



Back up on-premises ONTAP data

Cloud Manager

NetApp
January 21, 2022

Table of Contents

- Back up on-premises ONTAP data 1
 - Backing up on-premises ONTAP data to Amazon S3 1
 - Backing up on-premises ONTAP data to Azure Blob storage 13
 - Backing up on-premises ONTAP data to Google Cloud Storage 22
 - Backing up on-premises ONTAP data to StorageGRID 29

Back up on-premises ONTAP data

Backing up on-premises ONTAP data to Amazon S3

Complete a few steps to get started backing up data from your on-premises ONTAP systems to Amazon S3 storage.



In most cases you'll use Cloud Manager for all backup and restore operations. However, starting with ONTAP 9.9.1 you can initiate volume backup operations of your on-premises ONTAP clusters using ONTAP System Manager. [See how to use System Manager to back up your volumes to the cloud using Cloud Backup.](#)

A Beta feature released in January 2021 allows you to run compliance scans on the backed up volumes from your on-premises systems. Typically, compliance scans are free up to 1 TB of data, and then a cost for the service is applied for data over 1 TB. When combining Backup and Data Sense for your on-premises volumes, the cost for scans on those on-prem volumes is free. Learn more about how [Cloud Data Sense](#) can get your business applications and cloud environments privacy ready.

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.
 - The cluster is running ONTAP 9.7P5 or later.
 - The cluster has a SnapMirror license — it is included as part of the Premium Bundle or Data Protection Bundle.
 - The cluster must have the required network connections to S3 storage and to the Connector.
- The Connector must have the required network connections to S3 storage and to the cluster, and the required permissions.
- You have a valid AWS subscription for the object storage space where your backups will be located.
- You have an AWS Account with an access key and secret key, and the [required permissions](#) so the ONTAP cluster can back up and restore data.

2

Enable Cloud Backup on the system

Select the working environment and click **Enable** next to the Backup & Compliance service in the right-panel, and then follow the setup wizard.



3**Select the cloud provider and enter the provider details**

Select Amazon Web Services as your provider and then enter the provider details. You'll need to select the AWS Account and the region where you want to create the backups. You can also choose your own customer-managed key for data encryption instead of using the default Amazon S3 encryption key.

4**Select the cluster IPspace and optionally select an AWS PrivateLink connection**

Select the IPspace in the ONTAP cluster where the volumes reside. You can also choose to use an existing AWS PrivateLink configuration for a more secure connection to the VPC from your on-prem data center.

Name	VPC	Endpoint ID
<input type="radio"/> Private_Link_Name_001	vpce0-012345678901234567890 (Default)	vpce0-012345678901234567890
<input type="radio"/> Private_Link_Name_002	vpce0-012345678901234567890 (k8s)	vpce0-012345678901234567890

5**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 hourly, daily, 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 you want to retain.

By default, backups are stored in S3 Standard storage. If your cluster is using ONTAP 9.10.1 or greater, you can choose to tier backups to either S3 Glacier or S3 Glacier Deep Archive storage after a certain number of days for further cost optimization.

6

Select the volumes that you want to back up

Identify which volumes you want to back up from the cluster.

7

Activate Compliance scans on the backed up volumes (optional)

Choose whether you want to have Cloud Data Sense scan the volumes that are backed up in the cloud.

8

Restore your data, as needed


Choose to restore an entire backup to a new volume, or to restore individual files from the backup to an existing volume. You can restore data to a Cloud Volumes ONTAP system in AWS, or to an on-premises ONTAP system.

See [Restoring volume data from backup files](#) for details.

Requirements

Read the following requirements to make sure you have a supported configuration before you start backing up on-premises volumes to S3 storage.

The following image shows each component and the connections that you need to prepare between them:



¹ Cloud Restore instance is active only during single-file restore operations.

Note that when the Cloud Restore instance is deployed in the cloud, it is located in the same subnet as the Connector.

Preparing your ONTAP clusters

You need to discover your on-premises ONTAP clusters in Cloud Manager before you can start backing up volume data.

[Learn how to discover a cluster.](#)

ONTAP requirements

- ONTAP 9.7P5 and later.
- A SnapMirror license (included as part of the Premium Bundle or Data Protection Bundle).

Note: The "Hybrid Cloud Bundle" is not required when using the Cloud Backup service.

See how to [manage your cluster licenses](#).

- Time and time zone are set correctly.

See how to [configure your cluster time](#).

Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 from the intercluster LIF to Amazon S3 storage for backup and restore operations.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

- ONTAP requires an inbound connection from the Connector to the cluster management LIF. The Connector can reside in an AWS VPC.
- 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. [Learn more about IPspaces](#).

When you set up Cloud Backup, 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.

- The nodes' intercluster LIFs are able to access the object store.
- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- Note that if you are using a different IPspace than the Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow Cloud Backup service connections from ONTAP to object storage through port 443 and name resolution traffic from the storage VM to the DNS server over port 53 (TCP/UDP).

Creating or switching Connectors

A Connector is required to back up data to the cloud, and the Connector must be in an AWS VPC when backing up data to AWS S3 storage. You can't use a Connector that's deployed on-premises. You'll either need to create a new Connector or make sure that the currently selected Connector resides in the correct provider.

- [Learn about Connectors](#)
- [Creating a Connector in AWS](#)
- [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 Backup service over port 443 (HTTPS)
 - An HTTPS connection over port 443 to your S3 object storage
 - An HTTPS connection over port 443 to your ONTAP clusters
2. Enable a VPC Endpoint to S3. This is needed 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.

Supported regions

You can create backups from on-premises systems to Amazon S3 in all regions [where Cloud Volumes ONTAP is supported](#). You specify the region where the backups will be stored when you set up the service.

License requirements

Before your 30-day free trial of the Cloud Backup service expires, you need to subscribe to a pay-as-you-go (PAYGO) Cloud Manager Marketplace offering from AWS, or purchase and activate a Cloud Backup BYOL license from NetApp. These licenses are for the account and can be used across multiple systems.

- For Cloud Backup PAYGO licensing, you'll need a subscription to the [AWS Cloud Manager Marketplace](#) offering to continue using Cloud Backup. Billing for Cloud Backup is done through this subscription.
- For Cloud Backup BYOL licensing, you don't need a subscription. You need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#).

You need to have an AWS subscription for the object storage space where your backups will be located.

A SnapMirror license is required on the cluster. Note that the "Hybrid Cloud Bundle" is not required when using Cloud Backup.

Preparing Amazon S3 for backups

When you are using Amazon S3, you must configure permissions for the Connector to create and manage the S3 bucket, and you must configure permissions so the on-premises ONTAP cluster can read and write to the S3 bucket.

Steps

1. Confirm that the following S3 permissions (from the latest [Cloud Manager policy](#)) are part of the IAM role that provides the Connector with permissions:


```

{
  "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-*"
  ]
}

```

2. Add the following EC2 permissions to the IAM role that provides the Connector with permissions so that it can start, stop, and terminate the Cloud Restore instance:

```

"Action": [
  "ec2:DescribeInstanceTypeOfferings",
  "ec2:StartInstances",
  "ec2:StopInstances",
  "ec2:TerminateInstances"
]

```

3. During the Backup wizard you will be prompted to enter an access key and secret key. For that, you will need to create an IAM user with the following permissions. Cloud Backup passes these credentials on to the ONTAP cluster so that ONTAP can backup and restore data to the S3 bucket.

```
"s3:ListAllMyBuckets",  
"s3:ListBucket",  
"s3:GetBucketLocation",  
"s3:GetObject",  
"s3:PutObject",  
"s3:DeleteObject"
```

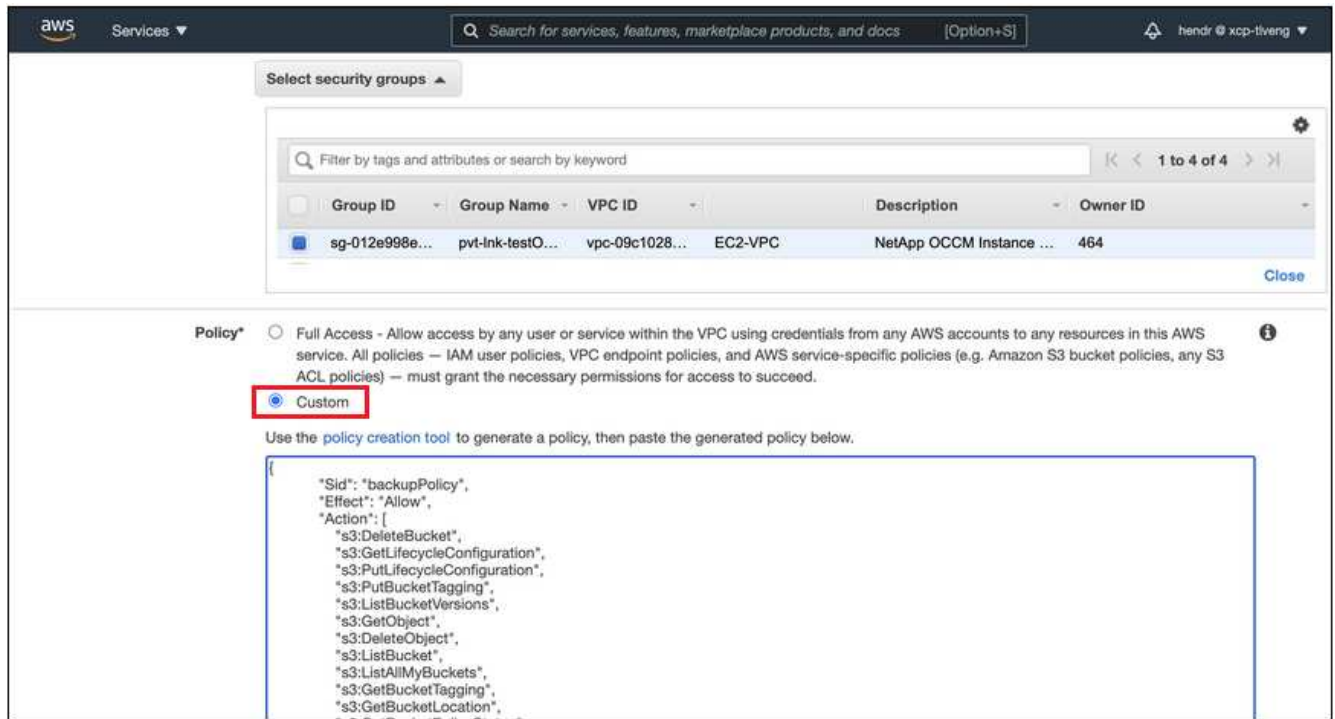
See the [AWS Documentation: Creating a Role to Delegate Permissions to an IAM User](#) for details.

4. If your virtual or physical network uses a proxy server for internet access, ensure that the Cloud Restore instance has outbound internet access to contact the following endpoints.

Endpoints	Purpose
http://amazonlinux.us-east-1.amazonaws.com/2/extras/docker/stable/x86_64/4bf88ee77c395ffe1e0c3ca68530dfb3a683ec65a4a1ce9c0ff394be50e922b2/	CentOS package for the Cloud Restore Instance AMI.
http://cloudmanagerinfraprod.azurecr.io https://cloudmanagerinfraprod.azurecr.io	Cloud Restore Instance image repository.

5. You can choose your own custom-managed keys for data encryption in the activation wizard instead of using the default Amazon S3 encryption keys. In this case you'll need to have the encryption managed keys already set up. [See how to use your own keys](#).
6. If you want to have a more secure connection over the public internet from your on-prem data center to the VPC, there is an option to select an AWS PrivateLink connection in the activation wizard. It is required if you are connecting your on-premises system via VPN/DirectConnect. In this case you'll need to have created an Interface endpoint configuration using the Amazon VPC console or the command line. [See details about using AWS PrivateLink](#).

Note that you'll also need to modify the security group configuration that is associated with the Cloud Manager Connector. You must change the policy to "Custom" (from "Full Access"), and you must add the permissions from the backup policy as shown earlier (above).



Enabling Cloud Backup

Enable Cloud Backup at any time directly from the on-premises working environment.

Steps

1. From the Canvas, select the working environment and click **Enable** next to the Backup & Compliance service in the right-panel.



2. Select Amazon Web Services as your provider and click **Next**.
3. Enter the provider details and click **Next**.
 - a. The AWS Account, the AWS Access Key, and the Secret Key used to store the backups.

The access key and secret key are for the user you created to give the ONTAP cluster access to the S3 bucket.

- b. The AWS region where the backups will be stored.
- c. Whether you'll use the default Amazon S3 encryption keys or choose your own customer-managed keys from your AWS account to manage encryption of your data. ([See how to use your own keys](#)).

4. Enter the networking details and click **Next**.

- a. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.
- b. Optionally, choose whether you'll use an AWS PrivateLink that you have previously configured. [See details about using an AWS PrivateLink.](#)

Name	VPC	Endpoint ID
<input type="radio"/> Private_Link_Name_001	vpce0-012345678901234567890 (Default)	vpce0-012345678901234567890
<input type="radio"/> Private_Link_Name_002	vpce0-012345678901234567890 (k8s)	vpce0-012345678901234567890

5. Enter the backup policy details and click **Next**.

- a. Define the backup schedule and choose the number of backups to retain. [See the list of existing policies you can choose.](#)
- b. When using ONTAP 9.10.1 and greater, you can choose to tier backups to either S3 Glacier or S3 Glacier Deep Archive storage after a certain number of days for further cost optimization. [Learn more about using archival tiers.](#)

Define Policy

Policy - Retention & Schedule Create a New Policy Select an Existing Policy

Hourly Number of backups to retain

Daily Number of backups to retain

Weekly Number of backups to retain

Monthly Number of backups to retain

Archival Policy Backups reside in S3 Standard storage for frequently accessed data. Optionally, you can tier backups to either S3 Glacier or S3 Glacier Deep Archive storage for further cost optimization.

Tier Backups to Archival

Archive after (Days) Storage Class

S3 Glacier

S3 Glacier Deep Archive

S3 Bucket Cloud Manager will create the S3 bucket for you. Wizard

6. Select the volumes that you want to back up.

- To back up all volumes, check the box in the title row (Volume Name).
- To back up individual volumes, check the box for each volume (Volume_1).

Select Volumes

57 Volumes 🔍

<input checked="" type="checkbox"/>	Volume Name	Volume Type	SVM Name	Used Capacity	Allocated Capacity	Backup Status
<input checked="" type="checkbox"/>	Volume_Name_1	RW	SVM_Name_1	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_2	RW	SVM_Name_2	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_3	RW	SVM_Name_3	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_4	DP	SVM_Name_4	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_5	RW	SVM_Name_5	0.25 TB	10 TB	⊖ Not Active

7. Click **Activate Backup** and Cloud Backup starts taking the initial backups of your volumes.

You are prompted whether you want to run compliance scans on the backed up volumes. Cloud Data Sense scans are free when you run them on the backed up volumes (except for the [cost of the deployed Cloud Data Sense instance](#)).

Activate Cloud Data Sense on your Backed Up Volumes

You have successfully activated Backup to Cloud on 57 Volumes in your working environment "On-Premises".

Cloud Data Sense

- Cloud Data Sense offers automated controls for data privacy regulations such as the GDPR, CCPA and more.
- Driven by powerful artificial intelligence algorithms, Cloud Data Sense gets your business application data and cloud environments privacy ready.

Go to Cloud Data Sense **Cancel**

8. Click **Go to Compliance** to activate compliance scans on the volumes. (If you choose **Close** and not to scan these backed up volumes, you can always [enable this functionality](#) later from Cloud Data Sense.)
 - If an instance of Cloud Data Sense is already deployed in your environment, you are directed to the Configuration page to select the volumes you want to scan in each on-premises working environment that has backups. See [how to choose the volumes](#).

Data Sense Governance Compliance Investigation Policies **Configuration**

(2/20) Working Environments

Filter by: CVO ANF S3 DB ONEDR **BACKUP** [Clear filters](#)

Working Environment 1 (back up)
Cloud Backup of ONTAP **BETA**

Map all backed up Volumes **Map & Classify all backed up Volumes**

Or select scanning type for each volume

- If Cloud Data Sense has not been deployed, you are directed to the Compliance page where you can choose to deploy Compliance in the cloud or in your premises. We strongly recommend deploying it in the cloud. Go [here](#) for installation requirements and instructions.

Data Sense

How does it work? [?](#)

Always-on Privacy & Compliance Controls

Automated controls for data privacy regulations - GDPR, CCPA, HIPAA and more. Driven by powerful artificial intelligence algorithms, Data Sense gets your business application data and cloud environments privacy ready.

Deploy Data Sense in the Cloud **Deploy Data Sense On-Premises**

[Learn about the differences between cloud deployment and on-premises deployment](#)

Compliance Status

Data Distribution

75% Non-Sensitive	20% Personal	5% Sensitive Personal
-------------------	--------------	-----------------------

28,000 Personal Files [View All](#)

7,000 Sensitive Personal Files [View All](#)

Email Address	2,700 Files	Health	2,700 Files
Credit Card	2,700 Files	Ethnicity	2,700 Files

After you have deployed Compliance you can choose the volumes you want to scan as described above.

Result

Cloud Backup backs up your volumes from the on-premises ONTAP system, and optionally, Cloud Data Sense runs compliance scans on the backed up volumes.

What's next?

You can [start and stop backups for volumes or change the backup schedule](#) and you can [restore entire volumes or individual files from a backup file](#).

You can also [view the results of the compliance scans](#) and review other features of Cloud Data Sense that can help you understand data context and identify sensitive data in your organization.



The scan results are not available immediately because Cloud Backup has to finish creating the backups before Cloud Data Sense can start compliance scans.

Backing up on-premises ONTAP data to Azure Blob storage

Complete a few steps to get started backing up data from your on-premises ONTAP systems to Azure Blob storage.



In most cases you'll use Cloud Manager for all backup and restore operations. However, starting with ONTAP 9.9.1 you can initiate volume backup operations of your on-premises ONTAP clusters using ONTAP System Manager. [See how to use System Manager to back up your volumes to the cloud using Cloud Backup](#).

A Beta feature released in January 2021 allows you to run compliance scans on the backed up volumes from your on-premises systems. Typically, compliance scans are free up to 1 TB of data, and then a cost for the service is applied for data over 1 TB. When combining Backup and Data Sense for your on-premises volumes, the cost for scans on those on-prem volumes is free. Learn more about how [Cloud Data Sense](#) can get your business applications and cloud environments privacy ready.

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 have discovered the on-premises cluster and added it to a working environment in Cloud Manager. See [Discovering ONTAP clusters](#) for details.
 - The cluster is running ONTAP 9.7P5 or later.
 - The cluster has a SnapMirror license — it is included as part of the Premium Bundle or Data Protection Bundle.
 - The cluster must have the required network connections to Blob storage and to the Connector.
- The Connector must have the required network connections to Blob storage and to the cluster, and the required permissions.
- You have a valid Azure subscription for the object storage space where your backups will be located.

2

Enable Cloud Backup on the system


Select the working environment and click **Enable** next to the Backup & Compliance service in the right-panel, and then follow the setup wizard.



3

Select the cloud provider and enter the provider details

Select Microsoft Azure as your provider and then enter the provider details. You'll need to select the Azure Subscription and the region where you want to create the backups. You can also choose your own customer-managed key for data encryption instead of using the default Microsoft-managed encryption key.



4

Select the cluster IPspace and optional use of a private VNet endpoint

Select the IPspace in the ONTAP cluster where the volumes reside. You can also choose to use an existing Azure Private Endpoint for a more secure connection to the VNet from your on-prem data center.



5

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 hourly, daily, 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 you want to retain.

By default, backups are stored in the Cool access tier. If your cluster is using ONTAP 9.10.1 or greater, you can choose to tier backups to Azure Archive storage after a certain number of days for further cost optimization.

Define Policy

Policy - Retention & Schedule Create a New Policy Select an Existing Policy

Hourly Number of backups to retain 24

Daily Number of backups to retain 30

Weekly Number of backups to retain 52

Monthly Number of backups to retain 12

Archival Policy

Backups reside in Cool Azure Blob storage for frequently accessed data.
Optionally, you can tier backups to Azure Archive storage for further cost optimization.

Tier Backups to Archival

Archive after (Days) 30 Access Tier Azure Archive

Storage Account Cloud Manager will create the storage account after you complete the wizard

6

Select the volumes that you want to back up

Identify which volumes you want to back up from the cluster.

7

Activate Compliance scans on the backed up volumes (optional)

Choose whether you want to have Cloud Data Sense scan the volumes that are backed up in the cloud.

8

Restore your data, as needed


Choose to restore an entire backup to a new volume, or to restore individual files from the backup to an existing volume. You can restore data to a Cloud Volumes ONTAP system in Azure, or to an on-premises ONTAP system.

See [Restoring volume data from backup files](#) for details.

Requirements

Read the following requirements to make sure you have a supported configuration before you start backing up on-premises volumes to Azure Blob storage.

The following image shows each component and the connections that you need to prepare between them:



¹ Cloud Restore instance is active only during single-file restore operations.

Note that when the Cloud Restore instance is deployed in the cloud, it is located in the same subnet as the Connector.

Preparing your ONTAP clusters

You need to discover your on-premises ONTAP clusters in Cloud Manager before you can start backing up volume data.

[Learn how to discover a cluster.](#)

ONTAP requirements

- ONTAP 9.7P5 and later.
- A SnapMirror license (included as part of the Premium Bundle or Data Protection Bundle).

Note: The "Hybrid Cloud Bundle" is not required when using the Cloud Backup service.

See how to [manage your cluster licenses](#).

- Time and time zone are set correctly.

See how to [configure your cluster time](#).

Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 from the intercluster LIF to Azure Blob storage for backup and restore operations.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

- ONTAP requires an inbound connection from the Connector to the cluster management LIF. The Connector can reside in an Azure VNet.
- 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. [Learn more about IPspaces](#).

When you set up Cloud Backup, 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.

- The nodes' and intercluster LIFs are able to access the object store.
- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- Note that if you are using a different IPspace than the Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow Cloud Backup service connections from ONTAP to object storage through port 443 and name resolution traffic from the storage VM to the DNS server over port 53 (TCP/UDP).

Creating or switching Connectors

A Connector is required to back up data to the cloud, and the Connector must be in an Azure VNet when backing up data to Azure Blob storage. You can't use a Connector that's deployed on-premises. You'll either need to create a new Connector or make sure that the currently selected Connector resides in the correct provider.

- [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 network where the Connector is installed enables the following connections:
 - An outbound internet connection to the Cloud Backup service over port 443 (HTTPS)
 - An HTTPS connection over port 443 to your Blob object storage
 - An HTTPS connection over port 443 to your ONTAP clusters
2. Enable a VNet Private Endpoint to Azure storage. This is needed 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.

Supported regions

You can create backups from on-premises systems to Azure Blob in all regions [where Cloud Volumes ONTAP is supported](#). You specify the region where the backups will be stored when you set up the service.

License requirements

Before your 30-day free trial of the Cloud Backup service expires, you need to subscribe to a pay-as-you-go (PAYGO) Cloud Manager Marketplace offering from Azure, or purchase and activate a Cloud Backup BYOL license from NetApp. These licenses are for the account and can be used across multiple systems.

- For Cloud Backup PAYGO licensing, you'll need a subscription to the [Azure](#) Cloud Manager Marketplace offering to continue using Cloud Backup. Billing for Cloud Backup is done through this subscription.
- For Cloud Backup BYOL licensing, you don't need a subscription. You need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#).

You need to have an Azure subscription for the object storage space where your backups will be located.

A SnapMirror license is required on the cluster. Note that the "Hybrid Cloud Bundle" is not required when using Cloud Backup.

Preparing Azure Blob storage for backups

1. If your virtual or physical network uses a proxy server for internet access, ensure that the Cloud Restore virtual machine has outbound internet access to contact the following endpoints.

Endpoints	Purpose
http://olcentgbl.trafficmanager.net https://olcentgbl.trafficmanager.net	Provides CentOS packages for the Cloud Restore virtual machine.
http://cloudmanagerinfraprod.azurecr.io https://cloudmanagerinfraprod.azurecr.io	Cloud Restore virtual machine image repository.

2. You use choose your own custom-managed keys for data encryption in the activation wizard instead of using the default Microsoft-managed encryption keys. In this case you will need to have the Azure Subscription, Key Vault name, and the Key. [See how to use your own keys](#).
3. If you want to have a more secure connection over the public internet from your on-prem data center to the VNet, there is an option to configure an Azure Private Endpoint in the activation wizard. In this case you will need to know the VNet and Subnet for this connection. [See details about using a Private Endpoint](#).

Enabling Cloud Backup

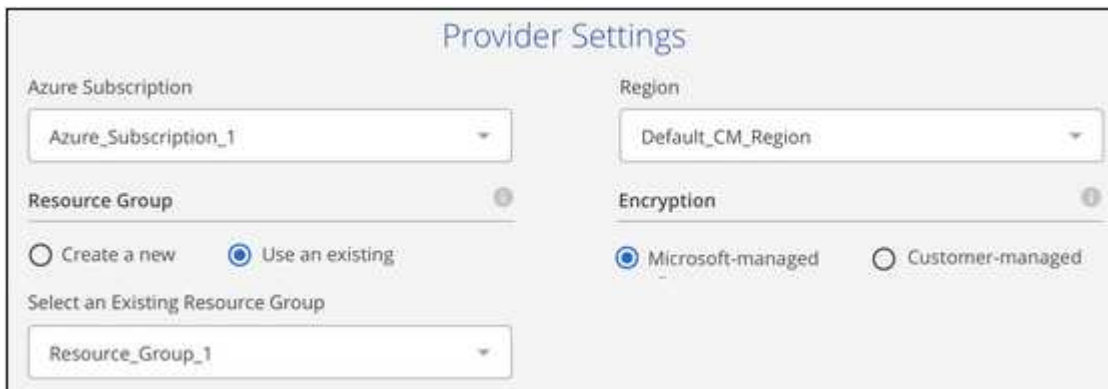
Enable Cloud Backup at any time directly from the on-premises working environment.

Steps

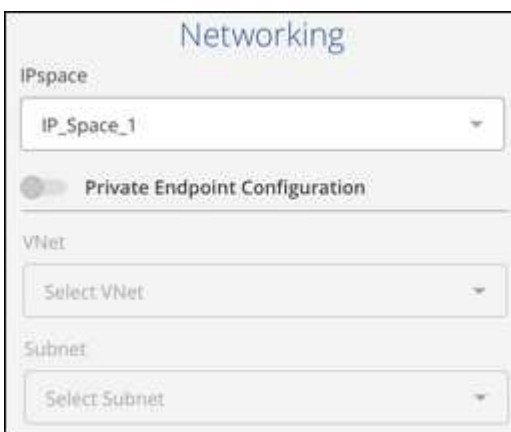
1. From the Canvas, select the working environment and click **Enable** next to the Backup & Compliance service in the right-panel.



2. Select Microsoft Azure as your provider and click **Next**.
3. Enter the provider details and click **Next**.
 - a. The Azure subscription used for backups and the Azure region where the backups will be stored.
 - b. The resource group that manages the Blob container - you can create a new resource group or select an existing resource group.
 - c. Whether you will use the default Microsoft-managed encryption key or choose your own customer-managed keys to manage encryption of your data. ([See how to use your own keys](#)).



4. Enter the networking details and click **Next**.
 - a. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.
 - b. Optionally, choose whether you will configure an Azure Private Endpoint. [See details about using a Private Endpoint](#).



5. Enter the backup policy details and click **Next**.
 - a. Define the backup schedule and choose the number of backups to retain. [See the list of existing policies you can choose](#).
 - b. When using ONTAP 9.10.1 and greater, you can choose to tier backups to Azure Archive storage after a certain number of days for further cost optimization. [Learn more about using archival tiers](#).

Define Policy

Policy - Retention & Schedule Create a New Policy Select an Existing Policy

Hourly Number of backups to retain: 24

 Daily Number of backups to retain: 30

 Weekly Number of backups to retain: 52

 Monthly Number of backups to retain: 12

Archival Policy

Backups reside in Cool Azure Blob storage for frequently accessed data. Optionally, you can tier backups to Azure Archive storage for further cost optimization.

Tier Backups to Archival

Archive after (Days): Access Tier:

Storage Account Cloud Manager will create the storage account after you complete the wizard

6. Select the volumes that you want to back up.

- To back up all volumes, check the box in the title row (Volume Name).
- To back up individual volumes, check the box for each volume (Volume_1).

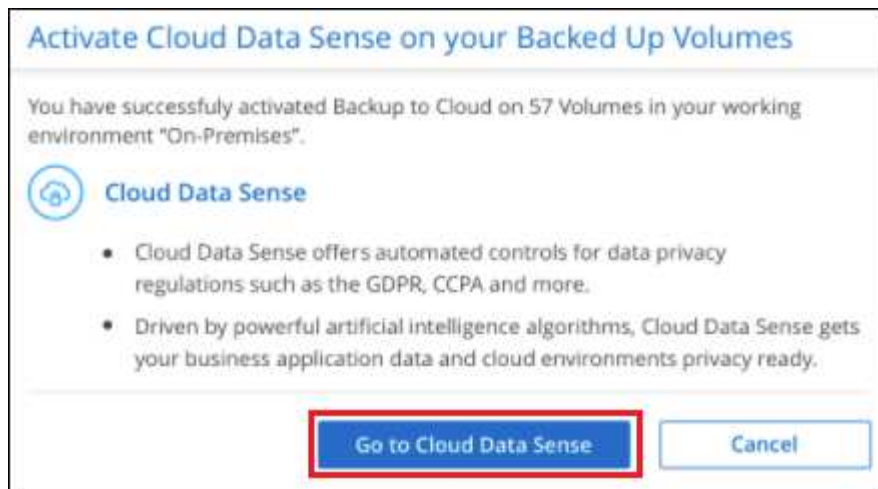
Select Volumes

57 Volumes 🔍

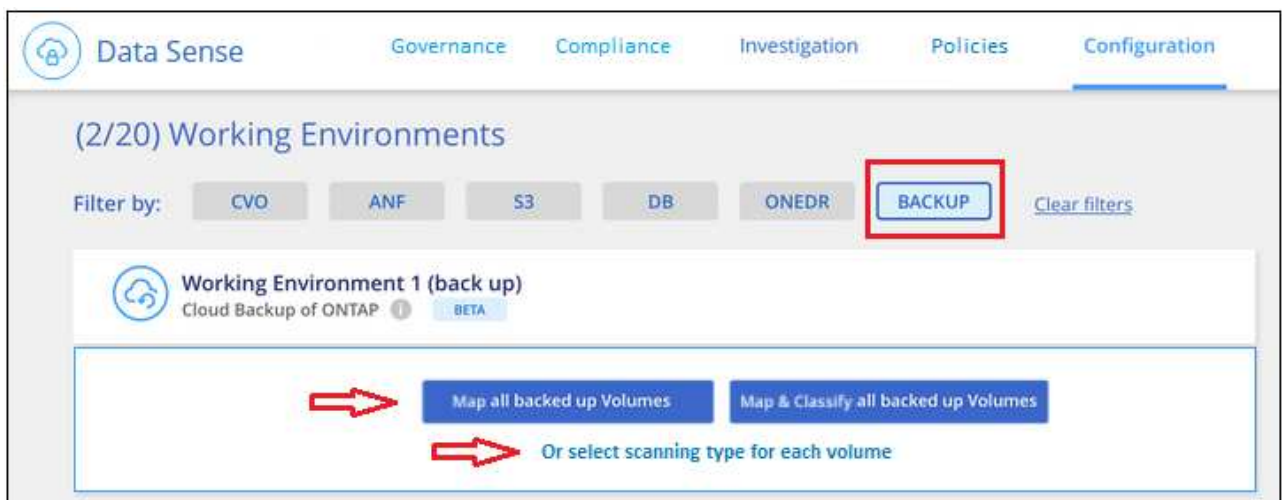
<input checked="" type="checkbox"/>	Volume Name	Volume Type	SVM Name	Used Capacity	Allocated Capacity	Backup Status
<input checked="" type="checkbox"/>	Volume_Name_1	RW	SVM_Name_1	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_2	RW	SVM_Name_2	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_3	RW	SVM_Name_3	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_4	DP	SVM_Name_4	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_5	RW	SVM_Name_5	0.25 TB	10 TB	⊖ Not Active

7. Click **Activate Backup** and Cloud Backup starts taking the initial backups of your volumes.

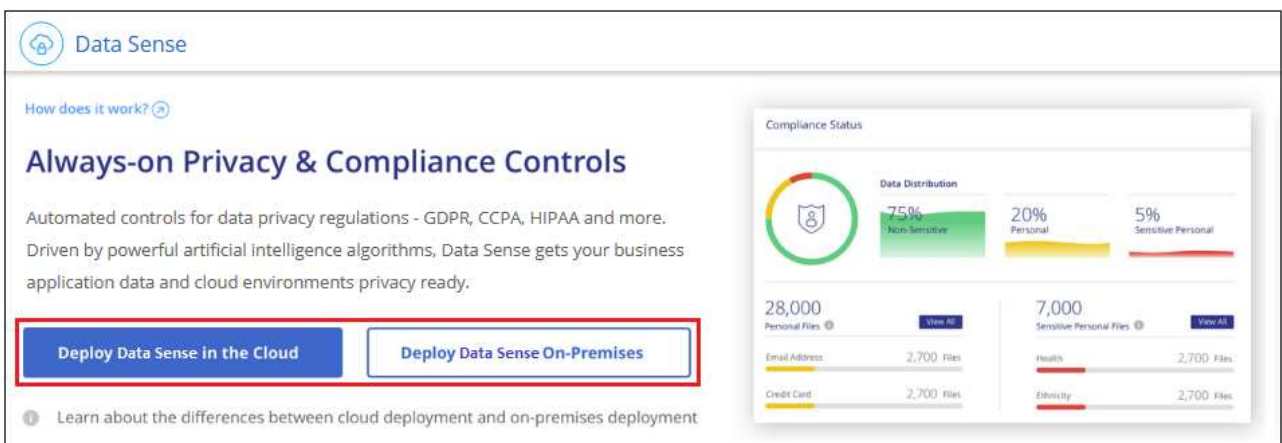
You are prompted whether you want to run compliance scans on the backed up volumes. Cloud Data Sense scans are free when you run them on the backed up volumes (except for the [cost of the deployed Cloud Data Sense instance](#)).



8. Click **Go to Compliance** to activate compliance scans on the volumes. (If you choose **Close** and not to scan these backed up volumes, you can always [enable this functionality](#) later from Cloud Data Sense.)
 - If an instance of Cloud Data Sense is already deployed in your environment, you are directed to the Configuration page to select the volumes you want to scan in each on-premises working environment that has backups. See [how to choose the volumes](#).



- If Cloud Data Sense has not been deployed, you are directed to the Compliance page where you can choose to deploy Compliance in the cloud or in your premises. We strongly recommend deploying it in the cloud. Go [here](#) for installation requirements and instructions.



After you have deployed Compliance you can choose the volumes you want to scan as described above.

Result

Cloud Backup backs up your volumes from the on-premises ONTAP system, and optionally, Cloud Data Sense runs compliance scans on the backed up volumes.

What's next?

You can [start and stop backups for volumes or change the backup schedule](#) and you can [restore entire volumes or individual files from a backup file](#).

You can also [view the results of the compliance scans](#) and review other features of Cloud Data Sense that can help you understand data context and identify sensitive data in your organization.



The scan results are not available immediately because Cloud Backup has to finish creating the backups before Cloud Data Sense can start compliance scans.

Backing up on-premises ONTAP data to Google Cloud Storage

Complete a few steps to get started backing up data from your on-premises ONTAP systems to Google Cloud Storage.



In most cases you'll use Cloud Manager for all backup and restore operations. However, starting with ONTAP 9.9.1 you can initiate volume backup operations of your on-premises ONTAP clusters using ONTAP System Manager. [See how to use System Manager to back up your volumes to the cloud using Cloud Backup](#).

A Beta feature released in January 2021 allows you to run compliance scans on the backed up volumes from your on-premises systems. Typically, compliance scans are free up to 1 TB of data, and then a cost for the service is applied for data over 1 TB. When combining Backup and Data Sense for your on-premises volumes, the cost for scans on those on-prem volumes is free. Learn more about how [Cloud Data Sense](#) can get your business applications and cloud environments privacy ready.

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.
 - The cluster is running ONTAP 9.7P5 or later.
 - The cluster has a SnapMirror license — it is included as part of the Premium Bundle or Data Protection Bundle.
 - The cluster must have the required network connections to Google storage and to the Connector.
- The Connector must have the required network connections to Google storage and to the cluster.

- You have a valid Google subscription for the object storage space where your backups will be located.
- You have a Google account with an access key and secret key so the ONTAP cluster can back up and restore data.

2 Enable Cloud Backup on the system

Select the working environment and click **Enable** next to the Backup & Compliance service in the right-panel, and then follow the setup wizard.

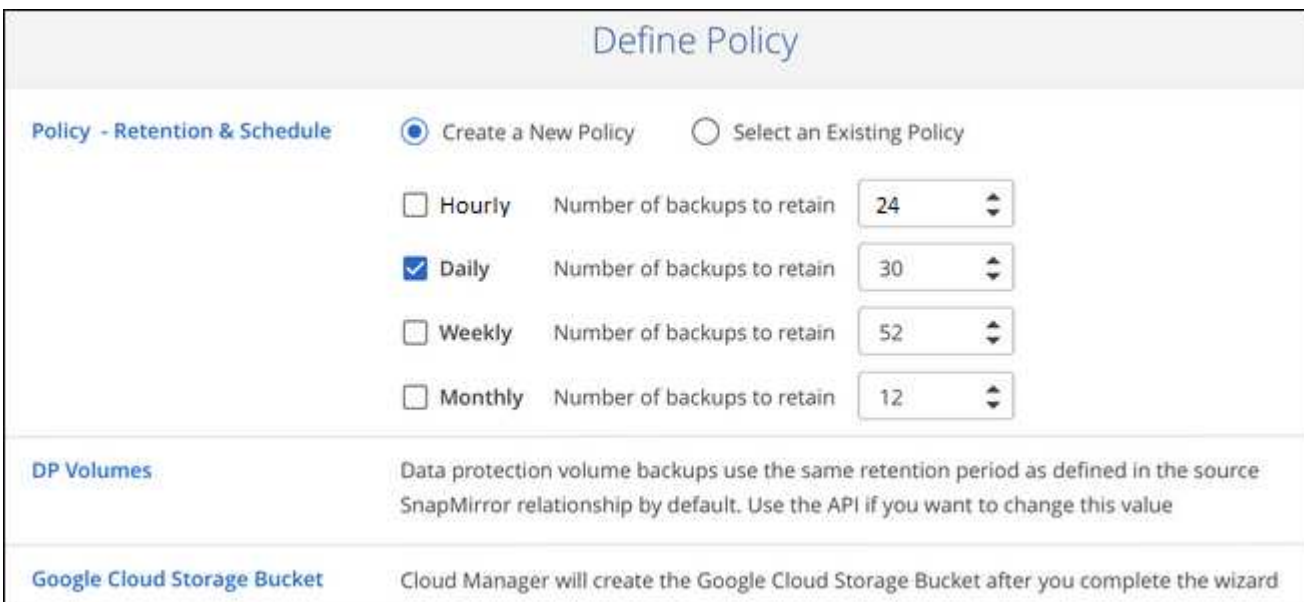


3 Select the cloud provider and enter the provider details

Select Google Cloud as your provider and then enter the provider details. You also need to specify the IPspace in the ONTAP cluster where the volumes reside.

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 hourly, daily, 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 you want to retain.



5 Select the volumes that you want to back up

Identify which volumes you want to back up from the cluster.

6

Activate Compliance scans on the backed up volumes (optional)

Choose whether you want to have Cloud Data Sense scan the volumes that are backed up in the cloud.

7

Restore your data, as needed

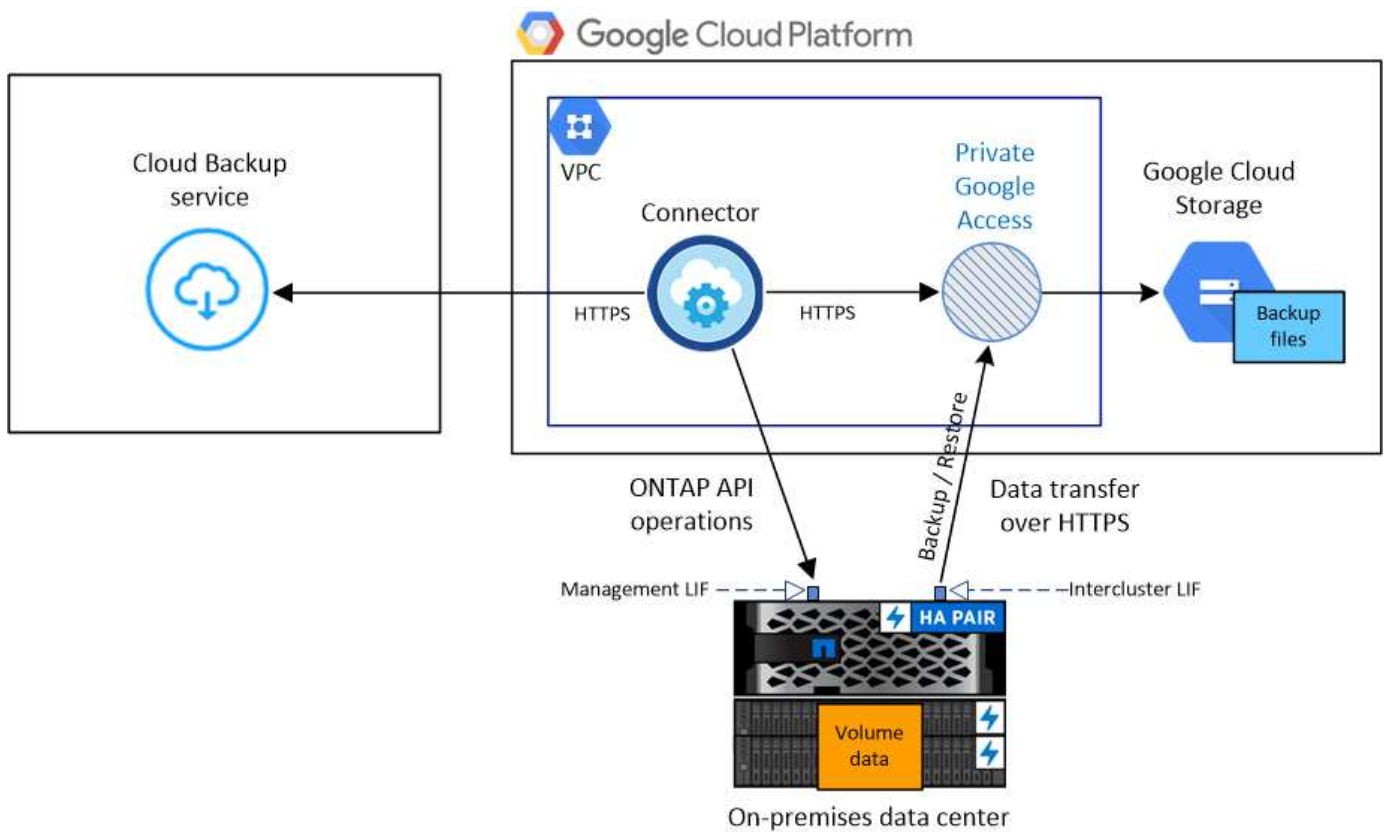
Restore a backup to a new volume. You can restore data to a Cloud Volumes ONTAP system in Google, or to an on-premises ONTAP system.

See [Restoring volume data from backup files](#) for details.

Requirements

Read the following requirements to make sure you have a supported configuration before you start backing up on-premises volumes to Google Cloud storage.

The following image shows each component and the connections that you need to prepare between them:



Note that the Cloud Restore instance is not shown in this diagram because single-file restore is not currently supported in GCP.

Preparing your ONTAP clusters

You need to discover your on-premises ONTAP clusters in Cloud Manager before you can start backing up volume data.

[Learn how to discover a cluster.](#)

ONTAP requirements

- ONTAP 9.7P5 and later.
- A SnapMirror license (included as part of the Premium Bundle or Data Protection Bundle).

Note: The "Hybrid Cloud Bundle" is not required when using the Cloud Backup service.

See how to [manage your cluster licenses](#).

- Time and time zone are set correctly.

See how to [configure your cluster time](#).

Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over port 443 from the intercluster LIF to Google Cloud storage for backup and restore operations.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

- ONTAP requires an inbound connection from the Connector to the cluster management LIF. The Connector can reside in a Google Cloud Platform VPC.
- 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. [Learn more about IPspaces](#).

When you set up Cloud Backup, 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.

- The nodes' intercluster LIFs are able to access the object store.
- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- Note that if you use are using a different IPspace than the Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow Cloud Backup service connections from ONTAP to object storage through port 443 and name resolution traffic from the storage VM to the DNS server over port 53 (TCP/UDP).

Creating or switching Connectors

A Connector is required to back up data to the cloud, and the Connector must be in a Google Cloud Platform VPC when backing up data to Google Cloud storage. You can't use a Connector that's deployed on-premises. You'll either need to create a new Connector or make sure that the currently selected Connector resides in the correct provider.

- [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 network where the Connector is installed enables the following connections:
 - An outbound internet connection to the Cloud Backup service over port 443 (HTTPS)
 - An HTTPS connection over port 443 to your Google Cloud storage
 - An HTTPS connection over port 443 to your ONTAP clusters
2. Enable Private Google Access on the subnet where you plan to deploy the Connector. [Private Google Access](#) is needed 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).

Supported regions

You can create backups from on-premises systems to Google Cloud storage in all regions [where Cloud Volumes ONTAP is supported](#). You specify the region where the backups will be stored when you set up the service.

License requirements

Before your 30-day free trial of the Cloud Backup service expires, you need to subscribe to a pay-as-you-go (PAYGO) Cloud Manager Marketplace offering from Google, or purchase and activate a Cloud Backup BYOL license from NetApp. These licenses are for the account and can be used across multiple systems.

- For Cloud Backup PAYGO licensing, you'll need a subscription to the [Google](#) Cloud Manager Marketplace offering to continue using Cloud Backup. Billing for Cloud Backup is done through this subscription.
- For Cloud Backup BYOL licensing, you don't need a subscription. You need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses](#).

You need to have a Google subscription for the object storage space where your backups will be located.

A SnapMirror license is required on the cluster. Note that the "Hybrid Cloud Bundle" is not required when using Cloud Backup.

Preparing Google Cloud Storage for backups

When you set up backup, you need to provide storage access keys for a service account that has Storage Admin permissions. A service account enables Cloud Backup to authenticate and access Cloud Storage buckets used to store backups. 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 Backup later when you configure the backup service.

Enabling Cloud Backup

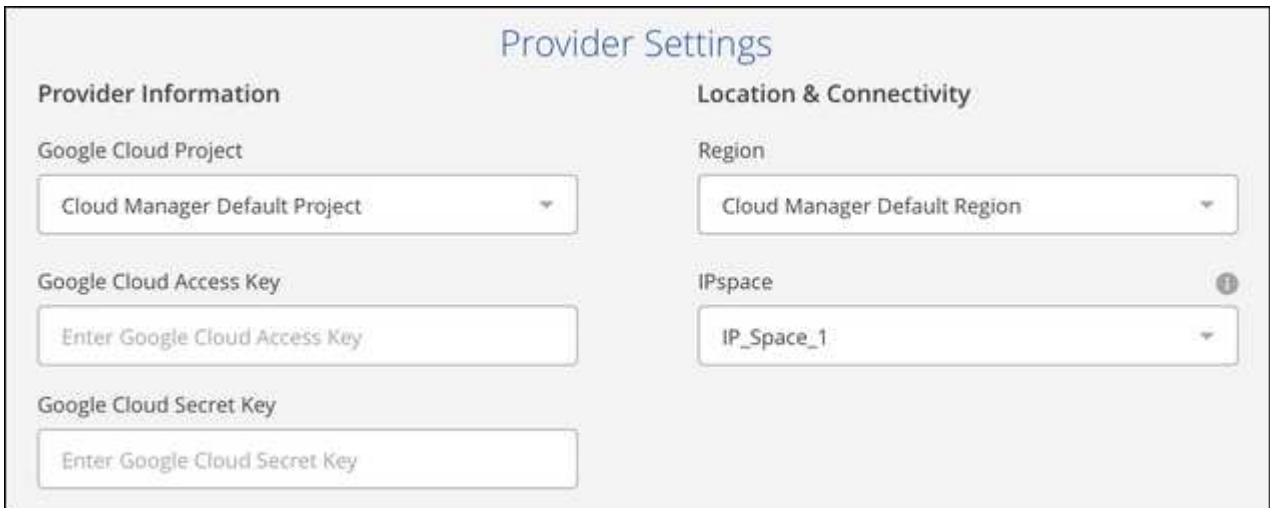
Enable Cloud Backup at any time directly from the on-premises working environment.

Steps

1. From the Canvas, select the working environment and click **Enable** next to the Backup & Compliance service in the right-panel.



2. Select Google Cloud as your provider and click **Next**.
3. Enter the provider details and click **Next**.
 - a. The Google Cloud Project where you want the Google Cloud Storage bucket to be created for backups. (The Project must have a Service Account that has the predefined Storage Admin role.)
 - b. The Google Access Key and Secret Key used to store the backups.
 - c. The Google region where the backups will be stored.
 - d. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.

A screenshot of a 'Provider Settings' form. The form is divided into two columns: 'Provider Information' and 'Location & Connectivity'. Under 'Provider Information', there are three fields: 'Google Cloud Project' (a dropdown menu with 'Cloud Manager Default Project' selected), 'Google Cloud Access Key' (a text input field with placeholder text 'Enter Google Cloud Access Key'), and 'Google Cloud Secret Key' (a text input field with placeholder text 'Enter Google Cloud Secret Key'). Under 'Location & Connectivity', there are two fields: 'Region' (a dropdown menu with 'Cloud Manager Default Region' selected) and 'IPspace' (a dropdown menu with 'IP_Space_1' selected). There is a small information icon to the right of the 'IPspace' field.

4. In the *Define Policy* page, select an existing backup schedule and retention value, or define a new backup policy, and click **Next**.

Define Policy

Policy - Retention & Schedule Create a New Policy Select an Existing Policy

Hourly Number of backups to retain

Daily Number of backups to retain

Weekly Number of backups to retain

Monthly 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

Google Cloud Storage Bucket Cloud Manager will create the Google Cloud Storage Bucket after you complete the wizard

See [the list of existing policies](#).

5. Select the volumes that you want to back up.

- To back up all volumes, check the box in the title row (Volume Name).
- To back up individual volumes, check the box for each volume (Volume_1).

Select Volumes

57 Volumes 🔍

<input checked="" type="checkbox"/>	Volume Name	Volume Type	SVM Name	Used Capacity	Allocated Capacity	Backup Status
<input checked="" type="checkbox"/>	Volume_Name_1	RW	SVM_Name_1	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_2	RW	SVM_Name_2	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_3	RW	SVM_Name_3	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_4	DP	SVM_Name_4	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_5	RW	SVM_Name_5	0.25 TB	10 TB	⊖ Not Active

6. Click **Activate Backup** and Cloud Backup starts taking the initial backups of your volumes.

Result

Cloud Backup starts taking the initial backups of each selected volume and the Volume Backup Dashboard is displayed so you can monitor the state of the backups.

What's next?

You can [start and stop backups for volumes or change the backup schedule](#) and you can [restore entire volumes from a backup file](#).

Backing up on-premises ONTAP data to StorageGRID

Complete a few steps to get started backing up data from your on-premises ONTAP systems to object storage in your NetApp StorageGRID systems.

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.
 - The cluster is running ONTAP 9.7P5 or later.
 - The cluster has a SnapMirror license — it is included as part of the Premium Bundle or Data Protection Bundle.
 - The cluster must have the required network connections to StorageGRID and to the Connector.
- You have a Connector installed on your premises.
 - Networking for the Connector enables an outbound HTTPS connection to the ONTAP cluster and to StorageGRID.
- You have purchased [and activated](#) a Cloud Backup BYOL license from NetApp.
- Your StorageGRID has version 10.3 or later with access keys that have S3 permissions.

2

Enable Cloud Backup on the system

Select the working environment and click **Enable** next to the Backup & Compliance service in the right-panel, and then follow the setup wizard.



3

Enter the StorageGRID details

Select StorageGRID as the provider, and then enter the StorageGRID server and service account details. You also need to specify the IPspace in the ONTAP cluster where the volumes reside.

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 hourly, daily, 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 you want to retain.

5

Select the volumes that you want to back up

Identify which volumes you want to back up from the cluster. An S3 bucket is created automatically in the service account indicated by the S3 access key and secret key you entered, and the backup files are stored there.

6

Restore your data, as needed

If necessary, choose the backup file to restore an entire backup to a new volume on an on-premises ONTAP system.

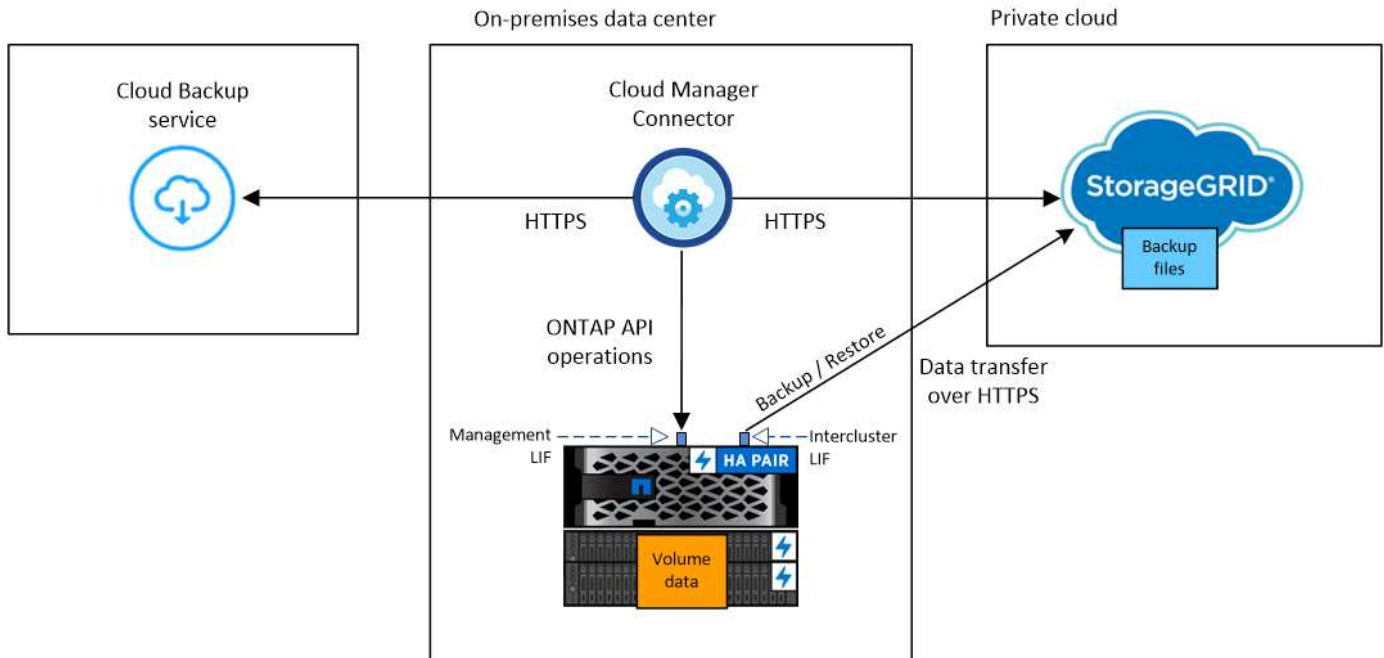
See [Restoring volume data from backup files](#) for details.

Requirements

Read the following requirements to make sure you have a supported configuration before you start backing up

on-premises volumes to StorageGRID.

The following image shows each component when backing up an on-prem ONTAP system to StorageGRID and the connections that you need to prepare between them:



Note that the Cloud Restore instance is not shown in this diagram because single-file restore is not currently supported when using StorageGRID.

Preparing your ONTAP clusters

You need to discover your on-premises ONTAP clusters in Cloud Manager before you can start backing up volume data.

[Learn how to discover a cluster.](#)

ONTAP requirements

- ONTAP 9.7P5 and later.
- A SnapMirror license (included as part of the Premium Bundle or Data Protection Bundle).

Note: The "Hybrid Cloud Bundle" is not required when using the Cloud Backup service.

See how to [manage your cluster licenses](#).

- Time and time zone are set correctly.

See how to [configure your cluster time](#).

Cluster networking requirements

- The ONTAP cluster initiates an HTTPS connection over a user-specified port from the intercluster LIF to StorageGRID for backup and restore operations. The port is configurable during backup setup.

ONTAP reads and writes data to and from object storage. The object storage never initiates, it just responds.

- ONTAP requires an inbound connection from the Connector to the cluster management LIF. The Connector must reside on your premises.
- 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. [Learn more about IPspaces.](#)

When you set up Cloud Backup, 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.

- The nodes' intercluster LIFs are able to access the object store.
- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM.](#)
- Note that if you use are using a different IPspace than the Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow Cloud Backup service connections from ONTAP to object storage through the port you specified (typically port 443) and name resolution traffic from the storage VM to the DNS server over port 53 (TCP/UDP).

Preparing StorageGRID

StorageGRID must meet the following requirements. See the [StorageGRID documentation](#) for more information.

Supported StorageGRID versions

StorageGRID 10.3 and later is supported.

S3 credentials

When you set up backup to StorageGRID, the backup wizard prompts you for an S3 access key and secret key for a service account. A service account enables Cloud Backup to authenticate and access the StorageGRID buckets used to store backups. The keys are required so that StorageGRID knows who is making the request.

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

When backing up 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](#)
- [Installing the Connector on a Linux host with internet access](#)
- [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 HTTPS connection over port 443 to StorageGRID
 - An HTTPS connection over port 443 to your ONTAP clusters
 - An outbound internet connection to the Cloud Backup service over port 443 (HTTPS)

License requirements

Before your 30-day free trial of the Cloud Backup service expires, you need to purchase and activate a Cloud Backup BYOL license from NetApp. This license is for the account and can be used across multiple systems.

You'll need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses.](#)



PAYGO licensing is not currently supported when backing up files to StorageGRID.

A SnapMirror license is required on the cluster. Note that the "Hybrid Cloud Bundle" is not required when using Cloud Backup.

Enabling Cloud Backup to StorageGRID

Enable Cloud Backup at any time directly from the on-premises working environment.

Steps

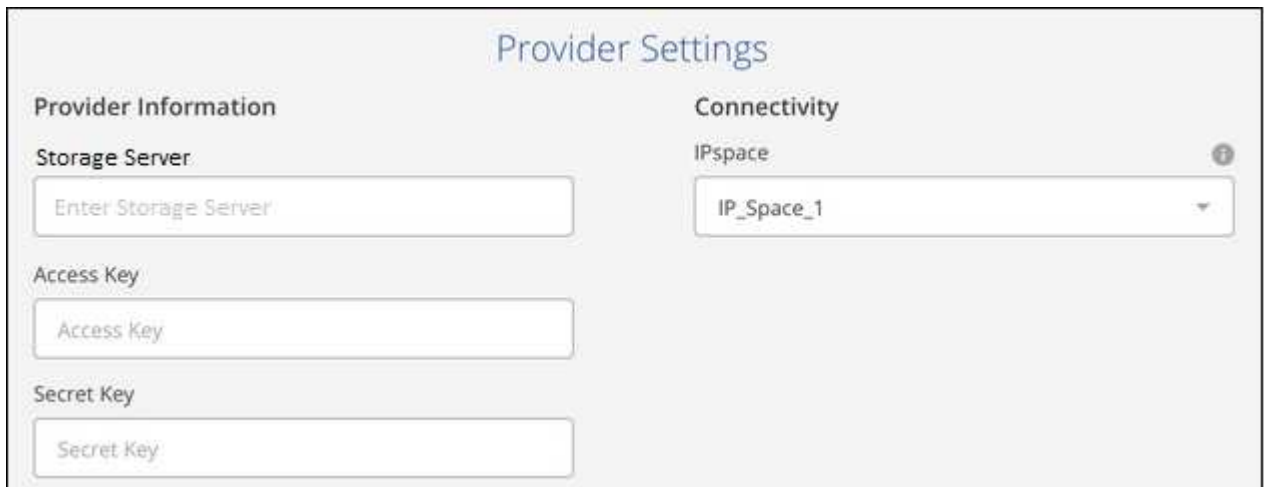
1. From the Canvas, select the on-premises working environment and click **Enable** next to the Backup & Compliance service in the right-panel.



2. Select **StorageGRID** as the provider, click **Next**, and then enter the provider details:
 - a. The FQDN of the StorageGRID server and the port that ONTAP should use for HTTPS communication with StorageGRID; for example: `s3.eng.company.com:8082`
 - b. The Access Key and the Secret Key used to access the bucket to store backups.
 - c. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.

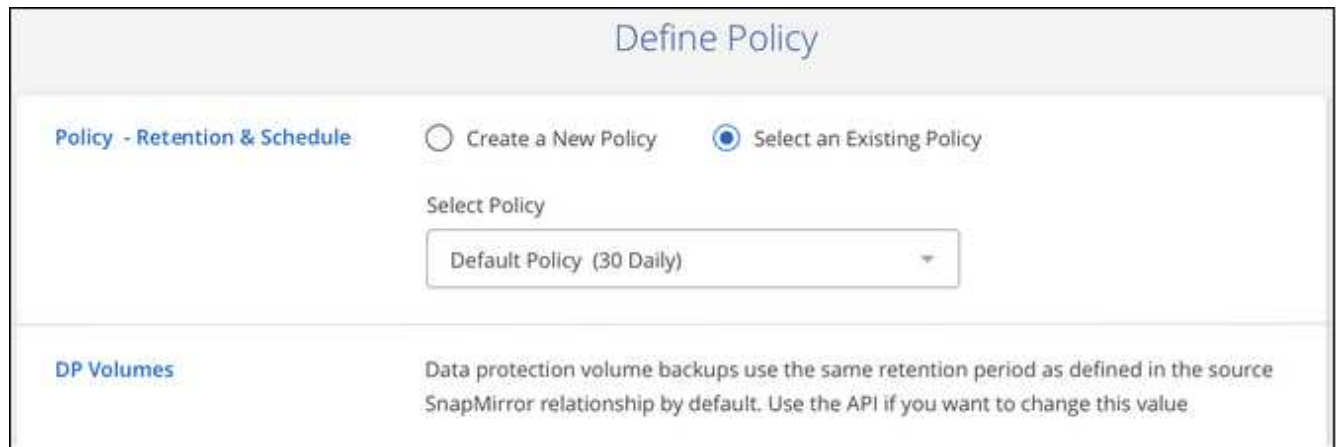
Selecting the correct IPspace ensures that Cloud Backup can set up a connection from ONTAP to your

StorageGRID object storage.



Note that you cannot change this information after the service has started.

3. In the *Define Policy* page, select the backup schedule and retention value and click **Next**.



See [the list of existing policies](#).

4. Select the volumes that you want to back up.
 - To back up all volumes, check the box in the title row (Volume Name).
 - To back up individual volumes, check the box for each volume (Volume_1).

Select Volumes

57 Volumes 🔍

<input checked="" type="checkbox"/>	Volume Name	Volume Type	SVM Name	Used Capacity	Allocated Capacity	Backup Status
<input checked="" type="checkbox"/>	Volume_Name_1	RW	SVM_Name_1	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_2	RW	SVM_Name_2	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_3	RW	SVM_Name_3	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_4	DP	SVM_Name_4	0.25 TB	10 TB	⊖ Not Active
<input checked="" type="checkbox"/>	Volume_Name_5	RW	SVM_Name_5	0.25 TB	10 TB	⊖ Not Active

5. Click **Activate Backup** and Cloud Backup starts taking the initial backups of each selected volume.

Result

An S3 bucket is created automatically in the service account indicated by the S3 access key and secret key you entered, and the backup files are stored there. The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

What's next?

You can [start and stop backups for volumes or change the backup schedule](#) and you can [restore entire volumes from a backup file](#).

Copyright Information

Copyright © 2022 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system-without prior written permission of the copyright owner.

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

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

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

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

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

Trademark Information

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