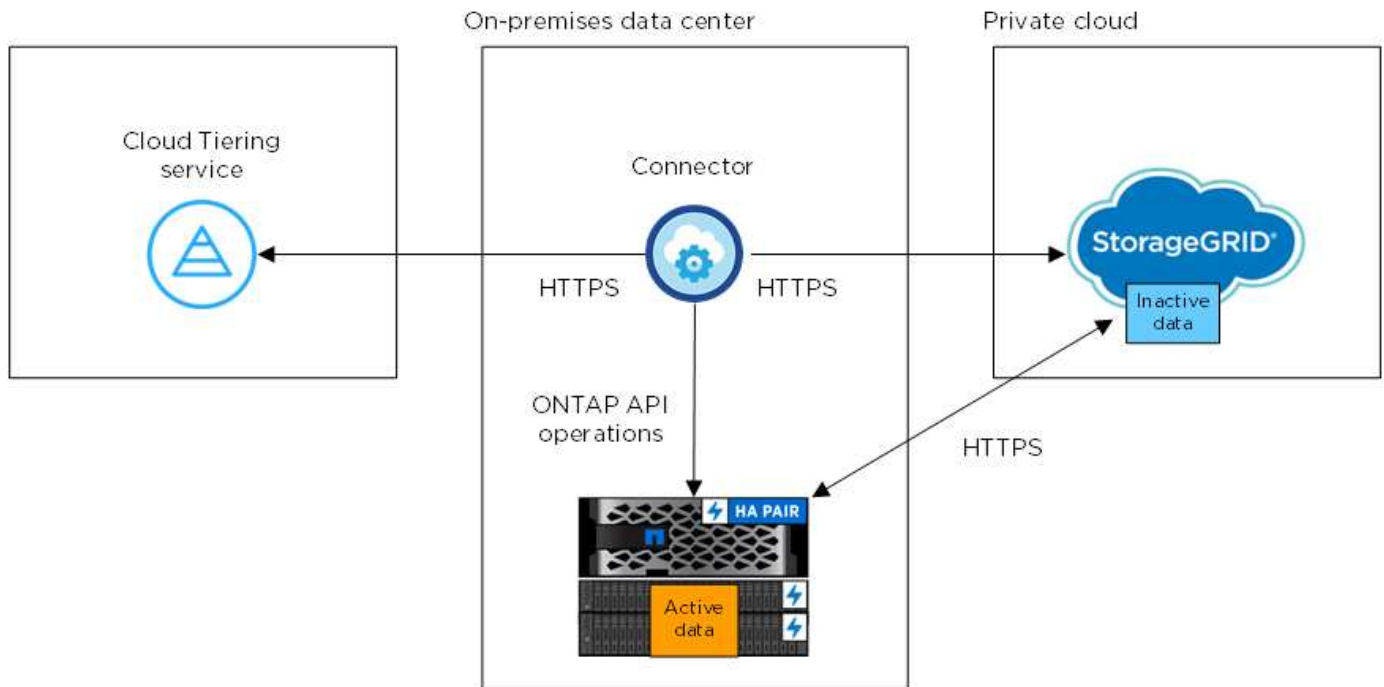
## 2 Set up tiering

Select an on-prem working environment, click **Setup Tiering** and follow the prompts to tier data to StorageGRID.

### Requirements

Verify support for your ONTAP cluster, set up your networking, and prepare your object storage.

The following image shows each component and the connections that you need to prepare between them:



Communication between the Connector and StorageGRID is for object storage setup only.

### Preparing your ONTAP clusters

Your ONTAP clusters must meet the following requirements when tiering data to StorageGRID.

### Supported ONTAP platforms

Cloud Tiering supports AFF systems and all-SSD aggregates on FAS systems.

### Supported ONTAP version

ONTAP 9.4 or later

## Licensing

A FabricPool license isn't required on the ONTAP cluster when tiering data to StorageGRID.

## Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over a user-specified port to StorageGRID (the port is configurable during tiering setup).

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

- An inbound connection is required from the Connector, which must reside on your premises.

A connection between the cluster and the Cloud Tiering service is not required.

- An intercluster LIF is required on each ONTAP node that hosts tiered volumes. The LIF must be associated with the *IPspace* that ONTAP should use to connect to object storage.

IPspaces enable network traffic segregation, allowing for separation of client traffic for privacy and security. [Learn more about IPspaces.](#)

When you set up data tiering, Cloud Tiering prompts you for the IPspace to use. You should choose the IPspace that each LIF is associated with. That might be the "Default" IPspace or a custom IPspace that you created.

## Supported volumes and aggregates

The total number of volumes that Cloud Tiering can tier might be less than the number of volumes on your ONTAP system. That's because volumes can't be tiered from some aggregates. For example, you can't tier data from SnapLock volumes or from MetroCluster configurations. Refer to ONTAP documentation for [functionality or features not supported by FabricPool.](#)



Cloud Tiering supports FlexGroup volumes, starting with ONTAP 9.5. Setup works the same as any other volume.

## Preparing StorageGRID

StorageGRID must meet the following requirements.

### Supported StorageGRID versions

StorageGRID 10.3 and later are supported.

### S3 credentials

When you set up tiering to StorageGRID, you need to provide Cloud Tiering with an S3 access key and secret key. Cloud Tiering uses the keys to access your buckets.

These access keys must be associated with a user who has the following permissions:

```
"s3:ListAllMyBuckets",  
"s3:ListBucket",  
"s3:GetObject",  
"s3:PutObject",  
"s3:DeleteObject",  
"s3:CreateBucket"
```

## Object versioning

You must not enable StorageGRID object versioning on the object store bucket.

### Creating or switching Connectors

A Connector is required to tier data to the cloud. When tiering data to StorageGRID, a Connector must be available on your premises. You'll either need to install a new Connector or make sure that the currently selected Connector resides on-prem.

- [Learn about Connectors](#)
- [Connector host requirements](#)
- [Installing the Connector on an existing Linux host](#)
- [Switching between Connectors](#)

### Preparing networking for the Connector

Ensure that the Connector has the required networking connections.

#### Steps

1. Ensure that the network where the Connector is installed enables the following connections:
  - An outbound internet connection to the Cloud Tiering service over port 443 (HTTPS)
  - An HTTPS connection over port 443 to StorageGRID
  - An HTTPS connection over port 443 to your ONTAP clusters

## Tiering inactive data from your first cluster to StorageGRID


After you prepare your environment, start tiering inactive data from your first cluster.

#### What you'll need

- [An on-premises working environment.](#)
- An AWS access key that has the required S3 permissions.

#### Steps


1. Select an on-prem cluster.
2. Click **Setup Tiering**.



You're now on the Tiering dashboard.

3. Click **Set up Tiering** next to the cluster.
4. Complete the steps on the **Tiering Setup** page:
  - a. **Choose your provider:** Select StorageGRID.
  - b. **Server:** Enter the FQDN of the StorageGRID server, enter the port that ONTAP should use for HTTPS communication with StorageGRID, and enter the access key and secret key for an AWS account that has the required S3 permissions.
  - c. **Bucket:** Add a new bucket or select an existing bucket for the tiered data.
  - d. **Cluster Network:** Select the IPspace that ONTAP should use to connect to object storage and click **Continue**.

Selecting the correct IPspace ensures that Cloud Tiering can set up a connection from ONTAP to your cloud provider's object storage.

5. Click **Continue** to select the volumes that you want to tier.
6. On the **Tier Volumes** page, set up tiering for each volume. Click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

[Learn more about volume tiering policies.](#)

Tier Volumes Learn how much you can save with each Tiering Policy

1 - 3 of 3 Volumes 🔍

Volume Name ↑	SVM Name ↑	Volume Size ↑	Used Size ↑	Cold Data ↑	Tier Status [3] ⇅	Tiering Policy ↑
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ... 70 %	✓ Tiered Volume	All user data
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ... 70 %	✓ Tiered Volume	Cold user data
vol3	svm_AFF1	325 GB	2.59 MB	0 B 0 %	✓ Tiered Volume	Cold snapshots

## Result

You've successfully set up data tiering from volumes on the cluster to StorageGRID.

## What's next?

You can add additional clusters or review information about the active and inactive data on the cluster. For details, see [Managing data tiering from your clusters](#).

# Set up licensing for Cloud Tiering

Pay for Cloud Tiering through a pay-as-you-go subscription, an ONTAP tiering license called *FabricPool*, or a combination of both. If you want to pay as you go, then you need to subscribe from the marketplace for the cloud provider to which you want to tier cold data. There's no need to subscribe from every marketplace.

A few notes before you read any further:

- If a FabricPool license is already installed on your cluster, then you're all set—there's nothing else that you need to do.
- If you've already subscribed to the Cloud Manager subscription in your cloud provider's marketplace, then you're automatically subscribed to Cloud Tiering, as well. You'll see an active subscription in the Cloud Tiering **Licensing** tab. You won't need to subscribe again.
- There are no charges when tiering data to StorageGRID. Neither a BYOL license or PAYGO registration is required.

[Learn more about how licensing works for Cloud Tiering.](#)

## Subscribing from the AWS Marketplace

Subscribe to Cloud Tiering from the AWS Marketplace to set up a pay-as-you-go subscription for data tiering from ONTAP clusters to AWS S3.

### Steps

1. In Cloud Manager, click **Tiering > Licensing**.
2. Click **Subscribe** under AWS Marketplace and then click **Continue**.
3. Subscribe from the AWS Marketplace, and then log back in to Cloud Central to complete the registration.

The following video shows the process:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_aws\\_tiering.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_aws_tiering.mp4) (video)

## Subscribing from the Azure Marketplace

Subscribe to Cloud Tiering from the Azure Marketplace to set up a pay-as-you-go subscription for data tiering from ONTAP clusters to Azure Blob storage.

### Steps

1. In Cloud Manager, click **Tiering > Licensing**.
2. Click **Subscribe** under Azure Marketplace and then click **Continue**.
3. Subscribe from the Azure Marketplace, and then log back in to Cloud Central to complete the registration.

The following video shows the process:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_azure\\_tiering.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_azure_tiering.mp4) (video)

## Subscribing from the GCP Marketplace

Subscribe to Cloud Tiering from the GCP Marketplace to set up a pay-as-you-go subscription for data tiering from ONTAP clusters to Google Cloud storage.

### Steps

1. In Cloud Manager, click **Tiering > Licensing**.
2. Click **Subscribe** under GCP Marketplace and then click **Continue**.
3. Subscribe from the GCP Marketplace, and then log back in to Cloud Central to complete the registration.

The following video shows the process:

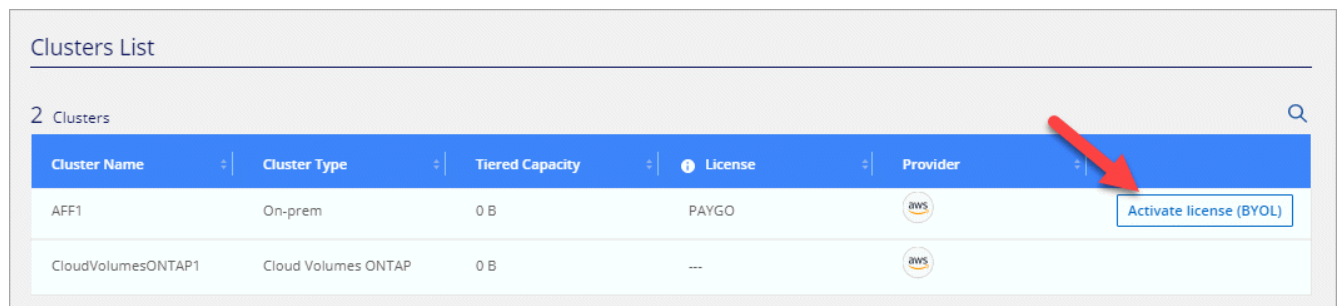
► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_gcp\\_tiering.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_gcp_tiering.mp4) (video)



## Adding a tiering license to ONTAP

Bring your own license by purchasing an ONTAP FabricPool license from NetApp.

### Steps

1. If you don't have a FabricPool license, [contact us to purchase one](#).
2. In Cloud Manager, click **Tiering > Licensing**.
3. In the Clusters List table, click **Activate license (BYOL)** for an on-prem ONTAP cluster.



Cluster Name	Cluster Type	Tiered Capacity	License	Provider	
AFF1	On-prem	0 B	PAYGO		<a href="#">Activate license (BYOL)</a>
CloudVolumesONTAP1	Cloud Volumes ONTAP	0 B	---		



4. Enter the serial number of the license and then enter the NetApp Support Site account that's associated with the serial number.
5. Click **Activate license**.

## Result

Cloud Tiering registers the license and installs it on the cluster.

## After you finish

If you purchase additional add-on capacity at a later time, the license on the cluster is automatically updated with the new capacity. There's no need to apply a new NetApp License File (NLF) to the cluster.


# Managing data tiering from your clusters

Now that you've set up data tiering from your ONTAP clusters, you can tier data from additional volumes, change a volume's tiering policy, and more.

## Tiering data from additional volumes

Set up data tiering for additional volumes at any time—for example, after creating a new volume.

### Steps

1. At the top of Cloud Manager, click **Tiering**.
2. From the **Cluster Dashboard**, click **Tier Volumes** for the cluster.
3. For each volume, click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

[Learn more about volume tiering policies.](#)

Tier Volumes							Learn how much you can save with each Tiering Policy
1 - 3 of 3 Volumes							Q
Volume Name	SVM Name	Volume Size	Used Size	Cold Data		Tier Status [3]	Tiering Policy
vol1	svm_AFF1	200 GB	3.8 MB	2.66 ...	70 %	Tiered Volume	All user data
vol2	svm_AFF1	400 GB	2.59 MB	1.81 ...	70 %	Tiered Volume	Cold user data
vol3	svm_AFF1	325 GB	2.59 MB	0 B	0 %	Tiered Volume	Cold snapshots



You don't need to configure the object storage because it was already configured when you initially set up tiering for the cluster. ONTAP will tier inactive data from these volumes to the same object store.


4. When you're done, click **Close**.

## Changing a volume's tiering policy

Changing the tiering policy for a volume changes how ONTAP tiers cold data to object storage. The change starts from the moment that you change the policy—it changes only the subsequent tiering behavior for the

volume.

### Steps

1. At the top of Cloud Manager, click **Tiering**.
2. From the **Cluster Dashboard**, click **Tier Volumes** for the cluster.
3. Click the  icon, select a tiering policy, optionally adjust the cooling days, and click **Apply**.

[Learn more about volume tiering policies.](#)

## Managing tiering settings on aggregates

Each aggregate has two settings that you can adjust: the tiering fullness threshold and whether inactive data reporting is enabled.

### Tiering fullness threshold

Setting the threshold to a lower number reduces the amount of data required to be stored on the performance tier before tiering takes place. This might be useful for large aggregates that contain little active data.

Setting the threshold to a higher number increases the amount of data required to be stored on the performance tier before tiering takes place. This might be useful for solutions designed to tier only when aggregates are near maximum capacity.

### Inactive data reporting


Inactive data reporting (IDR) uses a 31-day cooling period to determine which data is considered inactive. The amount of cold data that is tiered is dependent on the tiering policies set on volumes. This amount might be different than the amount of cold data detected by IDR using a 31-day cooling period.




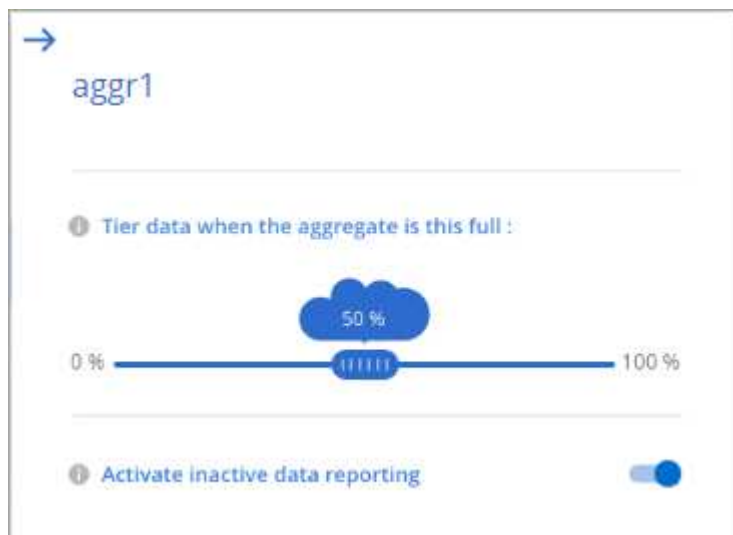
It's best to keep IDR enabled because it helps to identify your inactive data and savings opportunities. IDR must remain enabled if data tiering was enabled on an aggregate.

### Steps

1. At the top of Cloud Manager, click **Tiering**.
2. From the **Cloud Tiering** page, click the menu icon for a cluster and select **Manage Aggregates**.



3. On the **Manage Aggregates** page, click the  icon for an aggregate in the table.
4. Modify the fullness threshold and choose whether to enable or disable inactive data reporting.




5. Click **Apply**.

## Reviewing tiering info for a cluster

You might want to see how much data is in the cloud tier and how much data is on disks. Or, you might want to see the amount of hot and cold data on the cluster's disks. Cloud Tiering provides this information for each cluster.

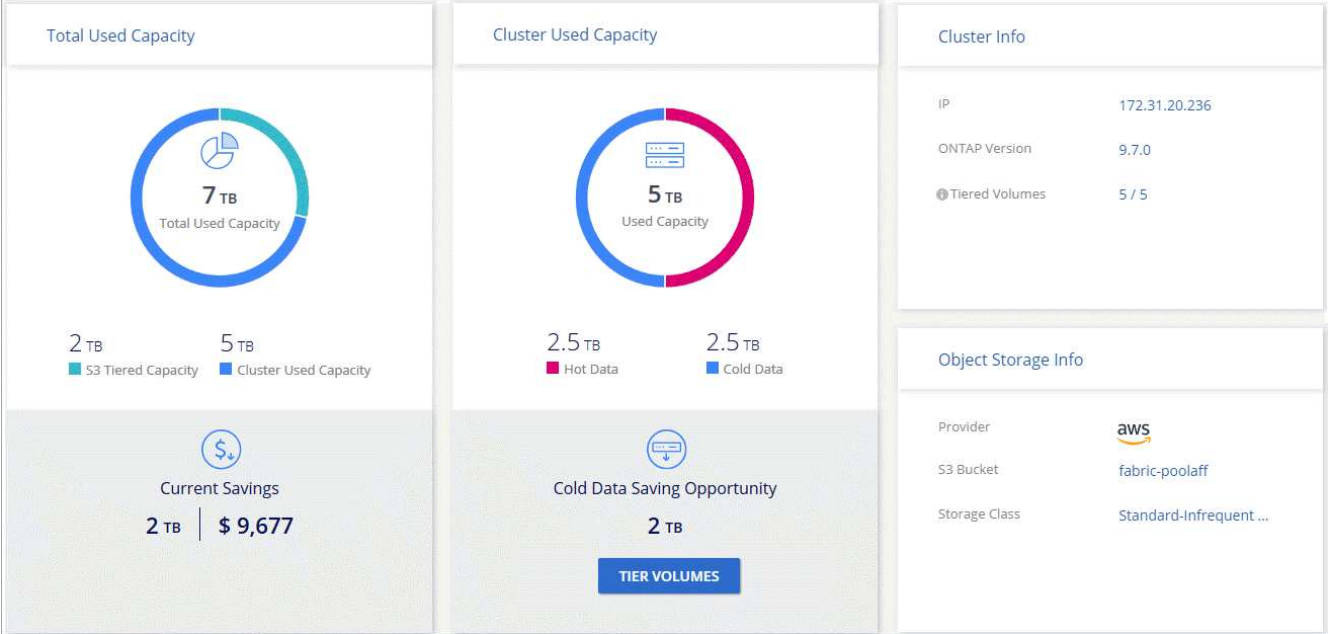
### Steps

1. At the top of Cloud Manager, click **Tiering**.
2. From the **Cluster Dashboard**, click **More info** for a cluster.



3. Review details about the cluster.

Here's an example:




## Fixing operational health

Failures can happen. When they do, Cloud Tiering displays a "Failed" operational health status on the Cluster Dashboard. The health reflects the status of the ONTAP system and Cloud Manager.

### Steps

1. Identify any clusters that have an operational health of "Failed."



2. Hover over the **i** icon to see the failure reason.
3. Correct the issue:
  - a. Verify that the ONTAP cluster is operational and that it has an inbound and outbound connection to your object storage provider.
  - b. Verify that Cloud Manager has outbound connections to the Cloud Tiering service, to the object store, and to the ONTAP clusters that it discovers.

## Cloud Tiering technical FAQ

This FAQ can help if you're just looking for a quick answer to a question.

### ONTAP

The following questions relate to ONTAP.

#### What are the requirements for my ONTAP cluster?

It depends on where you tier the cold data. Refer to the following:

- [Tiering data from on-premises ONTAP clusters to Amazon S3](#)
- [Tiering data from on-premises ONTAP clusters to Azure Blob storage](#)
- [Tiering data from on-premises ONTAP clusters to Google Cloud Storage](#)
- [Tiering data from on-premises ONTAP clusters to StorageGRID](#)

#### Does Cloud Tiering enable inactive data reporting?

Yes, Cloud Tiering enables inactive data reporting on each aggregate. This setting enables us to identify the amount of inactive data that can be tiered to low-cost object storage.

## Can I tier data from NAS volumes and SAN volumes?

You can use Cloud Tiering to tier data from NAS volumes to the public cloud and from SAN volumes to a private cloud using StorageGRID.

## What about Cloud Volumes ONTAP?

If you have Cloud Volumes ONTAP systems, you'll find them in the Cluster Dashboard so you see a full view of data tiering in your hybrid cloud infrastructure.

From the Cluster Dashboard, you can view tiering information similar to an on-prem ONTAP cluster: operational health, current savings, savings opportunities, details about volumes and aggregates, and more.

Cloud Volumes ONTAP systems are read-only from Cloud Tiering. You can't set up data tiering on Cloud Volumes ONTAP from Cloud Tiering. You'll still set up tiering the same way: from the working environment in Cloud Manager.

## Object storage

The following questions relate to object storage.

### Which object storage providers are supported?

Amazon S3, Azure Blob storage, Google Cloud Storage, and StorageGRID using the S3 protocol are supported.

### Can I use my own bucket/container?

Yes, you can. When you set up data tiering, you have the choice to add a new bucket/container or to select an existing bucket/container.

### Which regions are supported?

- [Supported AWS regions](#)
- [Supported Azure regions](#)
- [Supported Google Cloud regions](#)

### Which S3 storage classes are supported?

Cloud Tiering supports data tiering to the *Standard*, *Standard-Infrequent Access*, *One Zone-IA*, or *Intelligent* storage class. See [Supported S3 storage classes](#) for more details.

### Which Azure Blob access tiers are supported?

Cloud Tiering automatically uses the *Hot* access tier for your inactive data.

### Which storage classes are supported for Google Cloud Storage?

Cloud Tiering uses the *Standard* storage class for inactive data.

### Does Cloud Tiering use one object store for the entire cluster or one per aggregate?

One object store for the entire cluster.

## Can I apply policies to my object store to move data around independent of tiering?

No, Cloud Tiering does not support object lifecycle management rules that move or delete data from object stores.

## Connectors

The following questions relate to Connectors.

### Where does the Connector need to be installed?

- When tiering data to S3, a Connector can reside in an AWS VPC or on your premises.
- When tiering data to Blob storage, a Connector must reside in an Azure VNet.
- When tiering data to Google Cloud Storage, a Connector must reside in a Google Cloud Platform VPC.
- When tiering data to StorageGRID, a Connector must reside on an on premises Linux host.

## Networking

The following questions relate to networking.

### What are the networking requirements?

- The ONTAP cluster initiates an HTTPS connection over port 443 to your object storage provider.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

- For StorageGRID, the ONTAP cluster initiates an HTTPS connection over a user-specified port to StorageGRID (the port is configurable during tiering setup).
- A Connector needs an outbound HTTPS connection over port 443 to your ONTAP clusters, to the object store, and to the Cloud Tiering service.

For more details, see:

- [Tiering data from on-premises ONTAP clusters to Amazon S3](#)
- [Tiering data from on-premises ONTAP clusters to Azure Blob storage](#)
- [Tiering data from on-premises ONTAP clusters to Google Cloud Storage](#)
- [Tiering data from on-premises ONTAP clusters to StorageGRID](#)

## Permissions

The following questions relate to permissions.

### What permissions are required in AWS?

Permissions are required [to manage the S3 bucket](#).

### What permissions are required in Azure?

No extra permissions are needed outside of the permissions that you need to provide to Cloud Manager.

## What permissions are required in Google Cloud Platform?

Storage Admin permissions are needed for a service account that has storage access keys.

## What permissions are required for StorageGRID?

[S3 permissions are needed.](#)

# Reference

## Supported S3 storage classes and regions

Cloud Tiering supports several S3 storage classes and most regions.

### Supported S3 storage classes

Cloud Tiering can apply a lifecycle rule so the data transitions from the *Standard* storage class to another storage class after 30 days. You can choose from the following storage classes:

- Standard-Infrequent Access
- One Zone-IA
- Intelligent

If you choose Standard, then the data remains in that storage class.

[Learn about S3 storage classes.](#)

### Supported AWS regions

Cloud Tiering supports the following AWS regions.

#### Asia Pacific

- Mumbai
- Seoul
- Singapore
- Sydney
- Tokyo

#### Europe

- Frankfurt
- Ireland
- London
- Paris
- Stockholm



## North America

- Canada Central
- GovCloud (US-West) – starting with ONTAP 9.3
- US East (N. Virginia)
- US East (Ohio)
- US West (N. California)
- US West (Oregon)

## South America

- São Paulo

## Supported Azure Blob access tiers and regions

Cloud Tiering supports the *Hot* access tier and most regions.

### Supported Azure Blob access tiers

When you set up data tiering to Azure, Cloud Tiering automatically uses the *Hot* access tier for your inactive data.

### Supported Azure regions

Cloud Tiering supports the following Azure regions.

#### Africa

- South Africa North

#### Asia Pacific

- Australia East
- Australia Southeast
- East Asia
- Japan East
- Japan West
- Korea Central
- Korea South
- Southeast Asia

#### Europe

- France Central
- Germany Central
- Germany Northeast
- North Europe

- UK South
- UK West
- West Europe

#### **North America**

- Canada Central
- Canada East
- Central US
- East US
- East US 2
- North Central US
- South Central US
- West US
- West US 2
- West Central US

#### **South America**

- Brazil South

## **Supported Google Cloud storage classes and regions**

Cloud Tiering supports the Standard storage class and most Google Cloud regions.

### **Supported access tiers**

Cloud Tiering uses the *Standard* access tier for your inactive data.

### **Supported Google Cloud regions**

Cloud Tiering supports the following regions.

#### **Americas**

- Iowa
- Los Angeles
- Montreal
- N. Virginia
- Oregon
- Sao-Paulo
- South Carolina

#### **Asia Pacific**

- Hong Kong

- Mumbai
- Osaka
- Singapore
- Sydney
- Taiwan
- Tokyo

#### **Europe**

- Belgium
- Finland
- Frankfurt
- London
- Netherlands
- Zurich

# Viewing your Amazon S3 buckets

After you install a Connector in AWS, Cloud Manager can automatically discover information about the Amazon S3 buckets that reside in the AWS account where it's installed.

You can see details about your S3 buckets, including the region, access level, storage class, and whether the bucket is used with Cloud Volumes ONTAP for backups or data tiering. And you can scan the S3 buckets with Cloud Compliance.

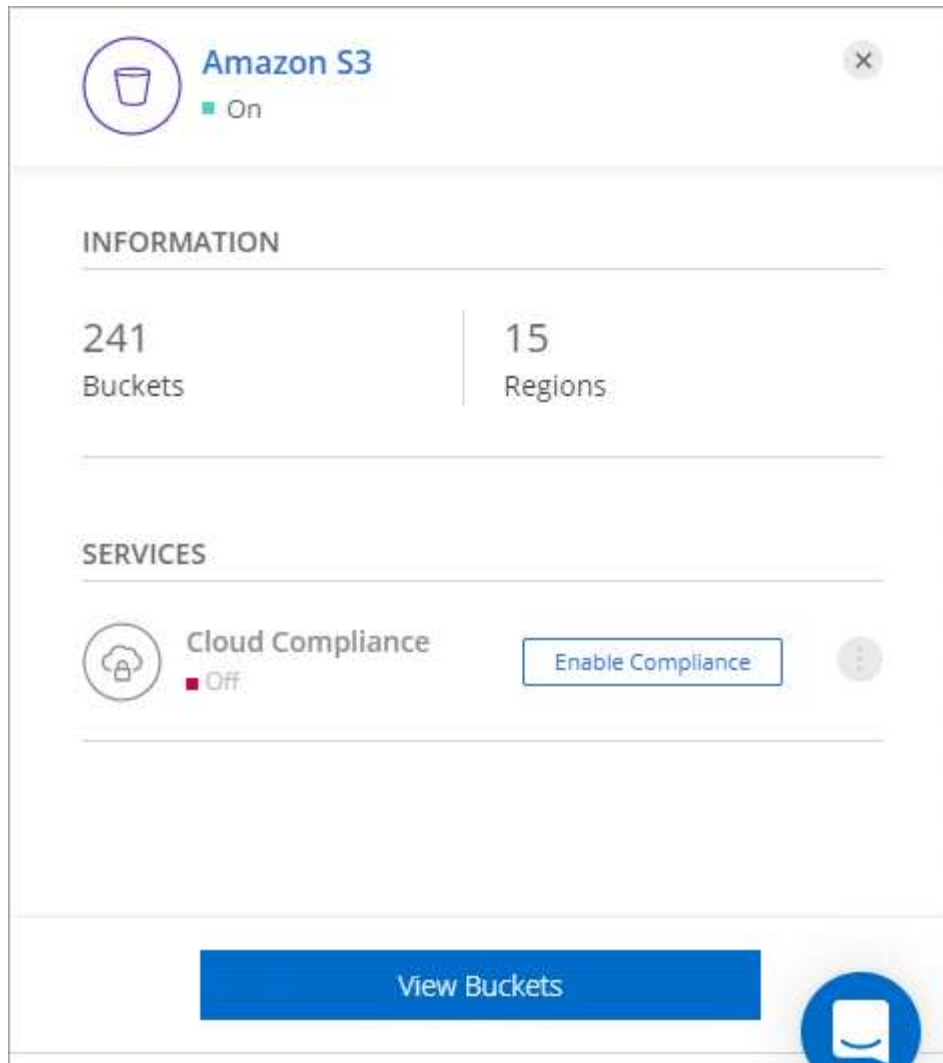
## Steps

1. [Install a Connector](#) in the AWS account where you want to view your Amazon S3 buckets.

You should automatically see an Amazon S3 working environment shortly after.



2. Click the working environment and select an action from the right pane.



3. Click **Enable Compliance** to scan the S3 buckets for personal and sensitive data.

For more details, see [Getting started with Cloud Compliance for Amazon S3](#).

4. Click **View Buckets** to view details about the S3 buckets in your AWS account.

# Administer Cloud Manager

## Finding your Cloud Manager system ID


To help you get started, your NetApp representative might ask you for your Cloud Manager system ID. The ID is typically used for licensing and troubleshooting purposes.

### What you'll need

You need to create a Connector before you can change Cloud Manager settings. [Learn how.](#)

### Steps

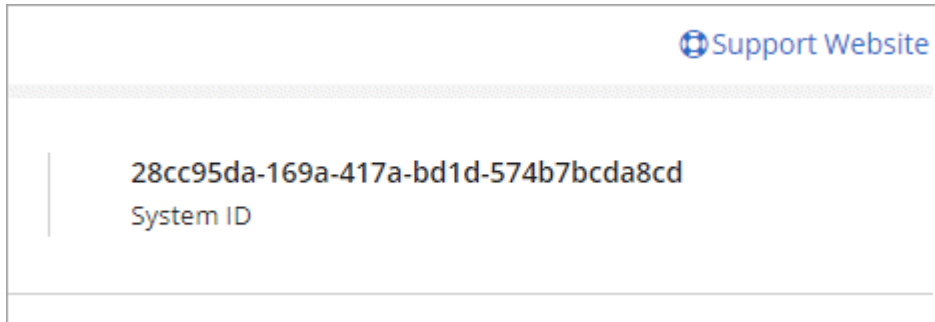
1. In the upper right of the Cloud Manager console, click the Settings icon.



2. Click **Support Dashboard**.

Your system ID appears in the top right.

### Example



## Manage Connectors

### Managing existing Connectors

After you create one or more Connectors, you can manage them by switching between Connectors, connecting to the local user interface running on a Connector, and more.

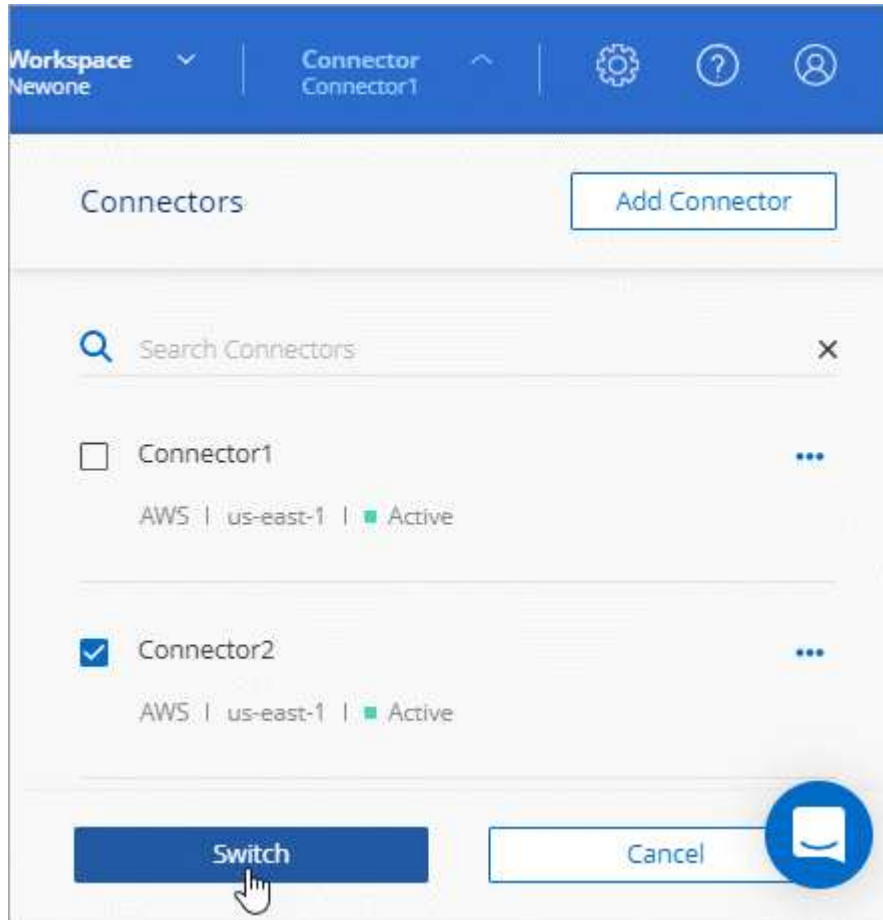
#### Switching between Connectors

If you have multiple Connectors, you can switch between them to see the Working Environments that are associated with a specific Connector.

For example, let's say that you're working in a multi-cloud environment. You might have one Connector in AWS and another in Google Cloud. You'd need to switch between those Connectors to manage the Cloud Volumes ONTAP systems running in those clouds.

## Step

1. Click the **Connector** drop-down, select another Connector, and then click **Switch**.



Cloud Manager refreshes and shows the Working Environments associated with the selected Connector.

## Accessing the local UI

While you should perform almost all tasks from the SaaS user interface, a local user interface is still available on the Connector. This interface is needed for a few tasks that need to be performed from the Connector itself:


- [Setting a proxy server](#)
- Installing a patch (you'll typically work with NetApp personnel to install a patch)
- Downloading AutoSupport messages (usually directed by NetApp personnel when you have issues)

## Steps

1. [Log in to the Cloud Manager SaaS interface](#) from a machine that has a network connection to the Connector instance.

If the Connector doesn't have a public IP address, you'll need a VPN connection or you'll need to connect from a jump host that's in the same network as the Connector.

2. Click the **Connector** drop-down, click the action menu for a Connector, and then click **Go to Local UI**.



The Cloud Manager interface running on the Connector loads in a new browser tab.

## Removing Connectors from Cloud Manager

If a Connector is inactive, you can remove it from the list of Connectors in Cloud Manager. You might do this if you deleted the Connector virtual machine or if you uninstalled the Connector software.

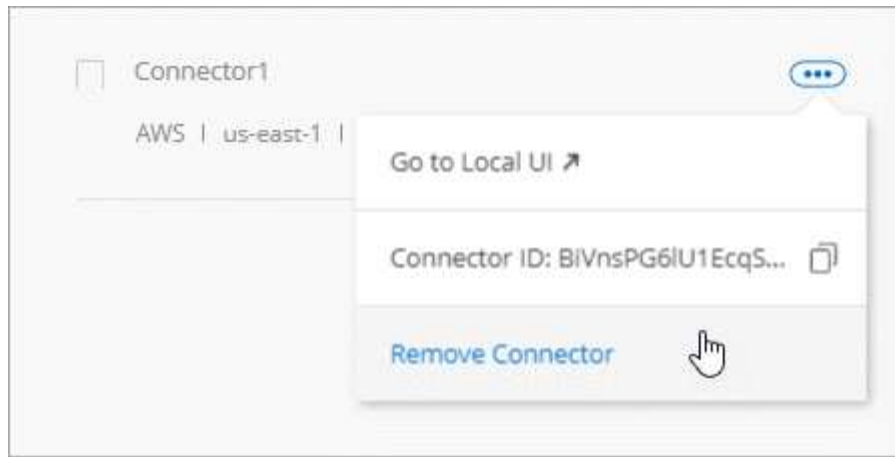
Note the following about removing a Connector:

- This action doesn't delete the virtual machine.
- This action can't be reverted—once you remove a Connector from Cloud Manager, you can't add it back to Cloud Manager.

### Steps

1. Click the Connector drop-down from the Cloud Manager header.
2. Click the action menu for an inactive Connector and click **Remove Connector**.





3. Enter the name of the Connector to confirm and then click Remove.

### Result

Cloud Manager removes the Connector from its records.

### Uninstalling the Connector software

The Connector includes an uninstallation script that you can use to uninstall the software to troubleshoot issues or to permanently remove the software from the host.

### Step

1. From the Linux host, run the uninstallation script:

```
/opt/application/netapp/cloudmanager/bin/uninstall.sh [silent]
```

*silent* runs the script without prompting you for confirmation.

### What about software upgrades?

The Connector automatically updates its software to the latest version, as long as it has [outbound internet access](#) to obtain the software update.

## More ways to create Connectors

### Connector host requirements

The Connector software must run on a host that meets specific operating system requirements, RAM requirements, port requirements, and so on.

### A dedicated host is required

The Connector is not supported on a host that is shared with other applications. The host must be a dedicated host.

### CPU

4 cores or 4 vCPUs

## RAM

14 GB

## AWS EC2 instance type

An instance type that meets the CPU and RAM requirements above. We recommend t3.xlarge and use that instance type when you deploy the Connector directly from Cloud Manager.

## Azure VM size

An instance type that meets the CPU and RAM requirements above. We recommend DS3 v2 and use that VM size when you deploy the Connector directly from Cloud Manager.

## GCP machine type

An instance type that meets the CPU and RAM requirements above. We recommend n1-standard-4 and use that machine type when you deploy the Connector directly from Cloud Manager.

## Supported operating systems

- CentOS 7.6
- CentOS 7.7
- Red Hat Enterprise Linux 7.6
- Red Hat Enterprise Linux 7.7

The Red Hat Enterprise Linux system must be registered with Red Hat Subscription Management. If it is not registered, the system cannot access repositories to update required 3rd party software during Connector installation.

The Connector is supported on English-language versions of these operating systems.

## Hypervisor

A bare metal or hosted hypervisor that is certified to run CentOS or Red Hat Enterprise Linux  
[Red Hat Solution: Which hypervisors are certified to run Red Hat Enterprise Linux?](#)

## Disk space in /opt

100 GB of space must be available

## Outbound internet access

Outbound internet access is required to install the Connector and for the Connector to manage resources and processes within your public cloud environment. For a list of endpoints, see [Networking requirements for the Connector](#).

## Creating a Connector from the AWS Marketplace

It's best to create a Connector directly from Cloud Manager, but you can launch a Connector from the AWS Marketplace, if you'd rather not specify AWS access keys. After you create and set up the Connector, Cloud Manager will automatically use it when you create new working environments.

## Steps

1. Create an IAM policy and role for the EC2 instance:
  - a. Download the Cloud Manager IAM policy from the following location:

## NetApp Cloud Manager: AWS, Azure, and GCP Policies

- b. From the IAM console, create your own policy by copying and pasting the text from the Cloud Manager IAM policy.
  - c. Create an IAM role with the role type Amazon EC2 and attach the policy that you created in the previous step to the role.
2. Now go to the [Cloud Manager page on the AWS Marketplace](#) to deploy Cloud Manager from an AMI.
- The IAM user must have AWS Marketplace permissions to subscribe and unsubscribe.
3. On the Marketplace page, click **Continue to Subscribe** and then click **Continue to Configuration**.

The screenshot displays the AWS Marketplace interface for NetApp Cloud Manager. The top navigation bar includes links for 'Delivery Methods', 'Solutions', 'Migration Mapping Assistant', 'Your Saved List' (with a '2' notification), 'Partners', 'Sell in AWS Marketplace', and 'Amazon Web Services Home'. The main content area features the NetApp logo and the product title 'Cloud Manager - Manual Installation without access keys'. Below the title, it states 'By: NetApp, Inc.' and 'Latest Version: 3.8.4'. A 'Continue to Subscribe' button is highlighted with a red arrow. To the right, a pricing box shows a 'Typical Total Price' of '\$0.226/hr'. Below the product overview, a 'Highlights' box contains a bullet point: 'See Product Overview for instructions on how to deploy NetApp Cloud Manager.' The bottom part of the screenshot shows the 'Continue to Configuration' button, also with a red arrow pointing to it.

4. Change any of the default options and click **Continue to Launch**.
5. Under **Choose Action**, select **Launch through EC2** and then click **Launch**.

These steps describe how to launch the instance from the EC2 Console because the console enables you to attach an IAM role to the Cloud Manager instance. This isn't possible using the **Launch from Website** action.

6. Follow the prompts to configure and deploy the instance:
  - **Choose Instance Type:** Depending on region availability, choose one of the supported instance types (t3.xlarge is recommended).

[Review the instance requirements.](#)

- **Configure Instance:** Select a VPC and subnet, choose the IAM role that you created in step 1, enable termination protection (recommended), and choose any other configuration options that meet your requirements.

<b>Number of instances</b> ⓘ	<input type="text" value="1"/>	<a href="#">Launch into Auto Scaling Group</a> ⓘ
<b>Purchasing option</b> ⓘ	<input type="checkbox"/> Request Spot instances	
<b>Network</b> ⓘ	<input type="text" value="vpc-a76d91c2   VPC4QA (default)"/>	<a href="#">Create new VPC</a>
<b>Subnet</b> ⓘ	<input type="text" value="subnet-39536c13   QASubnet1   us-east-1b"/> 155 IP Addresses available	<a href="#">Create new subnet</a>
<b>Auto-assign Public IP</b> ⓘ	<input type="text" value="Enable"/>	
<b>Placement group</b> ⓘ	<input type="checkbox"/> Add instance to placement group	
<b>Capacity Reservation</b> ⓘ	<input type="text" value="Open"/>	<a href="#">Create new Capacity Reservation</a>
<b>IAM role</b> ⓘ	<input type="text" value="Cloud_Manager"/>	<a href="#">Create new IAM role</a>
<b>CPU options</b> ⓘ	<input type="checkbox"/> Specify CPU options	
<b>Shutdown behavior</b> ⓘ	<input type="text" value="Stop"/>	
<b>Enable termination protection</b> ⓘ	<input checked="" type="checkbox"/> Protect against accidental termination	
<b>Monitoring</b> ⓘ	<input type="checkbox"/> Enable CloudWatch detailed monitoring <a href="#">Additional charges apply.</a>	

- **Add Storage:** Keep the default storage options.
- **Add Tags:** Enter tags for the instance, if desired.
- **Configure Security Group:** Specify the required connection methods for the Connector instance: SSH, HTTP, and HTTPS.
- **Review:** Review your selections and click **Launch**.

AWS launches the software with the specified settings. The Connector instance and software should be running in approximately five minutes.


7. Open a web browser from a host that has a connection to the Connector instance and enter the following URL:

http://ipaddress:80

8. After you log in, set up the Connector:
  - a. Specify the Cloud Central account to associate with the Connector.

[Learn about Cloud Central accounts.](#)

- b. Enter a name for the system.



### Result

The Connector is now installed and set up with your Cloud Central account. Cloud Manager will automatically use this Connector when you create new working environments. But if you have more than one Connector, you'll need to [switch between them](#).

### Creating a Connector from the Azure Marketplace

It's best to create a Connector directly from Cloud Manager, but you can launch a Connector from the Azure Marketplace, if you prefer. After you create and set up the Connector, Cloud Manager will automatically use it when you create new working environments.

## Creating a Connector in Azure

Deploy the Connector in Azure using the image in the Azure Marketplace and then log in to the Connector to specify your Cloud Central account.

### Steps

1. [Go to the Azure Marketplace page for Cloud Manager.](#)
2. Click **Get it now** and then click **Continue**.
3. From the Azure portal, click **Create** and follow the steps to configure the virtual machine.

Note the following as you configure the VM:

- Cloud Manager can perform optimally with either HDD or SSD disks.
- Choose a VM size that meets CPU and RAM requirements. We recommend DS3 v2.

[Review the VM requirements.](#)

- For the network security group, the Connector requires inbound connections using SSH, HTTP, and HTTPS.

[Learn more about security group rules for the Connector.](#)

- Under **Management**, enable **System assigned managed identity** for the Connector by selecting **On**.

This setting is important because a managed identity allows the Connector virtual machine to identify itself to Azure Active Directory without providing any credentials. [Learn more about managed identities for Azure resources.](#)

4. On the **Review + create** page, review your selections and click **Create** to start the deployment.

Azure deploys the virtual machine with the specified settings. The virtual machine and Connector software should be running in approximately five minutes.


5. Open a web browser from a host that has a connection to the Connector virtual machine and enter the following URL:

`http://ipaddress:80`

6. After you log in, set up the Connector:
  - a. Specify the Cloud Central account to associate with the Connector.

[Learn about Cloud Central accounts.](#)

- b. Enter a name for the system.



## Result

The Connector is now installed and set up. You must grant Azure permissions before users can deploy Cloud Volumes ONTAP in Azure.

## Granting Azure permissions

When you deployed the Connector in Azure, you should have enabled a [system-assigned managed identity](#). You must now grant the required Azure permissions by creating a custom role and then by assigning the role to the Connector virtual machine for one or more subscriptions.

## Steps

1. Create a custom role using the Cloud Manager policy:
  - a. Download the [Cloud Manager Azure policy](#).
  - b. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription from which users will create Cloud Volumes ONTAP systems.

## Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"
```



- c. Use the JSON file to create a custom role in Azure.

The following example shows how to create a custom role using the Azure CLI 2.0:

```
az role definition create --role-definition
C:\Policy_for_cloud_Manager_Azure_3.8.7.json
```

You should now have a custom role called Cloud Manager Operator that you can assign to the Connector virtual machine.

2. Assign the role to the Connector virtual machine for one or more subscriptions:
  - a. Open the **Subscriptions** service and then select the subscription in which you want to deploy Cloud Volumes ONTAP systems.
  - b. Click **Access control (IAM)**.
  - c. Click **Add > Add role assignment** and then add the permissions:
    - Select the **Cloud Manager Operator** role.



Cloud Manager Operator is the default name provided in the [Cloud Manager policy](#). If you chose a different name for the role, then select that name instead.

- Assign access to a **Virtual Machine**.
  - Select the subscription in which the Connector virtual machine was created.
  - Select the Connector virtual machine.
  - Click **Save**.
- d. If you want to deploy Cloud Volumes ONTAP from additional subscriptions, switch to that subscription and then repeat these steps.

## Result

The Connector now has the permissions that it needs to manage resources and processes within your public cloud environment. Cloud Manager will automatically use this Connector when you create new working environments. But if you have more than one Connector, you'll need to [switch between them](#).

## Installing the Connector software on an existing Linux host

The most common way to create a Connector is directly from Cloud Manager or from a cloud provider's marketplace. But you have the option to download and install the Connector software on an existing Linux host in your network or in the cloud.



If you want to create a Cloud Volumes ONTAP system in Google Cloud, then you must have a Connector running in Google Cloud, as well. You can't use a Connector that's running in another location.

## Requirements

- The host must meet [requirements for the Connector](#).
- A Red Hat Enterprise Linux system must be registered with Red Hat Subscription Management. If it is not registered, the system cannot access repositories to update required 3rd party software during installation.
- The Connector installer accesses several URLs during the installation process. You must ensure that outbound internet access is allowed to these endpoints:



- <http://dev.mysql.com/get/mysql-community-release-el7-5.noarch.rpm>
- <https://dl.fedoraproject.org/pub/epel/epel-release-latest-7.noarch.rpm>
- <https://s3.amazonaws.com/aws-cli/awscli-bundle.zip>

The host might try to update operating system packages during installation. The host can contact different mirroring sites for these OS packages.

### About this task

- Root privileges are not required to install the Connector.
- The installation installs the AWS command line tools (awscli) to enable recovery procedures from NetApp support.

If you receive a message that installing the awscli failed, you can safely ignore the message. The Connector can operate successfully without the tools.

- The installer that is available on the NetApp Support Site might be an earlier version. After installation, the Connector automatically updates itself if a new version is available.

### Steps

1. Download the Cloud Manager software from the [NetApp Support Site](#), and then copy it to the Linux host.

For help with connecting and copying the file to an EC2 instance in AWS, see [AWS Documentation: Connecting to Your Linux Instance Using SSH](#).

2. Assign permissions to execute the script.

### Example

```
chmod +x OnCommandCloudManager-V3.8.4.sh
```

3. Run the installation script:

```
./OnCommandCloudManager-V3.8.4.sh [silent] [proxy=ipaddress]  
[proxyport=port] [proxyuser=user_name] [proxypwd=password]
```

*silent* runs the installation without prompting you for information.

*proxy* is required if the host is behind a proxy server.

*proxyport* is the port for the proxy server.

*proxyuser* is the user name for the proxy server, if basic authentication is required.

*proxypwd* is the password for the user name that you specified.

4. Unless you specified the silent parameter, type **Y** to continue the script, and then enter the HTTP and HTTPS ports when prompted.

Cloud Manager is now installed. At the end of the installation, the Cloud Manager service (occm) restarts

twice if you specified a proxy server.

5. Open a web browser and enter the following URL:

`https://ipaddress:port`

*ipaddress* can be localhost, a private IP address, or a public IP address, depending on the configuration of the host. For example, if the Connector is in the public cloud without a public IP address, you must enter a private IP address from a host that has a connection to the Connector host.

*port* is required if you changed the default HTTP (80) or HTTPS (443) ports. For example, if the HTTPS port was changed to 8443, you would enter `https://ipaddress:8443`

6. Sign up at NetApp Cloud Central or log in.
7. After you log in, set up Cloud Manager:
  - a. Specify the Cloud Central account to associate with the Connector.  
[Learn about Cloud Central accounts.](#)
  - b. Enter a name for the system.



## Result

The Connector is now installed and set up with your Cloud Central account. Cloud Manager will automatically use this Connector when you create new working environments.

## After you finish

Set up permissions so Cloud Manager can manage resources and processes within your public cloud environment:

- AWS: [Set up an AWS account and then add it to Cloud Manager](#).
- Azure: [Set up an Azure account and then add it to Cloud Manager](#).
- GCP: Set up a service account that has the permissions that Cloud Manager needs to create and manage Cloud Volumes ONTAP systems in projects.
  1. [Create a role in GCP](#) that includes the permissions defined in the [Cloud Manager policy for GCP](#).
  2. [Create a GCP service account and apply the custom role that you just created](#).
  3. [Associate this service account with the Connector VM](#).
  4. If you want to deploy Cloud Volumes ONTAP in other projects, [grant access by adding the service account with the Cloud Manager role to that project](#). You'll need to repeat this step for each project.

## Default configuration for the Connector

If you need to troubleshoot the Connector, it might help to understand how it's configured.

- If you deployed the Connector from Cloud Manager (or directly from a cloud provider's marketplace), note the following:
  - In AWS, the user name for the EC2 Linux instance is `ec2-user`.
  - The operating system for the image is as follows:
    - AWS: Red Hat Enterprise Linux 7.5 (HVM)
    - Azure: Red Hat Enterprise Linux 7.6 (HVM)
    - GCP: CentOS 7.6

The operating system does not include a GUI. You must use a terminal to access the system.

- The Connector installation folder resides in the following location:

```
/opt/application/netapp/cloudmanager
```

- Log files are contained in the following folder:

```
/opt/application/netapp/cloudmanager/log
```

- The Cloud Manager service is named `occm`.
- The `occm` service is dependent on the MySQL service.

If the MySQL service is down, then the `occm` service is down too.

- Cloud Manager installs the following packages on the Linux host, if they are not already installed:
  - 7Zip
  - AWSCLI
  - Docker
  - Java

- Kubectl
- MySQL
- Tridentctl
- Pull
- Wget
- The Connector uses the following ports on the Linux host:
  - 80 for HTTP access
  - 443 for HTTPS access
  - 3306 for the Cloud Manager database
  - 8080 for the Cloud Manager API proxy
  - 8666 for the Service Manager API
  - 8777 for the Health-Checker Container Service API

## Manage credentials

### AWS

#### AWS credentials and permissions

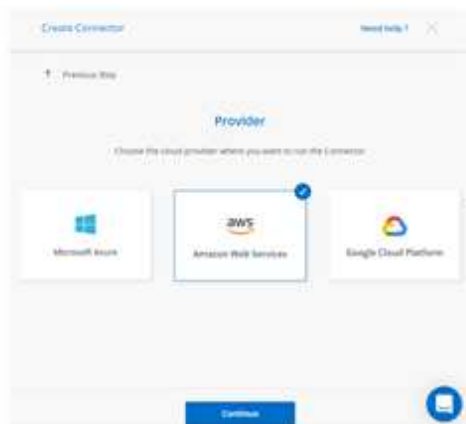
Cloud Manager enables you to choose the AWS credentials to use when deploying Cloud Volumes ONTAP. You can deploy all of your Cloud Volumes ONTAP systems using the initial AWS credentials, or you can add additional credentials.

#### Initial AWS credentials


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

When Cloud Manager launches the Connector 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 manage resources and processes within that AWS account. [Review how Cloud Manager uses the permissions](#).


#### Cloud Manager



#### AWS account

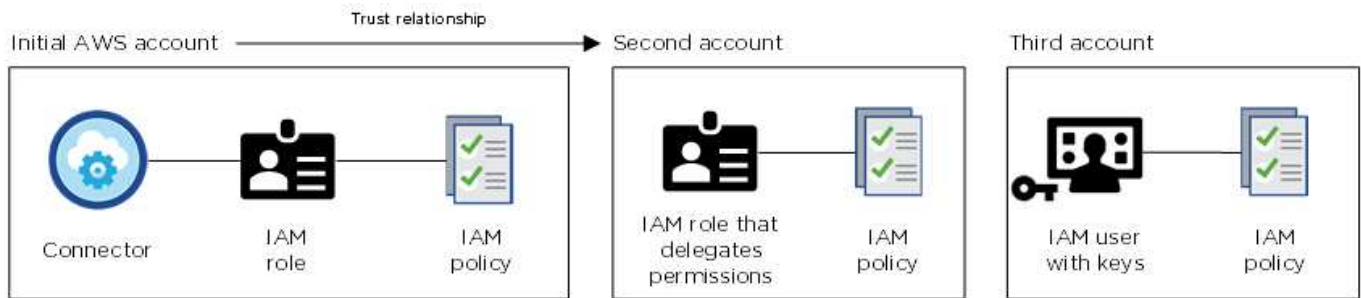


Cloud Manager selects these AWS credentials by default when you create a new working environment for Cloud Volumes ONTAP:




### Additional AWS credentials

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 account credentials 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 set of credentials, you can switch to them when creating a new working environment:



### What about Marketplace deployments and on-prem deployments?

The sections above describe the recommended deployment method for the Connector, which is from Cloud Manager. You can also deploy a Connector in AWS from the [AWS Marketplace](#) and you can [install the Connector 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.

### How can I securely rotate my AWS credentials?

As described above, Cloud Manager enables you to provide AWS credentials in a few ways: an IAM role associated with the Connector instance, by assuming an IAM role in a trusted account, or by providing AWS

access keys.

With the first two options, Cloud Manager uses the AWS Security Token Service to obtain temporary credentials that rotate constantly. This process is the best practice—it's automatic and it's secure.

If you provide Cloud Manager with AWS access keys, you should rotate the keys by updating them in Cloud Manager at a regular interval. This is a completely manual process.

## Managing AWS credentials and subscriptions for Cloud Manager

When you create a Cloud Volumes ONTAP system, you need to select the AWS credentials and subscription to use with that system. If you manage multiple AWS subscriptions, you can assign each one of them to different AWS credentials from the Credentials page.

Before you add AWS credentials to Cloud Manager, you need to provide the required permissions to that account. The permissions enable Cloud Manager to manage resources and processes within that AWS account. How you provide the permissions depends on whether you want to provide Cloud Manager with AWS keys or the ARN of a role in a trusted account.



When you deployed a Connector from Cloud Manager, Cloud Manager automatically added AWS credentials for the account in which you deployed the Connector. This initial account is not added if you manually installed the Connector software on an existing system. [Learn about AWS credentials and permissions.](#)

## Choices

- [Granting permissions by providing AWS keys](#)
- [Granting permissions by assuming IAM roles in other accounts](#)

### How can I securely rotate my AWS credentials?

Cloud Manager enables you to provide AWS credentials in a few ways: an IAM role associated with the Connector instance, by assuming an IAM role in a trusted account, or by providing AWS access keys. [Learn more about AWS credentials and permissions.](#)

With the first two options, Cloud Manager uses the AWS Security Token Service to obtain temporary credentials that rotate constantly. This process is the best practice, it's automatic and it's secure.

If you provide Cloud Manager with AWS access keys, you should rotate the keys by updating them in Cloud Manager at a regular interval. This is a completely manual process.

## Granting permissions by providing AWS keys

If you want to provide Cloud Manager with AWS keys for an IAM user, then you need to grant the required permissions to that user. The Cloud Manager IAM policy defines the AWS actions and resources that Cloud Manager is allowed to use.

## Steps

1. Download the Cloud Manager IAM policy from the [Cloud Manager Policies page](#).
2. From the IAM console, create your own policy by copying and pasting the text from the Cloud Manager

IAM policy.

[AWS Documentation: Creating IAM Policies](#)

3. Attach the policy to an IAM role or an IAM user.
  - [AWS Documentation: Creating IAM Roles](#)
  - [AWS Documentation: Adding and Removing IAM Policies](#)

## Result

The account now has the required permissions. [You can now add it to Cloud Manager.](#)

## Granting permissions by assuming IAM roles in other accounts

You can set up a trust relationship between the source AWS account in which you deployed the Connector instance and other AWS accounts by using IAM roles. You would then provide Cloud Manager with the ARN of the IAM roles from the trusted accounts.

## Steps

1. Go to the target account where you want to deploy Cloud Volumes ONTAP and create an IAM role by selecting **Another AWS account**.

Be sure to do the following:

- Enter the ID of the account where the Connector instance resides.
- Attach the Cloud Manager IAM policy, which is available from the [Cloud Manager Policies page](#).

2. Go to the source account where the Connector instance resides and select the IAM role that is attached to the instance.
  - a. Click **Attach policies** and then click **Create policy**.
  - b. Create a policy that includes the "sts:AssumeRole" action and the ARN of the role that you created in the target account.

## Example

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": "sts:AssumeRole",
    "Resource": "arn:aws:iam::ACCOUNT-B-ID:role/ACCOUNT-B-ROLENAME"
  }
}
```

## Result

The account now has the required permissions. [You can now add it to Cloud Manager.](#)



## Adding AWS credentials to Cloud Manager

After you provide an AWS account with the required permissions, you can add the credentials for that account to Cloud Manager. This enables you to launch Cloud Volumes ONTAP systems in that account.

### Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.



2. Click **Add Credentials** and select **AWS**.
3. Provide AWS keys or the ARN of a trusted IAM role.
4. Confirm that the policy requirements have been met and click **Continue**.
5. Choose the pay-as-you-go subscription that you want to associate with the credentials, or click **Add Subscription** if you don't have one yet.

To create a pay-as-you-go Cloud Volumes ONTAP system, AWS credentials must be associated with a subscription to Cloud Volumes ONTAP from the AWS Marketplace.

6. Click **Add**.

### Result

You can now switch to a different set of credentials from the Details and Credentials page when creating a new working environment:

## Edit Account & Add Subscription

### Credentials

Keys   Account ID: [REDACTED]
<b>Instance Profile   Account ID: [REDACTED]</b>
QA Subscription

### Associate Subscription to Credentials

To create a pay-as-you-go Cloud Volumes ONTAP system, you need to select AWS credentials that are associated with a subscription to Cloud Volumes ONTAP from the AWS Marketplace.

[+ Add Subscription](#)

Apply

Cancel

### Associating an AWS subscription to credentials

After you add your AWS credentials to Cloud Manager, you can associate an AWS Marketplace subscription with those credentials. The subscription enables you to create a pay-as-you-go Cloud Volumes ONTAP system, and to use other NetApp cloud services.

There are two scenarios in which you might associate an AWS Marketplace subscription after you've already added the credentials to Cloud Manager:


- You didn't associate a subscription when you initially added the credentials to Cloud Manager.
- You want to replace an existing AWS Marketplace subscription with a new subscription.

### What you'll need

You need to create a Connector before you can change Cloud Manager settings. [Learn how.](#)

### Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.
2. Hover over a set of credentials and click the action menu.
3. From the menu, click **Associate Subscription**.



4. Select a subscription from the down-down list or click **Add Subscription** and follow the steps to create a new subscription.

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_aws.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_aws.mp4) (video)

## Azure

### Azure credentials and permissions


Cloud Manager enables you to choose the Azure credentials to use when deploying Cloud Volumes ONTAP. You can deploy all of your Cloud Volumes ONTAP systems using the initial Azure credentials, or you can add additional credentials.

#### Initial Azure credentials


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

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

## Cloud Manager



## Azure account



Cloud Manager selects these Azure credentials by default when you create a new working environment for Cloud Volumes ONTAP:

Details & Credentials			
Managed Service Ide...	OCCM QA1	<span style="color: orange;">ⓘ</span> <i>No subscription is associated</i>	<button>Edit Credentials</button>
Credential Name	Azure Subscription	Marketplace Subscription	


### Additional Azure subscriptions for a managed identity

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


### Additional Azure credentials

If you want to deploy Cloud Volumes ONTAP using different Azure credentials, 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:


Initial Azure account



Second account



Third account



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

After you add another set of credentials, you can switch to them when creating a new working environment:

## Edit Account & Add Subscription

### Credentials

cloud-manager-app | Application ID: 57c42424-88a0-480a.

**Managed Service Identity**

OCCM QA1 (Default) ▼

### What about Marketplace deployments and on-prem deployments?

The sections above describe the recommended deployment method for the Connector, which is from NetApp Cloud Central. You can also deploy a Connector in Azure from the [Azure Marketplace](#), and you can [install the Connector 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 the Connector, and then provide permissions for any additional accounts.

For on-premises deployments, you can't set up a managed identity for the Connector, but you can provide permissions just like you would for additional accounts by using a service principal.

### Managing Azure credentials and subscriptions for Cloud Manager

When you create a Cloud Volumes ONTAP system, you need to select the Azure credentials and Marketplace subscription to use with that system. If you manage multiple Azure Marketplace subscriptions, you can assign each one of them to different Azure credentials from the Credentials page.

There are two ways to manage Azure credentials in Cloud Manager. First, if you want to deploy Cloud Volumes ONTAP in different Azure accounts, then you need to provide the required permissions and add the credentials to Cloud Manager. The second way is to associate additional subscriptions with the Azure managed identity.



When you deploy a Connector from Cloud Manager, Cloud Manager automatically adds the Azure account in which you deployed the Connector. An initial account is not added if you manually installed the Connector software on an existing system. [Learn about Azure accounts and permissions](#).


### Granting Azure permissions using a service principal

Cloud Manager needs permissions to perform actions in Azure. You can grant the required permissions to an Azure account by creating and setting up a service principal in Azure Active Directory and by obtaining the

Azure credentials that Cloud Manager needs.

### About this task

The following image depicts how Cloud Manager obtains permissions to perform operations in Azure. A service principal object, which is tied to one or more Azure subscriptions, represents Cloud Manager in Azure Active Directory and is assigned to a custom role that allows the required permissions.



### Steps

1. [Create an Azure Active Directory application.](#)
2. [Assign the application to a role.](#)
3. [Add Windows Azure Service Management API permissions.](#)
4. [Get the application ID and directory ID.](#)
5. [Create a client secret.](#)

### Creating an Azure Active Directory application


Create an Azure Active Directory (AD) application and service principal that Cloud Manager can use for role-based access control.

### Before you begin

You must have the right permissions in Azure to create an Active Directory application and to assign the application to a role. For details, refer to [Microsoft Azure Documentation: Required permissions.](#)

### Steps

1. From the Azure portal, open the **Azure Active Directory** service.



2. In the menu, click **App registrations**.
3. Click **New registration**.
4. Specify details about the application:
  - **Name**: Enter a name for the application.
  - **Account type**: Select an account type (any will work with Cloud Manager).
  - **Redirect URI**: Select **Web** and then enter any URL—for example, `https://url`
5. Click **Register**.

## Result

You've created the AD application and service principal.

## Assigning the application to a role

You must bind the service principal to one or more Azure subscriptions and assign it the custom "OnCommand Cloud Manager Operator" role so Cloud Manager has permissions in Azure.

## Steps

1. Create a custom role:
  - a. Download the [Cloud Manager Azure policy](#).
  - b. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription from which users will create Cloud Volumes ONTAP systems.

## Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

- c. Use the JSON file to create a custom role in Azure.

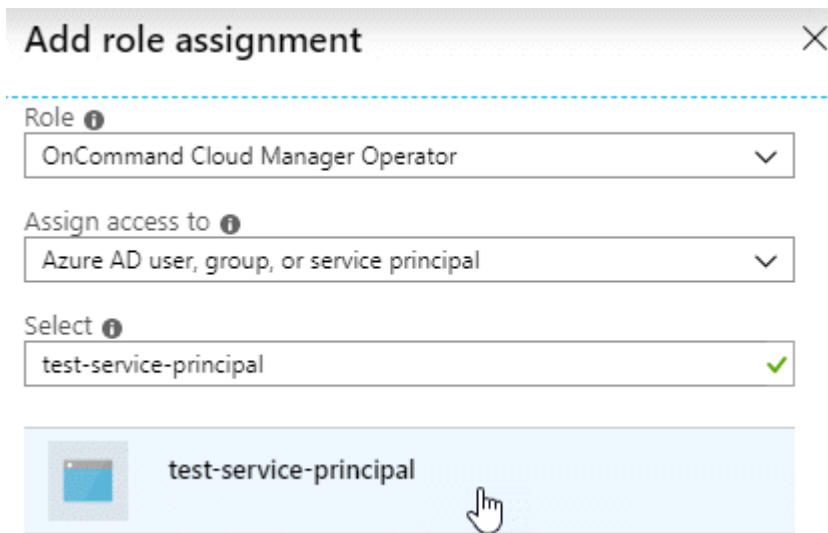
The following example shows how to create a custom role using the Azure CLI 2.0:

```
az role definition create --role-definition
```

C:\Policy\_for\_cloud\_Manager\_Azure\_3.8.7.json

You should now have a custom role called *Cloud Manager Operator*.

2. Assign the application to the role:
  - a. From the Azure portal, open the **Subscriptions** service.
  - b. Select the subscription.
  - c. Click **Access control (IAM) > Add > Add role assignment**.
  - d. Select the **Cloud Manager Operator** role.
  - e. Keep **Azure AD user, group, or service principal** selected.
  - f. Search for the name of the application (you can't find it in the list by scrolling).



The screenshot shows the 'Add role assignment' dialog box. It has a title bar with a close button (X). Below the title bar, there are three dropdown menus. The first is labeled 'Role' and has 'OnCommand Cloud Manager Operator' selected. The second is labeled 'Assign access to' and has 'Azure AD user, group, or service principal' selected. The third is labeled 'Select' and has 'test-service-principal' selected with a green checkmark. Below the dropdowns, there is a list of search results for 'test-service-principal'. The first result is highlighted in blue and has a hand cursor pointing to it.

- g. Select the application and click **Save**.

The service principal for Cloud Manager now has the required Azure permissions for that subscription.

If you want to deploy Cloud Volumes ONTAP from multiple Azure subscriptions, then you must bind the service principal to each of those subscriptions. Cloud Manager enables you to select the subscription that you want to use when deploying Cloud Volumes ONTAP.

## Adding Windows Azure Service Management API permissions

The service principal must have "Windows Azure Service Management API" permissions.

### Steps

1. In the **Azure Active Directory** service, click **App registrations** and select the application.
2. Click **API permissions > Add a permission**.
3. Under **Microsoft APIs**, select **Azure Service Management**.




## Request API permissions

### Select an API

[Microsoft APIs](#) [APIs my organization uses](#) [My APIs](#)


#### Commonly used Microsoft APIs

<b>Microsoft Graph</b> Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10. Access Azure AD, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint. 		
<b>Azure Batch</b> Schedule large-scale parallel and HPC applications in the cloud	<b>Azure Data Catalog</b> Programmatic access to Data Catalog resources to register, annotate and search data assets	<b>Azure Data Explorer</b> Perform ad-hoc queries on terabytes of data to build near real-time and complex analytics solutions
<b>Azure Data Lake</b> Access to storage and compute for big data analytic scenarios	<b>Azure DevOps</b> Integrate with Azure DevOps and Azure DevOps server	<b>Azure Import/Export</b> Programmatic control of import/export jobs
<b>Azure Key Vault</b> Manage your key vaults as well as the keys, secrets, and certificates within your Key Vaults	<b>Azure Rights Management Services</b> Allow validated users to read and write protected content	<b>Azure Service Management</b> Programmatic access to much of the functionality available through the Azure portal
<b>Azure Storage</b> Secure, massively scalable object and data lake storage for unstructured and semi-structured data	<b>Customer Insights</b> Create profile and interaction models for your products	<b>Data Export Service for Microsoft Dynamics 365</b> Export data from Microsoft Dynamics CRM organization to an external destination

4. Click **Access Azure Service Management as organization users** and then click **Add permissions**.

## Request API permissions

[< All APIs](#)

 Azure Service Management  
<https://management.azure.com/> [Docs](#)

What type of permissions does your application require?

### Delegated permissions

Your application needs to access the API as the signed-in user.

### Application permissions

Your application runs as a background service or daemon without a signed-in user.

Select permissions

[expand all](#)

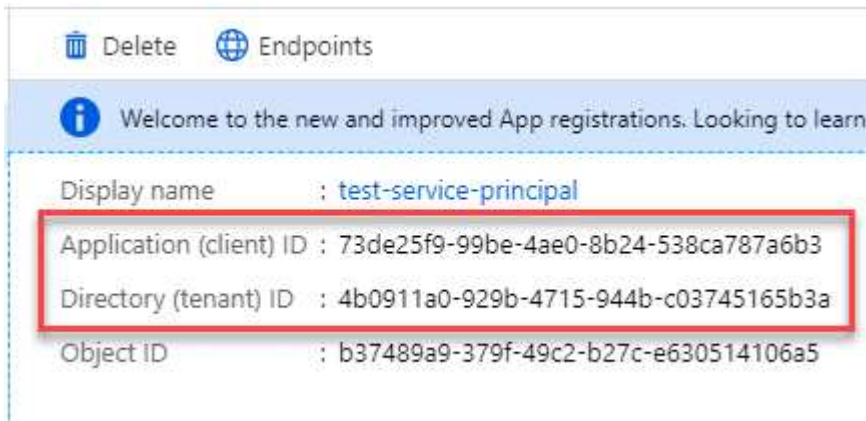
PERMISSION	ADMIN CONSENT REQUIRED
<input checked="" type="checkbox"/> <b>user_impersonation</b> Access Azure Service Management as organization users (preview) 	-

## Getting the application ID and directory ID

When you add the Azure account to Cloud Manager, you need to provide the application (client) ID and the directory (tenant) ID for the application. Cloud Manager uses the IDs to programmatically sign in.

### Steps

1. In the **Azure Active Directory** service, click **App registrations** and select the application.
2. Copy the **Application (client) ID** and the **Directory (tenant) ID**.



The screenshot shows the details of an application registration in Azure Active Directory. At the top, there are 'Delete' and 'Endpoints' icons. Below that is a blue banner with an information icon and the text 'Welcome to the new and improved App registrations. Looking to learn...'. The main content area lists several identifiers: 'Display name : test-service-principal', 'Application (client) ID : 73de25f9-99be-4ae0-8b24-538ca787a6b3', 'Directory (tenant) ID : 4b0911a0-929b-4715-944b-c03745165b3a', and 'Object ID : b37489a9-379f-49c2-b27c-e630514106a5'. The 'Application (client) ID' and 'Directory (tenant) ID' are highlighted with a red rectangular box.

## Creating a client secret

You need to create a client secret and then provide Cloud Manager with the value of the secret so Cloud Manager can use it to authenticate with Azure AD.



When you add the account to Cloud Manager, Cloud Manager refers to the client secret as the Application Key.

### Steps

1. Open the **Azure Active Directory** service.
2. Click **App registrations** and select your application.
3. Click **Certificates & secrets > New client secret**.
4. Provide a description of the secret and a duration.
5. Click **Add**.
6. Copy the value of the client secret.

### Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

DESCRIPTION	EXPIRES	VALUE	
test secret	8/16/2020	*sZ1jSe2By:D*-ZRov4NLfdAcY7:+0vA	Copy to clipboard

### Result

Your service principal is now setup and you should have copied the application (client) ID, the directory (tenant) ID, and the value of the client secret. You need to enter this information in Cloud Manager when you add an Azure account.

### Adding Azure credentials to Cloud Manager

After you provide an Azure account with the required permissions, you can add the credentials for that account to Cloud Manager. This enables you to launch Cloud Volumes ONTAP systems in that account.

### What you'll need

You need to create a Connector before you can change Cloud Manager settings. [Learn how.](#)

### Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.



2. Click **Add Credentials** and select **Microsoft Azure**.
3. Enter information about the Azure Active Directory service principal that grants the required permissions:
  - Application (client) ID: See [Getting the application ID and directory ID](#).
  - Directory (tenant) ID: See [Getting the application ID and directory ID](#).
  - Client Secret: See [Creating a client secret](#).
4. Confirm that the policy requirements have been met and then click **Continue**.
5. Choose the pay-as-you-go subscription that you want to associate with the credentials, or click **Add Subscription** if you don't have one yet.


To create a pay-as-you-go Cloud Volumes ONTAP system, Azure credentials must be associated with a

subscription to Cloud Volumes ONTAP from the Azure Marketplace.

6. Click **Add**.

## Result

You can now switch to different set of credentials from the Details and Credentials page [when creating a new working environment](#):



## Associating an Azure Marketplace subscription to credentials

After you add your Azure credentials to Cloud Manager, you can associate an Azure Marketplace subscription to those credentials. The subscription enables you to create a pay-as-you-go Cloud Volumes ONTAP system, and to use other NetApp cloud services.

There are two scenarios in which you might associate an Azure Marketplace subscription after you've already added the credentials to Cloud Manager:


- You didn't associate a subscription when you initially added the credentials to Cloud Manager.
- You want to replace an existing Azure Marketplace subscription with a new subscription.

## What you'll need

You need to create a Connector before you can change Cloud Manager settings. [Learn how](#).

## Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.
2. Hover over a set of credentials and click the action menu.
3. From the menu, click **Associate Subscription**.



4. Select a subscription from the down-down list or click **Add Subscription** and follow the steps to create a new subscription.

The following video starts from the context of the working environment wizard, but shows you the same workflow after you click **Add Subscription**:

► [https://docs.netapp.com/us-en/occm38//media/video\\_subscribing\\_azure.mp4](https://docs.netapp.com/us-en/occm38//media/video_subscribing_azure.mp4) (video)

#### Associating additional Azure subscriptions with a managed identity

Cloud Manager enables you to choose the Azure credentials and Azure subscription in which you want to deploy Cloud Volumes ONTAP. You can't select a different Azure subscription for the managed identity profile unless you associate the [managed identity](#) with those subscriptions.

#### About this task

A managed identity is [the initial Azure account](#) when you deploy a Connector from Cloud Manager. When you deployed the Connector, Cloud Manager created the Cloud Manager Operator role and assigned it to the Connector virtual machine.

#### Steps

1. Log in to the Azure portal.
2. Open the **Subscriptions** service and then select the subscription in which you want to deploy Cloud Volumes ONTAP.
3. Click **Access control (IAM)**.

- a. Click **Add > Add role assignment** and then add the permissions:

- Select the **Cloud Manager Operator** role.




Cloud Manager Operator is the default name provided in the [Cloud Manager policy](#). If you chose a different name for the role, then select that name instead.

- Assign access to a **Virtual Machine**.
- Select the subscription in which the Connector virtual machine was created.
- Select the Connector virtual machine.
- Click **Save**.

4. Repeat these steps for additional subscriptions.

#### Result

When you create a new working environment, you should now have the ability to select from multiple Azure subscriptions for the managed identity profile.



## GCP

### 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.

#### Project and permissions for Cloud Manager


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

Two sets of permissions must be in place before you deploy a Connector directly from Cloud Manager:

1. You need to deploy a Connector using a Google account that has permissions to launch the Connector VM instance from Cloud Manager.
2. When deploying the Connector, 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 the Connector, or in a different project. To deploy Cloud Volumes ONTAP in a different project, you need to first add the Connector service account and role to that project.

- [Learn how to set up service account \(see step 2\).](#)
- [Learn how to deploy Cloud Volumes ONTAP in GCP and select a project.](#)

### Account for data tiering



Cloud Manager requires a GCP account for Cloud Volumes ONTAP 9.6, but not for 9.7 and later. If you want to use data tiering with Cloud Volumes ONTAP 9.7, then follow step 4 in [Getting started with Cloud Volumes ONTAP in Google Cloud Platform](#).

Adding a Google Cloud account to Cloud Manager is required to enable data tiering on a Cloud Volumes ONTAP 9.6 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.](#)

### Managing GCP credentials and subscriptions for Cloud Manager

You can manage two types of Google Cloud Platform credentials from Cloud Manager: the credentials that are associated with the Connector VM instance and storage access keys used with a Cloud Volumes ONTAP 9.6 system for [data tiering](#).



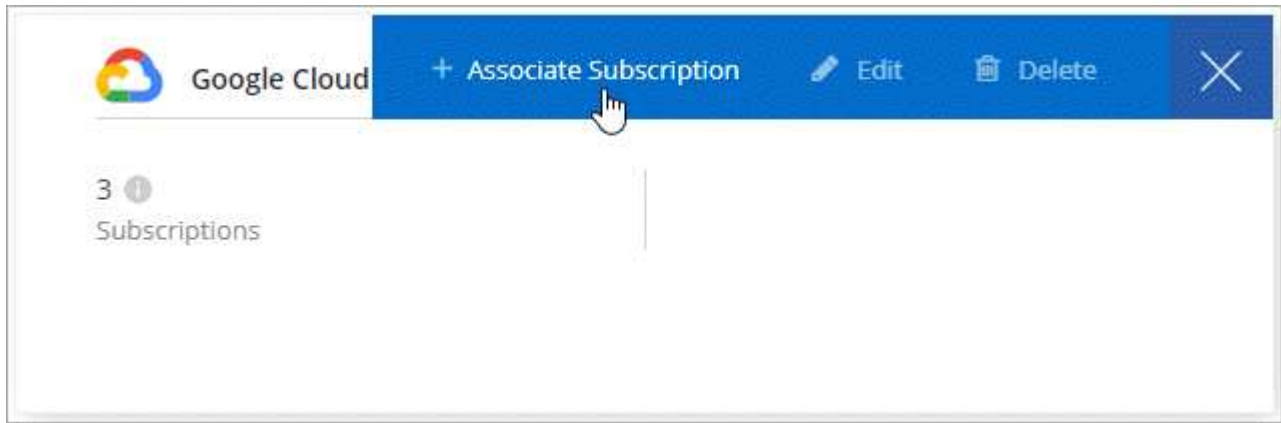
## Associating a Marketplace subscription with GCP credentials

When you deploy a Connector in GCP, Cloud Manager creates a default set of credentials that are associated with the Connector VM instance. These are the credentials that Cloud Manager uses to deploy Cloud Volumes ONTAP.

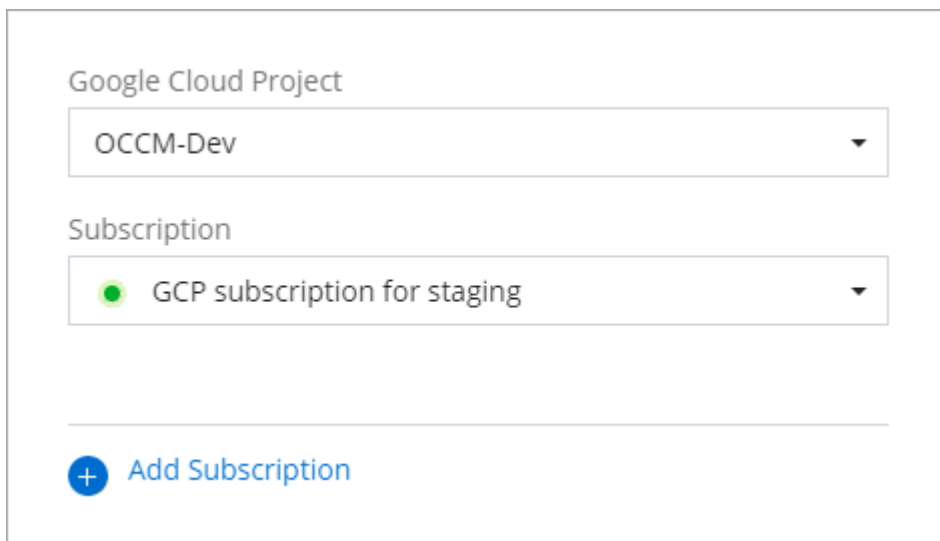
At any time, you can change the Marketplace subscription that's associated with these credentials. The subscription enables you to create a pay-as-you-go Cloud Volumes ONTAP system, and to use other NetApp cloud services.

### Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.
2. Hover over a set of credentials and click the action menu.
3. From the menu, click **Associate Subscription**.



4. Select a Google Cloud project and subscription from the down-down list or click **Add Subscription** and follow the steps to create a new subscription.

A screenshot of a form in the Google Cloud console. It has two dropdown menus. The first is labeled 'Google Cloud Project' and has 'OCCM-Dev' selected. The second is labeled 'Subscription' and has 'GCP subscription for staging' selected, which is preceded by a green dot. Below the dropdowns is a blue button with a plus sign and the text 'Add Subscription'.

5. Click **Associate**.

## Setting up and adding GCP accounts for data tiering with Cloud Volumes ONTAP 9.6

If you want to enable a Cloud Volumes ONTAP 9.6 system for [data tiering](#), 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.



If you want to use data tiering with Cloud Volumes ONTAP 9.7, then follow step 4 in [Getting started with Cloud Volumes ONTAP in Google Cloud Platform](#).

## Setting up a service account and access keys for Google Cloud Storage

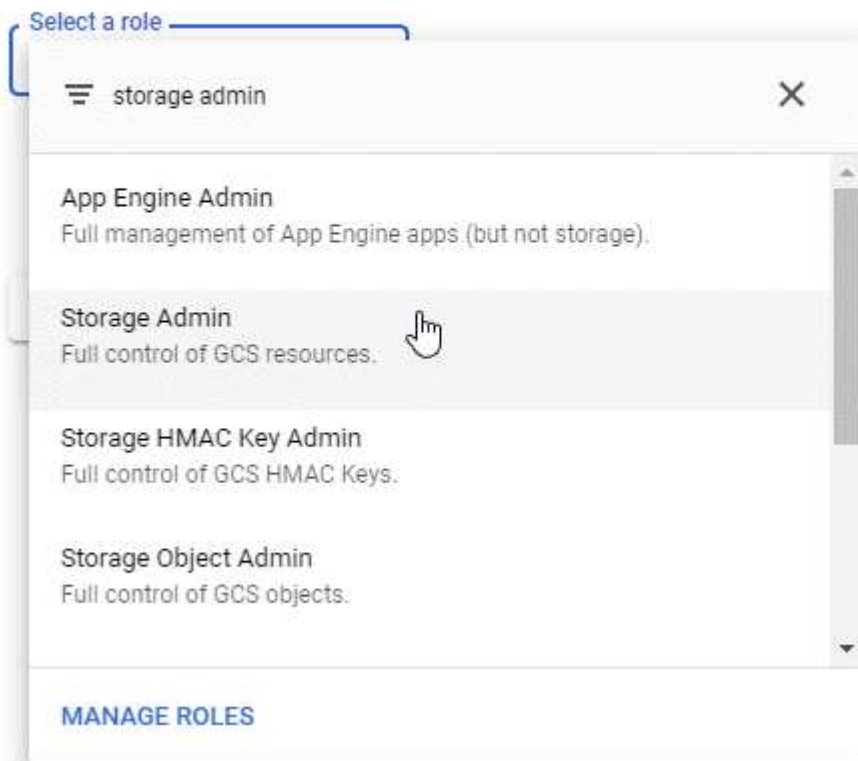
A service account enables Cloud Manager to authenticate and access Cloud Storage buckets used for data tiering. The keys are required so that Google Cloud Storage knows who is making the request.

### Steps

1. Open the GCP IAM console and [create a service account that has the Storage Admin role](#).

### Service account permissions (optional)

Grant this service account access to My Project 99247 so that it has permission to complete specific actions on the resources in your project. [Learn more](#)



2. Go to [GCP Storage Settings](#).
3. If you're prompted, select a project.
4. Click the **Interoperability** tab.
5. If you haven't already done so, click **Enable interoperability access**.
6. Under **Access keys for service accounts**, click **Create a key for a service account**.
7. Select the service account that you created in step 1.

## Select a service account

Email	Name	Keys
<input checked="" type="radio"/> data-tiering-for-netapp@top-monitor-250116.iam.gserviceaccount.com	data tiering for netapp	—

[CANCEL](#) [CREATE KEY](#) | [CREATE NEW ACCOUNT](#)

8. Click **Create Key**.
9. Copy the access key and secret.

You'll need to enter this information in Cloud Manager when you add the GCP account for data tiering.

### Adding a GCP account to Cloud Manager

Now that you have an access key for a service account, you can add it to Cloud Manager.

#### What you'll need

You need to create a Connector before you can change Cloud Manager settings. [Learn how](#).

#### Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.



2. Click **Add Credentials** and select **Google Cloud**.
3. Enter the access key and secret for the service account.

The keys enable Cloud Manager to set up a Cloud Storage bucket for data tiering.

4. Confirm that the policy requirements have been met and then click **Create Account**.

#### What's next?

You can now enable data tiering on individual volumes on a Cloud Volumes ONTAP 9.6 system when you create, modify, or replicate them. For details, see [Tiering inactive data to low-cost object storage](#).

But before you do, be sure that the subnet in which Cloud Volumes ONTAP resides is configured for Private Google Access. For instructions, refer to [Google Cloud Documentation: Configuring Private Google Access](#).

### Adding NetApp Support Site accounts to Cloud Manager

Adding your NetApp Support Site account to Cloud Manager is required to deploy a

BYOL system. It's also required to register pay-as-you-go systems and to upgrade ONTAP software.

Watch the following video to learn how to add NetApp Support Site accounts to Cloud Manager. Or scroll down to read the steps.



### What you'll need

You need to create a Connector before you can change Cloud Manager settings. [Learn how.](#)

### Steps

1. If you don't have a NetApp Support Site account yet, [register for one](#).
2. In the upper right of the Cloud Manager console, click the Settings icon, and select **Credentials**.



3. Click **Add Credentials** and select **NetApp Support Site**.
4. Specify a name for the account and then enter the user name and password.
  - The account must be a customer-level account (not a guest or temp account).
  - If you plan to deploy BYOL systems:
    - The account must be authorized to access the serial numbers of the BYOL systems.
    - If you purchased a secure BYOL subscription, then a secure NSS account is required.
5. Click **Create Account**.

### What's next?

Users can now select the account when creating new Cloud Volumes ONTAP systems and when registering existing systems.

- [Launching Cloud Volumes ONTAP in AWS](#)
- [Launching Cloud Volumes ONTAP in Azure](#)
- [Registering pay-as-you-go systems](#)
- [Learn how Cloud Manager manages license files](#)

## Managing users, workspaces, Connectors, and subscriptions

After you perform initial setup, you might need to administer your account settings later by managing users, workspaces, Connectors, and subscriptions.


[Learn more about how Cloud Central accounts work.](#)

### Adding users


Associate Cloud Central users with the Cloud Central account so those users can create and manage working environments in Cloud Manager.

#### Steps


1. If the user hasn't already done so, ask the user to go to [NetApp Cloud Central](#) and sign up.
2. From the top of Cloud Manager, click the **Account** drop-down.



3. Click **Manage Account** next to the currently selected account.



4. From the Users tab, click **Associate User**.
5. Enter the user's email address and select a role for the user:
  - **Account Admin**: Can perform any action in Cloud Manager.
  - **Workspace Admin**: Can create and manage resources in assigned workspaces.
  - **Compliance Viewer**: Can only view compliance information and generate reports for workspaces that they have permission to access.
6. If you selected Workspace Admin or Compliance Viewer, select one or more workspaces to associate with that user.



## Associate User

To add a user to your NetApp Cloud Account, that user must already have signed up at [NetApp Cloud Central](#). Enter the email address that they used when signing up with Cloud Central.

User's Email

Role

Associate User to Workspaces

7. Click **Associate User**.

### Result


The user should receive an email from NetApp Cloud Central titled "Account Association." The email includes the information needed to access Cloud Manager.

## Removing users


Disassociating a user makes it so they can no longer access the resources in a Cloud Central account.

### Steps

1. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.



2. From the Users tab, click the action menu in the row that corresponds to the user.



3. Click **Disassociate User** and click **Disassociate** to confirm.

### Result

The user can no longer access the resources in this Cloud Central account.


## Managing a Workspace Admin's workspaces

You can associate and disassociate Workspace Admins with workspaces at any time. Associating the user enables them to create and view the working environments in that workspace.

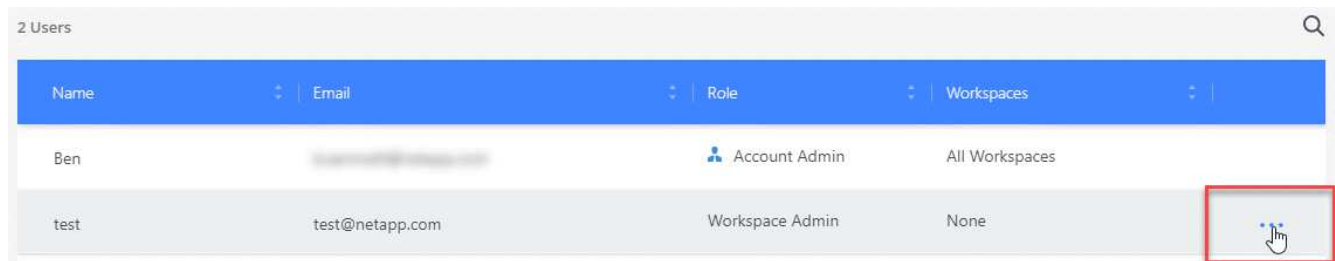
### Steps

1. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.





2. From the Users tab, click the action menu in the row that corresponds to the user.



3. Click **Manage Workspaces**.

4. Select the workspaces to associate with the user and click **Apply**.

## Result

The user can now access those workspaces from Cloud Manager, as long as the Connector was also associated with the workspaces.

## Managing workspaces

Manage your workspaces by creating, renaming, and deleting them. Note that you can't delete a workspace if it contains any resources. It must be empty.

### Steps

1. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.
2. Click **Workspaces**.
3. Choose one of the following options:
  - Click **Add New Workspace** to create a new workspace.
  - Click **Rename** to rename the workspace.
  - Click **Delete** to delete the workspace.



## Managing a Connector's workspaces

You need to associate the Connector with workspaces so Workspace Admins can access those workspaces from Cloud Manager.

If you only have Account Admins, then associating the Connector with workspaces isn't required. Account Admins have the ability to access all workspaces in Cloud Manager by default.

[Learn more about users, workspaces, and Connectors.](#)

### Steps

1. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.
2. Click **Connector**.
3. Click **Manage Workspaces** for the Connector that you want to associate.
4. Select the workspaces to associate with the Connector and click **Apply**.

## Managing subscriptions

After you subscribe from a cloud provider's marketplace, each subscription is available from the Account Settings widget. You have the option to rename a subscription and to disassociate the subscription from one or more accounts.

For example, let's say that you have two accounts and each is billed through separate subscriptions. You might disassociate a subscription from one of the accounts so the users in that account don't accidentally choose the wrong subscription when creating a Cloud Volume ONTAP working environment.

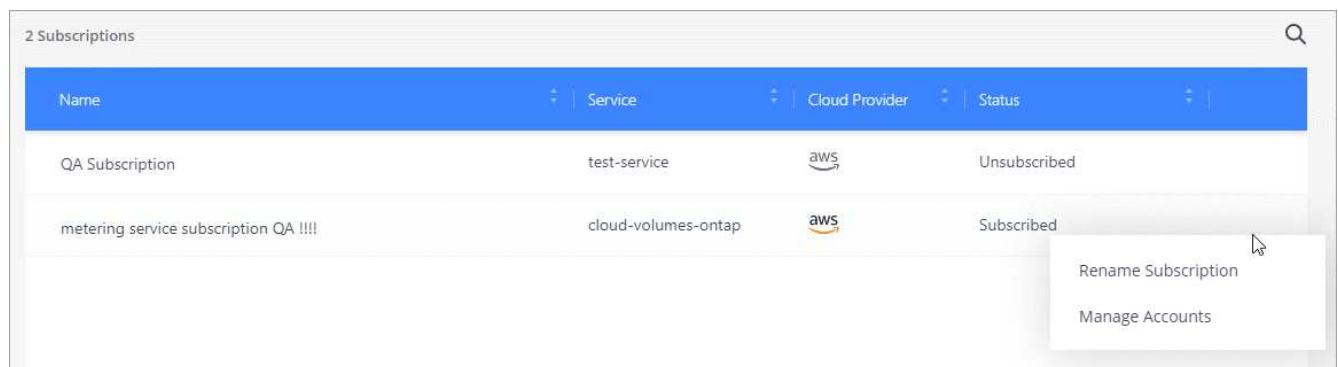
[Learn more about subscriptions.](#)

### Steps

1. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.
2. Click **Subscriptions**.

You'll only see the subscriptions that are associated with the account that you're currently viewing.

3. Click the action menu in the row that corresponds to the subscription that you want to manage.



4. Choose to rename the subscription or to manage the accounts that are associated with the subscription.

## Changing the account name

Change your account name at any time to change it to something meaningful for you.

### Steps

1. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.
2. In the **Overview** tab, click the edit icon next to the account name.
3. Type a new account name and click **Save**.

## Enabling or disabling the SaaS platform

We don't recommend disabling the SaaS platform unless you need to in order to comply with your company's security policies. Disabling the SaaS platform limits your ability to use NetApp's integrated cloud services.

The following services aren't available from Cloud Manager if you disable the SaaS platform:

- Cloud Compliance
- Kubernetes
- Cloud Tiering
- Global File Cache
- Monitoring (Cloud Insights)

### Steps

1. From the top of Cloud Manager, click the **Account** drop-down and click **Manage Account**.
2. In the **Overview** tab, toggle the option to enable use the SaaS platform.

## Managing an HTTPS certificate for secure access

By default, Cloud Manager uses a self-signed certificate for HTTPS access to the web console. You can install a certificate signed by a certificate authority (CA), which provides better security protection than a self-signed certificate.

### Before you get started

You need to create a Connector before you can change Cloud Manager settings. [Learn how](#).

### Installing an HTTPS certificate

Install a certificate signed by a CA for secure access.

### Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **HTTPS Setup**.



- In the HTTPS Setup page, install a certificate by generating a certificate signing request (CSR) or by installing your own CA-signed certificate:

Option	Description
Generate a CSR	<ol style="list-style-type: none"> <li>Enter the host name or DNS of the Connector host (its Common Name), and then click <b>Generate CSR</b>.  Cloud Manager displays a certificate signing request.</li> <li>Use the CSR to submit an SSL certificate request to a CA.  The certificate must use the Privacy Enhanced Mail (PEM) Base-64 encoded X.509 format.</li> <li>Copy the contents of the signed certificate, paste it in the Certificate field, and then click <b>Install</b>.</li> </ol>
Install your own CA-signed certificate	<ol style="list-style-type: none"> <li>Select <b>Install CA-signed certificate</b>.</li> <li>Load both the certificate file and the private key and then click <b>Install</b>.  The certificate must use the Privacy Enhanced Mail (PEM) Base-64 encoded X.509 format.</li> </ol>

### Result

Cloud Manager now uses the CA-signed certificate to provide secure HTTPS access. The following image shows a Cloud Manager system that is configured for secure access:

#### Cloud Manager HTTPS certificate

Expiration:

 Oct 27, 2016 05:13:28 am

Issuer:

CN=localhost, O=NetApp, OU=Tel-Aviv,  
EMAILADDRESS=admin@example.com

Subject:

EMAILADDRESS= admin@example.com ,  
OU=Tel-Aviv, O=NetApp, CN=localhost

 [View Certificate](#)

 [Renew HTTPS Certificate](#)

### Renewing the Cloud Manager HTTPS certificate

You should renew the Cloud Manager HTTPS certificate before it expires to ensure secure access to the Cloud Manager web console. If you do not renew the certificate before it expires, a warning appears when users access the web console using HTTPS.

## Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **HTTPS Setup**.

Details about the Cloud Manager certificate displays, including the expiration date.

2. Click **Renew HTTPS Certificate** and follow the steps to generate a CSR or install your own CA-signed certificate.

## Result

Cloud Manager uses the new CA-signed certificate to provide secure HTTPS access.

# Removing Cloud Volumes ONTAP working environments

The Account Admin can remove a Cloud Volumes ONTAP working environment to move it to another system or to troubleshoot discovery issues.

## About this task

Removing a Cloud Volumes ONTAP working environment removes it from Cloud Manager. It does not delete the Cloud Volumes ONTAP system. You can later rediscover the working environment.

Removing a working environment from Cloud Manager enables you to do the following:

- Rediscover it in another workspace
- Rediscover it from another Cloud Manager system
- Rediscover it if you had problems during the initial discovery

## Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Tools**.



2. From the Tools page, click **Launch**.
3. Select the Cloud Volumes ONTAP working environment that you want to remove.
4. On the Review and Approve page, click **Go**.

## Result

Cloud Manager removes the working environment. Users can rediscover this working environment from the Working Environments page at any time.

# Configuring a Connector to use a proxy server

If your corporate policies dictate that you use a proxy server for all HTTP communication to the internet, then you must configure your Connectors to use that proxy server. The proxy server can be in the cloud or in your network.

When you configure a Connector to use a proxy server, that Connector and the Cloud Volumes ONTAP

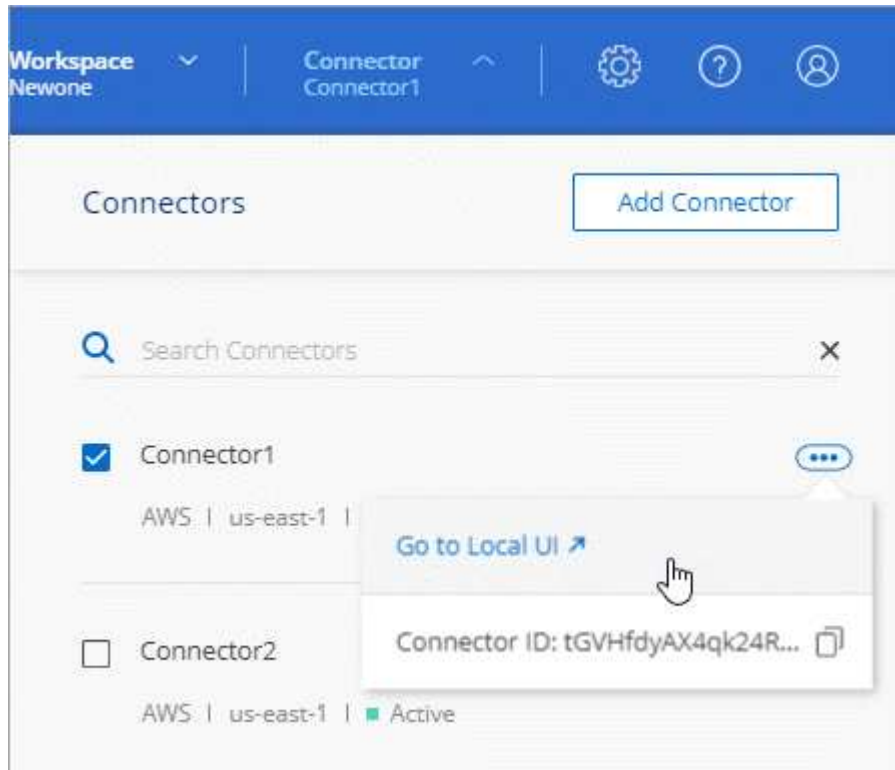
systems that it manages (including any HA mediators), all use the proxy server.

## Steps

1. [Log in to the Cloud Manager SaaS interface](#) from a machine that has a network connection to the Connector instance.

If the Connector doesn't have a public IP address, you'll need a VPN connection or you'll need to connect from a jump host that's in the same network as the Connector.

2. Click the **Connector** drop-down and then click **Go to local UI** for a specific Connector.



The Cloud Manager interface running on the Connector loads in a new browser tab.

3. In the upper right of the Cloud Manager console, click the Settings icon, and select **Cloud Manager Settings**.



4. Under HTTP Proxy, enter the server using the syntax `http://address:port`, specify a user name and password if basic authentication is required for the server, and then click **Save**.



Cloud Manager doesn't support passwords that include the @ character.

## Result

After you specify the proxy server, new Cloud Volumes ONTAP systems are automatically configured to use the proxy server when sending AutoSupport messages. If you didn't specify the proxy server before users create Cloud Volumes ONTAP systems, then they must use System Manager to manually set the proxy server

in the AutoSupport options for each system.

## Overriding CIFS locks for Cloud Volumes ONTAP HA in Azure

The Account Admin can enable a setting in Cloud Manager that prevents issues with Cloud Volumes ONTAP storage failover during Azure maintenance events. When you enable this setting, Cloud Volumes ONTAP vetoes CIFS locks and resets active CIFS sessions.

### About this task

Microsoft Azure schedules periodic maintenance events on its virtual machines. When a maintenance event occurs on a node in a Cloud Volumes ONTAP HA pair, the HA pair initiates storage takeover. If there are active CIFS sessions during this maintenance event, the locks on CIFS files can prevent storage failover.

If you enable this setting, Cloud Volumes ONTAP will veto the locks and reset the active CIFS sessions. As a result, the HA pair can complete storage failover during these maintenance events.



This process might be disruptive to CIFS clients. Data that is not committed from CIFS clients could be lost.

### What you'll need

You need to create a Connector before you can change Cloud Manager settings. [Learn how.](#)

### Steps

1. In the upper right of the Cloud Manager console, click the Settings icon, and select **Cloud Manager Settings**.



2. Under **HA CIFS Locks**, select the checkbox and click **Save**.

## Reference

### Roles

The Account Admin, Workspace Admin, and Cloud Compliance Viewer roles provide specific permissions to users.

Task	Account Admin	Workspace Admin	Cloud Compliance Viewer
Manage working environments	Yes	Yes	No
Enable services on working environments	Yes	Yes	No

Task	Account Admin	Workspace Admin	Cloud Compliance Viewer
View data replication status	Yes	Yes	No
View the timeline	Yes	Yes	No
Switch between workspaces	Yes	Yes	Yes
View Compliance scan results	Yes	Yes	Yes
Delete working environments	Yes	No	No
Connect Kubernetes clusters to working environments	Yes	No	No
Receive the Cloud Volumes ONTAP report	Yes	No	No
Create Connectors	Yes	No	No
Manage Cloud Central accounts	Yes	No	No
Manage credentials	Yes	No	No
Modify Cloud Manager settings	Yes	No	No
View and manage the Support Dashboard	Yes	No	No
Remove working environments from Cloud Manager	Yes	No	No
Install an HTTPS certificate	Yes	No	No

#### Related links

- [Setting up workspaces and users in the Cloud Central account](#)
- [Managing workspaces and users in the Cloud Central account](#)

## How Cloud Manager uses cloud provider permissions

Cloud Manager requires permissions to perform actions in your cloud provider. These permissions are included in [the policies provided by NetApp](#). You might want to understand what Cloud Manager does with these permissions.

### What Cloud Manager does with AWS permissions

Cloud Manager uses an AWS account to make API calls to several AWS services, including EC2, S3, CloudFormation, IAM, the Security Token Service (STS), and the Key Management Service (KMS).

Actions	Purpose
"ec2:StartInstances", "ec2:StopInstances", "ec2:DescribeInstances", "ec2:DescribeInstanceStatus", "ec2:RunInstances", "ec2:TerminateInstances", "ec2:ModifyInstanceAttribute",	Launches a Cloud Volumes ONTAP instance and stops, starts, and monitors the instance.
"ec2:DescribeInstanceAttribute",	Verifies that enhanced networking is enabled for supported instance types.
"ec2:DescribeRouteTables", "ec2:DescribeImages",	Launches a Cloud Volumes ONTAP HA configuration.
"ec2:CreateTags",	Tags every resource that Cloud Manager creates with the "WorkingEnvironment" and "WorkingEnvironmentId" tags. Cloud Manager uses these tags for maintenance and cost allocation.
"ec2:CreateVolume", "ec2:DescribeVolumes", "ec2:ModifyVolumeAttribute", "ec2:AttachVolume", "ec2>DeleteVolume", "ec2:DetachVolume",	Manages the EBS volumes that Cloud Volumes ONTAP uses as back-end storage.
"ec2:CreateSecurityGroup", "ec2>DeleteSecurityGroup", "ec2:DescribeSecurityGroups", "ec2:RevokeSecurityGroupEgress", "ec2:AuthorizeSecurityGroupEgress", "ec2:AuthorizeSecurityGroupIngress", "ec2:RevokeSecurityGroupIngress",	Creates predefined security groups for Cloud Volumes ONTAP.
"ec2:CreateNetworkInterface", "ec2:DescribeNetworkInterfaces", "ec2>DeleteNetworkInterface", "ec2:ModifyNetworkInterfaceAttribute",	Creates and manages network interfaces for Cloud Volumes ONTAP in the target subnet.
"ec2:DescribeSubnets", "ec2:DescribeVpcs",	Gets the list of destination subnets and security groups, which is needed when creating a new working environment for Cloud Volumes ONTAP.
"ec2:DescribeDhcpOptions",	Determines DNS servers and the default domain name when launching Cloud Volumes ONTAP instances.
"ec2:CreateSnapshot", "ec2>DeleteSnapshot", "ec2:DescribeSnapshots",	Takes snapshots of EBS volumes during initial setup and whenever a Cloud Volumes ONTAP instance is stopped.
"ec2:GetConsoleOutput",	Captures the Cloud Volumes ONTAP console, which is attached to AutoSupport messages.
"ec2:DescribeKeyPairs",	Obtains the list of available key pairs when launching instances.



Actions	Purpose
"ec2:DescribeRegions",	Gets a list of available AWS regions.
"ec2:DeleteTags", "ec2:DescribeTags",	Manages tags for resources associated with Cloud Volumes ONTAP instances.
"cloudformation:CreateStack", "cloudformation>DeleteStack", "cloudformation:DescribeStacks", "cloudformation:DescribeStackEvents", "cloudformation:ValidateTemplate",	Launches Cloud Volumes ONTAP instances.
"iam:PassRole", "iam:CreateRole", "iam>DeleteRole", "iam:PutRolePolicy", "iam:CreateInstanceProfile", "iam>DeleteRolePolicy", "iam:AddRoleToInstanceProfile", "iam:RemoveRoleFromInstanceProfile", "iam:DeleteInstanceProfile",	Launches a Cloud Volumes ONTAP HA configuration.
"iam:ListInstanceProfiles", "sts:DecodeAuthorizationMessage", "ec2:AssociateIamInstanceProfile", "ec2:DescribeIamInstanceProfileAssociations", "ec2:DisassociateIamInstanceProfile",	Manages instance profiles for Cloud Volumes ONTAP instances.
"s3:GetBucketTagging", "s3:GetBucketLocation", "s3:ListAllMyBuckets", "s3:ListBucket"	Obtains information about AWS S3 buckets so Cloud Manager can integrate with the NetApp Data Fabric Cloud Sync service.
"s3:CreateBucket", "s3>DeleteBucket", "s3:GetLifecycleConfiguration", "s3:PutLifecycleConfiguration", "s3:PutBucketTagging", "s3:ListBucketVersions", "s3:GetBucketPolicyStatus", "s3:GetBucketPublicAccessBlock", "s3:GetBucketAcl", "s3:GetBucketPolicy", "s3:PutBucketPublicAccessBlock"	Manages the S3 bucket that a Cloud Volumes ONTAP system uses as a capacity tier for data tiering.
"kms:List*", "kms:ReEncrypt*", "kms:Describe*", "kms:CreateGrant",	Enables data encryption of Cloud Volumes ONTAP using the AWS Key Management Service (KMS).
"ce:GetReservationUtilization", "ce:GetDimensionValues", "ce:GetCostAndUsage", "ce:GetTags"	Obtains AWS cost data for Cloud Volumes ONTAP.

Actions	Purpose
"ec2:CreatePlacementGroup", "ec2>DeletePlacementGroup"	When you deploy an HA configuration in a single AWS Availability Zone, Cloud Manager launches the two HA nodes and the mediator in an AWS spread placement group.
"ec2:DescribeReservedInstancesOfferings"	Cloud Manager uses the permission as part of Cloud Compliance deployment to choose which instance type to use.
"s3:DeleteBucket", "s3:GetLifecycleConfiguration", "s3:PutLifecycleConfiguration", "s3:PutBucketTagging", "s3:ListBucketVersions", "s3:GetObject", "s3:ListBucket", "s3:ListAllMyBuckets", "s3:GetBucketTagging", "s3:GetBucketLocation", "s3:GetBucketPolicyStatus", "s3:GetBucketPublicAccessBlock", "s3:GetBucketAcl", "s3:GetBucketPolicy", "s3:PutBucketPublicAccessBlock"	Cloud Manager uses these permissions when you enable the Backup to S3 service.

### What Cloud Manager does with Azure permissions

The Cloud Manager Azure policy includes the permissions that Cloud Manager needs to deploy and manage Cloud Volumes ONTAP in Azure.

Actions	Purpose
"Microsoft.Compute/locations/operations/read", "Microsoft.Compute/locations/vmSizes/read", "Microsoft.Compute/operations/read", "Microsoft.Compute/virtualMachines/instanceView/read", "Microsoft.Compute/virtualMachines/powerOff/action", "Microsoft.Compute/virtualMachines/read", "Microsoft.Compute/virtualMachines/restart/action", "Microsoft.Compute/virtualMachines/start/action", "Microsoft.Compute/virtualMachines/deallocate/action", "Microsoft.Compute/virtualMachines/vmSizes/read", "Microsoft.Compute/virtualMachines/write",	Creates Cloud Volumes ONTAP and stops, starts, deletes, and obtains the status of the system.
"Microsoft.Compute/images/write", "Microsoft.Compute/images/read",	Enables Cloud Volumes ONTAP deployment from a VHD.

Actions	Purpose
"Microsoft.Compute/disks/delete", "Microsoft.Compute/disks/read", "Microsoft.Compute/disks/write", "Microsoft.Storage/checknameavailability/read", "Microsoft.Storage/operations/read", "Microsoft.Storage/storageAccounts/listkeys/action", "Microsoft.Storage/storageAccounts/read", "Microsoft.Storage/storageAccounts/regeneratekey/action", "Microsoft.Storage/storageAccounts/write", "Microsoft.Storage/storageAccounts/delete", "Microsoft.Storage/usages/read",	Manages Azure storage accounts and disks, and attaches the disks to Cloud Volumes ONTAP.
"Microsoft.Network/networkInterfaces/read", "Microsoft.Network/networkInterfaces/write", "Microsoft.Network/networkInterfaces/join/action",	Creates and manages network interfaces for Cloud Volumes ONTAP in the target subnet.
"Microsoft.Network/networkSecurityGroups/read", "Microsoft.Network/networkSecurityGroups/write", "Microsoft.Network/networkSecurityGroups/join/action",	Creates predefined network security groups for Cloud Volumes ONTAP.
"Microsoft.Resources/subscriptions/locations/read", "Microsoft.Network/locations/operationResults/read", "Microsoft.Network/locations/operations/read", "Microsoft.Network/virtualNetworks/read", "Microsoft.Network/virtualNetworks/checkIpAvailability/read", "Microsoft.Network/virtualNetworks/subnets/read", "Microsoft.Network/virtualNetworks/subnets/virtualMachines/read", "Microsoft.Network/virtualNetworks/virtualMachines/read", "Microsoft.Network/virtualNetworks/subnets/join/action",	Gets network information about regions, the target VNet and subnet, and adds Cloud Volumes ONTAP to VNets.
"Microsoft.Network/virtualNetworks/subnets/write", "Microsoft.Network/routeTables/join/action",	Enables VNet service endpoints for data tiering.
"Microsoft.Resources/deployments/operations/read", "Microsoft.Resources/deployments/read", "Microsoft.Resources/deployments/write",	Deploys Cloud Volumes ONTAP from a template.

Actions	Purpose
"Microsoft.Resources/deployments/operations/read", "Microsoft.Resources/deployments/read", "Microsoft.Resources/deployments/write", "Microsoft.Resources/resources/read", "Microsoft.Resources/subscriptions/operationresults/read", "Microsoft.Resources/subscriptions/resourceGroups/delete", "Microsoft.Resources/subscriptions/resourceGroups/read", "Microsoft.Resources/subscriptions/resourcegroups/resources/read", "Microsoft.Resources/subscriptions/resourceGroups/write",	Creates and manages resource groups for Cloud Volumes ONTAP.
"Microsoft.Compute/snapshots/write", "Microsoft.Compute/snapshots/read", "Microsoft.Compute/disks/beginGetAccess/action"	Creates and manages Azure managed snapshots.
"Microsoft.Compute/availabilitySets/write", "Microsoft.Compute/availabilitySets/read",	Creates and manages availability sets for Cloud Volumes ONTAP.
"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/read", "Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/write"	Enables programmatic deployments from the Azure Marketplace.
"Microsoft.Network/loadBalancers/read", "Microsoft.Network/loadBalancers/write", "Microsoft.Network/loadBalancers/delete", "Microsoft.Network/loadBalancers/backendAddressPools/read", "Microsoft.Network/loadBalancers/backendAddressPools/join/action", "Microsoft.Network/loadBalancers/frontendIPConfigurations/read", "Microsoft.Network/loadBalancers/loadBalancingRules/read", "Microsoft.Network/loadBalancers/probes/read", "Microsoft.Network/loadBalancers/probes/join/action",	Manages an Azure load balancer for HA pairs.
"Microsoft.Authorization/locks/*"	Enables management of locks on Azure disks.
"Microsoft.Authorization/roleDefinitions/write", "Microsoft.Authorization/roleAssignments/write", "Microsoft.Web/sites/*"	Manages failover for HA pairs.

Actions	Purpose
"Microsoft.Network/privateEndpoints/write", "Microsoft.Storage/storageAccounts/PrivateEndpointConnectionsApproval/action", "Microsoft.Storage/storageAccounts/privateEndpointConnections/read", "Microsoft.Network/privateEndpoints/read", "Microsoft.Network/privateDnsZones/write", "Microsoft.Network/privateDnsZones/virtualNetworkLinks/write", "Microsoft.Network/virtualNetworks/join/action", "Microsoft.Network/privateDnsZones/A/write", "Microsoft.Network/privateDnsZones/read", "Microsoft.Network/privateDnsZones/virtualNetworkLinks/read",	Enables the management of private endpoints. Private endpoints are used when connectivity isn't provided to outside the subnet. Cloud Manager creates the storage account for HA with only internal connectivity within the subnet.
"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/delete",	Enables Cloud Manager to delete volumes for Azure NetApp Files.
"Microsoft.Resources/deployments/operationStatuses/read"	Azure requires this permission for some virtual machine deployments (it depends on the underlying physical hardware that's used during deployment).
"Microsoft.Resources/deployments/operationStatuses/read", "Microsoft.Insights/Metrics/Read", "Microsoft.Compute/virtualMachines/extensions/write", "Microsoft.Compute/virtualMachines/extensions/read", "Microsoft.Compute/virtualMachines/extensions/delete", "Microsoft.Compute/virtualMachines/delete", "Microsoft.Network/networkInterfaces/delete", "Microsoft.Network/networkSecurityGroups/delete", "Microsoft.Resources/deployments/delete",	Enables you to use Global File Cache.
"Microsoft.Compute/diskEncryptionSets/read"	Enables Cloud Manager to encrypt Azure managed disks on single node Cloud Volumes ONTAP systems using external keys from another account. This feature is supported using APIs.

### What Cloud Manager does with GCP permissions

The Cloud Manager policy for GCP includes the permissions that Cloud Manager needs to deploy and manage Cloud Volumes ONTAP.

Actions	Purpose
- compute.disks.create - compute.disks.createSnapshot - compute.disks.delete - compute.disks.get - compute.disks.list - compute.disks.setLabels - compute.disks.use	To create and manage disks for Cloud Volumes ONTAP.

Actions	Purpose
<ul style="list-style-type: none"> <li>- compute.firewalls.create</li> <li>- compute.firewalls.delete</li> <li>- compute.firewalls.get</li> <li>- compute.firewalls.list</li> </ul>	To create firewall rules for Cloud Volumes ONTAP.
<ul style="list-style-type: none"> <li>- compute.globalOperations.get</li> </ul>	To get the status of operations.
<ul style="list-style-type: none"> <li>- compute.images.get</li> <li>- compute.images.getFromFamily</li> <li>- compute.images.list</li> <li>- compute.images.useReadOnly</li> </ul>	To get images for VM instances.
<ul style="list-style-type: none"> <li>- compute.instances.attachDisk</li> <li>- compute.instances.detachDisk</li> </ul>	To attach and detach disks to Cloud Volumes ONTAP.
<ul style="list-style-type: none"> <li>- compute.instances.create</li> <li>- compute.instances.delete</li> </ul>	To create and delete Cloud Volumes ONTAP VM instances.
<ul style="list-style-type: none"> <li>- compute.instances.get</li> </ul>	To list VM instances.
<ul style="list-style-type: none"> <li>- compute.instances.getSerialPortOutput</li> </ul>	To get console logs.
<ul style="list-style-type: none"> <li>- compute.instances.list</li> </ul>	To retrieve the list of instances in a zone.
<ul style="list-style-type: none"> <li>- compute.instances.setDeletionProtection</li> </ul>	To set deletion protection on the instance.
<ul style="list-style-type: none"> <li>- compute.instances.setLabels</li> </ul>	To add labels.
<ul style="list-style-type: none"> <li>- compute.instances.setMachineType</li> </ul>	To change the machine type for Cloud Volumes ONTAP.
<ul style="list-style-type: none"> <li>- compute.instances.setMetadata</li> </ul>	To add metadata.
<ul style="list-style-type: none"> <li>- compute.instances.setTags</li> </ul>	To add tags for firewall rules.
<ul style="list-style-type: none"> <li>- compute.instances.start</li> <li>- compute.instances.stop</li> <li>- compute.instances.updateDisplayDevice</li> </ul>	To start and stop Cloud Volumes ONTAP.
<ul style="list-style-type: none"> <li>- compute.machineTypes.get</li> </ul>	To get the numbers of cores to check quotas.
<ul style="list-style-type: none"> <li>- compute.projects.get</li> </ul>	To support multi-projects.
<ul style="list-style-type: none"> <li>- compute.snapshots.create</li> <li>- compute.snapshots.delete</li> <li>- compute.snapshots.get</li> <li>- compute.snapshots.list</li> <li>- compute.snapshots.setLabels</li> </ul>	To create and manage persistent disk snapshots.
<ul style="list-style-type: none"> <li>- compute.networks.get</li> <li>- compute.networks.list</li> <li>- compute.regions.get</li> <li>- compute.regions.list</li> <li>- compute.subnetworks.get</li> <li>- compute.subnetworks.list</li> <li>- compute.zoneOperations.get</li> <li>- compute.zones.get</li> <li>- compute.zones.list</li> </ul>	To get the networking information needed to create a new Cloud Volumes ONTAP virtual machine instance.

Actions	Purpose
<ul style="list-style-type: none"> <li>- deploymentmanager.compositeTypes.get</li> <li>- deploymentmanager.compositeTypes.list</li> <li>- deploymentmanager.deployments.create</li> <li>- deploymentmanager.deployments.delete</li> <li>- deploymentmanager.deployments.get</li> <li>- deploymentmanager.deployments.list</li> <li>- deploymentmanager.manifests.get</li> <li>- deploymentmanager.manifests.list</li> <li>- deploymentmanager.operations.get</li> <li>- deploymentmanager.operations.list</li> <li>- deploymentmanager.resources.get</li> <li>- deploymentmanager.resources.list</li> <li>- deploymentmanager.typeProviders.get</li> <li>- deploymentmanager.typeProviders.list</li> <li>- deploymentmanager.types.get</li> <li>- deploymentmanager.types.list</li> </ul>	<p>To deploy the Cloud Volumes ONTAP virtual machine instance using Google Cloud Deployment Manager.</p>
<ul style="list-style-type: none"> <li>- logging.logEntries.list</li> <li>- logging.privateLogEntries.list</li> </ul>	<p>To get stack log drives.</p>
<ul style="list-style-type: none"> <li>- resourcemanager.projects.get</li> </ul>	<p>To support multi-projects.</p>
<ul style="list-style-type: none"> <li>- storage.buckets.create</li> <li>- storage.buckets.delete</li> <li>- storage.buckets.get</li> <li>- storage.buckets.list</li> <li>- storage.buckets.update</li> </ul>	<p>To create and manage a Google Cloud Storage bucket for data tiering.</p>
<ul style="list-style-type: none"> <li>- cloudkms.cryptoKeyVersions.useToEncrypt</li> <li>- cloudkms.cryptoKeys.get</li> <li>- cloudkms.cryptoKeys.list</li> <li>- cloudkms.keyRings.list</li> </ul>	<p>To use customer-managed encryption keys from the Cloud Key Management Service with Cloud Volumes ONTAP.</p>
<ul style="list-style-type: none"> <li>- compute.instances.setServiceAccount</li> <li>- iam.serviceAccounts.getIamPolicy</li> <li>- iam.serviceAccounts.list</li> </ul>	<p>To set a service account on the Cloud Volumes ONTAP instance. This service account provides permissions for data tiering to a Google Cloud Storage bucket.</p>

## AWS Marketplace pages for Cloud Manager and Cloud Volumes ONTAP

Several offerings are available in the AWS Marketplace for Cloud Manager and Cloud Volumes ONTAP. If need help understanding the purpose of each page, read the descriptions below.

In all cases, remember that you can't launch Cloud Volumes ONTAP in AWS from the AWS Marketplace. You need to launch it directly from Cloud Manager.

Goal	AWS Marketplace page to use	More information
<p>Enable the use of Cloud Volumes ONTAP PAYGO, Cloud Tiering, Cloud Compliance, and other add-on services</p>	<p><a href="#">Cloud Manager - Deploy &amp; Manage NetApp Cloud Data Services</a></p>	<p>This subscription enables charging for the PAYGO version of Cloud Volumes ONTAP 9.6 and later. It also enables charging for Cloud Tiering, Cloud Compliance, and other add-on services.</p> <p>You should subscribe to this offering when Cloud Manager prompts you and redirects you to the page. Cloud Manager prompts you in the Working Environment wizard or when you add new credentials in the Settings.</p> <p>This page doesn't enable you to launch Cloud Manager in AWS. That should be done from <a href="#">NetApp Cloud Central</a>, or alternatively using the AMI listed in row 3 of this table.</p>
<p>Enable the use of Cloud Volumes ONTAP PAYGO, Cloud Tiering, Cloud Compliance, and other add-on services <i>using an annual contract</i></p>	<p><a href="#">Cloud Manager (Contracts) - Deploy &amp; Manage NetApp Cloud Data Services</a></p>	<p>This subscription is an alternative to the subscription in the first row. It enables you to get an annual upfront payment for the listings. It's mostly for NetApp partners.</p>
<p>Deploy Cloud Manager from the AWS Marketplace using an AMI</p>	<p><a href="#">Cloud Manager - Manual installation without access keys</a></p>	<p>We recommend that you launch Cloud Manager in AWS from <a href="#">NetApp Cloud Central</a>, but you can launch it from this AWS Marketplace page, if you prefer.</p>
<p>Enable deployment of Cloud Volumes ONTAP PAYGO (9.5 or earlier)</p>	<ul style="list-style-type: none"> <li>• <a href="#">Cloud Volumes ONTAP for AWS</a></li> <li>• <a href="#">Cloud Volumes ONTAP for AWS - High Availability</a></li> </ul>	<p>These AWS Marketplace pages enable you to subscribe to the single node or HA versions of Cloud Volumes ONTAP PAYGO for versions 9.5 and earlier.</p> <p>Starting with version 9.6, you need to subscribe through the AWS Marketplace page listed in row 1 of this table for PAYGO deployments.</p>



# Use APIs and automation

## Automation resources for infrastructure as code

Use the resources on this page to get help integrating Cloud Manager and Cloud Volumes ONTAP with your [infrastructure as code](#).

DevOps teams use a variety of tools to automate the setup of new environments, which allows them to treat infrastructure as code. One such tool is Terraform. We have developed a Terraform provider that DevOps teams can use with Cloud Manager to automate and integrate Cloud Volumes ONTAP with infrastructure as code.

[View the netapp-cloudmmanager provider.](#)

### Related links

- [NetApp Cloud Blog: Using Cloud Manager REST APIs with Federated Access](#)
- [NetApp Cloud Blog: Cloud Automation with Cloud Volumes ONTAP and REST](#)
- [NetApp Cloud Blog: Automated Data Cloning for Cloud-Based Testing of Software Applications](#)
- [NetApp Blog: Infrastructure-As-Code \(IaC\) Accelerated with Ansible + NetApp](#)
- [NetApp thePub: Configuration Management & Automation with Ansible](#)
- [NetApp thePub: Roles for Ansible ONTAP use](#)

# Where to get help and find more information

You can get help and find more information about Cloud Manager and Cloud Volumes ONTAP through various resources, including videos, forums, and support.

- [NetApp Cloud Volumes ONTAP Support](#)

Access support resources to get help and troubleshoot issues with Cloud Volumes ONTAP.

- [Videos for Cloud Manager and Cloud Volumes ONTAP](#)

Watch videos that show you how to deploy and manage Cloud Volumes ONTAP and how to replicate data across your hybrid cloud.

- [Policies for Cloud Manager](#)

Download JSON files that include the permissions that Cloud Manager needs to perform actions in a cloud provider.

- [Cloud Manager API Developer Guide](#)

Read an overview of the APIs, examples of how to use them, and an API reference.

- Training for Cloud Volumes ONTAP

- [Cloud Volumes ONTAP Fundamentals](#)
- [Cloud Volumes ONTAP Deployment and Management for Azure](#)
- [Cloud Volumes ONTAP Deployment and Management for AWS](#)

- Technical reports

- [NetApp Technical Report 4383: Performance Characterization of Cloud Volumes ONTAP in Amazon Web Services with Application Workloads](#)
- [NetApp Technical Report 4671: Performance Characterization of Cloud Volumes ONTAP in Azure with Application Workloads](#)
- [NetApp Technical Report 4816: Performance Characterization of Cloud Volumes ONTAP for Google Cloud](#)

- SVM disaster recovery

SVM disaster recovery is the asynchronous mirroring of SVM data and configuration from a source SVM to a destination SVM. You can quickly activate a destination SVM for data access if the source SVM is no longer available.

- [Cloud Volumes ONTAP 9 SVM Disaster Recovery Preparation Express Guide](#)

Describes how to quickly configure a destination SVM in preparation for disaster recovery.

- [Cloud Volumes ONTAP 9 SVM Disaster Recovery Express Guide](#)

Describes how to quickly activate a destination SVM after a disaster, and then reactivate the source SVM.

- [FlexCache Volumes for Faster Data Access Power Guide](#)

Describes how to create and manage FlexCache volumes in the same cluster or different cluster as the origin volume for accelerating data access.

- [Security advisories](#)

Identify known vulnerabilities (CVEs) for NetApp products, including ONTAP. Note that you can remediate security vulnerabilities for Cloud Volumes ONTAP by following ONTAP documentation.

- [ONTAP 9 Documentation Center](#)

Access product documentation for ONTAP, which can help you as you use Cloud Volumes ONTAP.

- [NetApp Community: Cloud Data Services](#)

Connect with peers, ask questions, exchange ideas, find resources, and share best practices.

- [NetApp Cloud Central](#)

Find information about additional NetApp products and solutions for the cloud.

- [NetApp Product Documentation](#)

Search NetApp product documentation for instructions, resources, and answers.

# Earlier versions of Cloud Manager documentation

Documentation for previous releases of Cloud Manager is available in case you're not running the latest version.

- [Cloud Manager 3.7](#)
- [Cloud Manager 3.6](#)
- [Cloud Manager 3.5](#)
- [Cloud Manager 3.4](#)
- [Cloud Manager 3.3](#)
- [Cloud Manager 3.2](#)

# Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

## Copyright

<http://www.netapp.com/us/legal/copyright.aspx>

## Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

## Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/us/media/patents-page.pdf>

## Privacy policy

<https://www.netapp.com/us/legal/privacypolicy/index.aspx>

## Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

- [Notice for Cloud Manager 3.8.7](#)
- [Notice for Cloud Manager 3.8.6](#)
- [Notice for Cloud Manager 3.8.5](#)
- [Notice for Cloud Manager 3.8.4](#)
- [Notice for Cloud Manager 3.8.3](#)
- [Notice for Cloud Manager 3.8.2](#)
- [Notice for Cloud Manager 3.8.1](#)
- [Notice for Cloud Manager 3.8](#)
- [Notice for the Cloud Backup Service](#)
- [Notice for Global File Cache](#)
- [Notice for Cloud Compliance](#)

## 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.

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.