



Back up and restore ONTAP data

BlueXP backup and recovery

NetApp

September 12, 2024

This PDF was generated from <https://docs.netapp.com/us-en/bluexp-backup-recovery/concept-ontap-backup-to-cloud.html> on September 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Back up and restore ONTAP data 1
 - Protect your ONTAP volume data using BlueXP backup and recovery 1
 - Plan your protection journey 10
 - Manage backup policies for ONTAP volumes 17
 - Backup-to-object policy options 21
 - Manage backup-to-object storage options in the Advanced Settings page 31
 - Back up Cloud Volumes ONTAP data to Amazon S3 34
 - Back up Cloud Volumes ONTAP data to Azure Blob storage 46
 - Back up Cloud Volumes ONTAP data to Google Cloud Storage 56
 - Back up on-premises ONTAP data to Amazon S3 68
 - Back up on-premises ONTAP data to Azure Blob storage 84
 - Back up on-premises ONTAP data to Google Cloud Storage 96
 - Back up on-premises ONTAP data to ONTAP S3 108
 - Back up on-premises ONTAP data to StorageGRID 118
 - Manage backups for your ONTAP systems 128
 - Restore ONTAP data from backup files 147


Back up and restore ONTAP data

Protect your ONTAP volume data using BlueXP backup and recovery

The BlueXP backup and recovery service provides backup and restore capabilities for protection and long-term archive of your ONTAP volume data. You can implement a 3-2-1 strategy where you have 3 copies of your source data on 2 different storage systems along with 1 copy in the cloud.

After activation, backup and recovery creates block-level, incremental forever backups that are stored on another ONTAP cluster and in object storage in the cloud. In addition to your source volume, you'll have a:

- Snapshot copy of the volume on the source system
- Replicated volume on a different storage system
- Backup of the volume in object storage



BlueXP backup and recovery leverages NetApp's SnapMirror data replication technology to ensure that all the backups are fully synchronized by creating Snapshot copies and transferring them to the backup locations.

The benefits of the 3-2-1 approach include:

- Multiple data copies provide multi-layer protection against both internal (insider) and external cybersecurity threats.
- Multiple media types ensure failover viability in the case of physical or logical failure of one media type.
- The onsite copy facilitates quick restores, with the offsite copies at the ready just in case the onsite copy is compromised.

When necessary, you can restore an entire *volume*, a *folder*, or one or more *files*, from any of the backup copies to the same or different working environment.

Features

Replication features:

- Replicate data between ONTAP storage systems to support backup and disaster recovery.
- Ensure the reliability of your DR environment with high availability.
- Native ONTAP in-flight encryption set up via Pre-Shared Key (PSK) between the two systems.
- Copied data is immutable until you make it writable and ready to use.
- Replication is self-healing in the event of a transfer failure.
- When compared to the [BlueXP replication service](#), the replication in BlueXP backup and recovery includes the following features:
 - Replicate multiple FlexVol volumes at a time to a secondary system.
 - Restore a replicated volume to the source system or to a different system using the UI.
 - Manage replication policies

See [Replication limitations](#) for a list of replication features that are unavailable with BlueXP backup and recovery.

Backup-to-object features:

- Back up independent copies of your data volumes to low-cost object storage.
- Apply a single backup policy to all volumes in a cluster, or assign different backup policies to volumes that have unique recovery point objectives.
- Create a backup policy to be applied to all future volumes created in the cluster.
- Make immutable backup files so they are locked and protected for the retention period.
- Scan backup files for possible ransomware attack - and remove/replace infected backups automatically.
- Tier older backup files to archival storage to save costs.
- Delete the backup relationship so you can archive unneeded source volumes while retaining volume backups.
- Back up from cloud to cloud, and from on-premises systems to public or private cloud.
- Backup data is secured with AES-256 bit encryption at-rest and TLS 1.2 HTTPS connections in-flight.
- Use your own customer-managed keys for data encryption instead of using the default encryption keys from your cloud provider.
- Support for up to 4,000 backups of a single volume.

Restore features:

- Restore data from a specific point in time from local Snapshot copies, replicated volumes, or backed up volumes in object storage.
- Restore a volume, a folder, or individual files, to the source system or to a different system.
- Restore data to a working environment using a different subscription/account or that is in a different region.
- Perform a *quick restore* of a volume from cloud storage to a Cloud Volumes ONTAP system or to an on-premises system; perfect for disaster recovery situations where you need to provide access to a volume as soon as possible.
- Restore data on a block level, placing the data directly in the location you specify, all while preserving the original ACLs.
- Browse and search file catalogs for easy selection of individual folders and files for single file restore.

Supported working environments for backup and restore operations

BlueXP backup and recovery supports ONTAP working environments and public and private cloud providers.

Supported regions

BlueXP backup and recovery is supported with Cloud Volumes ONTAP in many Amazon Web Services, Microsoft Azure, and Google Cloud regions.

[Learn more using the Global Regions Map](#)

Supported backup destinations

BlueXP backup and recovery enables you to back up ONTAP volumes from the following source working environments to the following secondary working environments and object storage in public and private cloud providers. Snapshot copies reside on the source working environment.

Source Working Environment	Secondary Working Environment (Replication)	Destination Object Store (Backup)
Cloud Volumes ONTAP in AWS	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Amazon S3
Cloud Volumes ONTAP in Azure	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Azure Blob
Cloud Volumes ONTAP in Google	Cloud Volumes ONTAP in Google On-premises ONTAP system	Google Cloud Storage
On-premises ONTAP system	Cloud Volumes ONTAP On-premises ONTAP system	Amazon S3 Azure Blob Google Cloud Storage NetApp StorageGRID ONTAP S3

Supported restore destinations

You can restore ONTAP data from a backup file that resides in a secondary working environment (a replicated volume) or in object storage (a backup file) to the following working environments. Snapshot copies reside on the source working environment and can be restored only to that same system.

Backup File Location		Destination Working Environment
Object Store (Backup)	Secondary System (Replication)	
Amazon S3	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Cloud Volumes ONTAP in AWS On-premises ONTAP system
Azure Blob	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Cloud Volumes ONTAP in Azure On-premises ONTAP system

Backup File Location		Destination Working Environment
Google Cloud Storage	Cloud Volumes ONTAP in Google On-premises ONTAP system	Cloud Volumes ONTAP in Google On-premises ONTAP system
NetApp StorageGRID	On-premises ONTAP system Cloud Volumes ONTAP	On-premises ONTAP system
ONTAP S3	On-premises ONTAP system Cloud Volumes ONTAP	On-premises ONTAP system

Note that references to "on-premises ONTAP systems" includes FAS, AFF, and ONTAP Select systems.

Supported volumes

BlueXP backup and recovery supports the following types of volumes:

- FlexVol read-write volumes
- FlexGroup volumes (requires ONTAP 9.12.1 or later)
- SnapLock Enterprise volumes (requires ONTAP 9.11.1 or later)
- SnapMirror data protection (DP) destination volumes

See the sections on [Backup and Restore limitations](#) for additional requirements and limitations.

Cost

There are two types of costs associated with using BlueXP backup and recovery with ONTAP systems: resource charges and service charges. Both of these charges are for the backup to object portion of the service.

There is no charge to create Snapshot copies or replicated volumes - other than the disk space required to store the Snapshot copies and replicated volumes.

Resource charges

Resource charges are paid to the cloud provider for object storage capacity and for writing and reading backup files to the cloud.

- For Backup to object storage, you pay your cloud provider for object storage costs.

Since BlueXP backup and recovery preserves the storage efficiencies of the source volume, you pay the cloud provider object storage costs for the data *after* ONTAP efficiencies (for the smaller amount of data after deduplication and compression have been applied).

- For restoring data using Search & Restore, certain resources are provisioned by your cloud provider, and there is per-TiB cost associated with the amount of data that is scanned by your search requests. (These resources are not needed for Browse & Restore.)
 - In AWS, [Amazon Athena](#) and [AWS Glue](#) resources are deployed in a new S3 bucket.
 - In Azure, an [Azure Synapse workspace](#) and [Azure Data Lake Storage](#) are provisioned in your storage account to store and analyze your data.

- In Google, a new bucket is deployed, and the [Google Cloud BigQuery services](#) are provisioned on an account/project level.
- If you plan to restore volume data from a backup file that has been moved to archival object storage, then there's an additional per-GiB retrieval fee and per-request fee from the cloud provider.
- If you plan to scan a backup file for ransomware during the process of restoring volume data (if you have enabled DataLock and Ransomware Protection for your cloud backups), then you'll incur extra egress costs from your cloud provider as well.

Service charges

Service charges are paid to NetApp and cover both the cost to *create* backups to object storage and to *restore* volumes, or files, from those backups. You pay only for the data that you protect in object storage, calculated by the source logical used capacity (*before* ONTAP efficiencies) of ONTAP volumes which are backed up to object storage. This capacity is also known as Front-End Terabytes (FETB).

There are three ways to pay for the Backup service. The first option is to subscribe from your cloud provider, which enables you to pay per month. The second option is to get an annual contract. The third option is to purchase licenses directly from NetApp. Read the [Licensing](#) section for details.

Licensing

BlueXP backup and recovery is available with the following consumption models:

- **BYOL:** A license purchased from NetApp that can be used with any cloud provider.
- **PAYGO:** An hourly subscription from your cloud provider's marketplace.
- **Annual:** An annual contract from your cloud provider's marketplace.

A Backup license is required only for backup and restore from object storage. Creating Snapshot copies and replicated volumes do not require a license.

Bring your own license

BYOL is term-based (1, 2, or 3 years) *and* capacity-based in 1 TiB increments. You pay NetApp to use the service for a period of time, say 1 year, and for a maximum amount capacity, say 10 TiB.

You'll receive a serial number that you enter in the BlueXP digital wallet page to enable the service. When either limit is reached, you'll need to renew the license. The Backup BYOL license applies to all source systems associated with your [BlueXP account](#).

[Learn how to manage your BYOL licenses.](#)

Pay-as-you-go subscription

BlueXP backup and recovery offers consumption-based licensing in a pay-as-you-go model. After subscribing through your cloud provider's marketplace, you pay per GiB for data that's backed up — there's no up-front payment. You are billed by your cloud provider through your monthly bill.

[Learn how to set up a pay-as-you-go subscription.](#)

Note that a 30-day free trial is available when you initially sign up with a PAYGO subscription.

Annual contract

When you use AWS, two annual contracts are available for 1, 2, or 3 year terms:

- A "Cloud Backup" plan that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP data.
- A "CVO Professional" plan that enables you to bundle Cloud Volumes ONTAP and BlueXP backup and recovery. This includes unlimited backups for Cloud Volumes ONTAP volumes charged against this license (backup capacity is not counted against the license).

When you use Azure, two annual contracts are available for 1, 2, or 3 year terms:

- A "Cloud Backup" plan that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP data.
- A "CVO Professional" plan that enables you to bundle Cloud Volumes ONTAP and BlueXP backup and recovery. This includes unlimited backups for Cloud Volumes ONTAP volumes charged against this license (backup capacity is not counted against the license).

When you use GCP, you can request a private offer from NetApp, and then select the plan when you subscribe from the Google Cloud Marketplace during BlueXP backup and recovery activation.

[Learn how to set up annual contracts.](#)


How BlueXP backup and recovery works

When you enable BlueXP backup and recovery on a Cloud Volumes ONTAP or on-premises ONTAP system, the service performs a full backup of your data. After the initial backup, all additional backups are incremental, which means that only changed blocks and new blocks are backed up. This keeps network traffic to a minimum. Backup to object storage is built on top of the [NetApp SnapMirror Cloud technology](#).



Any actions taken directly from your cloud provider environment to manage or change cloud backup files may corrupt the files and will result in an unsupported configuration.

The following image shows the relationship between each component:



This diagram shows volumes being replicated to a Cloud Volumes ONTAP system, but volumes could be replicated to an on-premises ONTAP system as well.

Where backups reside

Backups reside in different locations based on the type of backup:

- *Snapshot copies* reside on the source volume in the source working environment.
- *Replicated volumes* reside on the secondary storage system - a Cloud Volumes ONTAP or on-premises ONTAP system.
- *Backup copies* are stored in an object store that BlueXP creates in your cloud account. There's one object store per cluster/working environment, and BlueXP names the object store as follows: "netapp-backup-clusteruuiid". Be sure not to delete this object store.
 - In AWS, BlueXP enables the [Amazon S3 Block Public Access feature](#) on the S3 bucket.
 - In Azure, BlueXP uses a new or existing resource group with a storage account for the Blob container. BlueXP [blocks public access to your blob data](#) by default.
 - In GCP, BlueXP uses a new or existing project with a storage account for the Google Cloud Storage bucket.
 - In StorageGRID, BlueXP uses an existing tenant account for the S3 bucket.
 - In ONTAP S3, BlueXP uses an existing user account for the S3 bucket.

If you want to change the destination object store for a cluster in the future, you'll need to [unregister BlueXP backup and recovery for the working environment](#), and then enable BlueXP backup and recovery using the new cloud provider information.

Customizable backup schedule and retention settings

When you enable BlueXP backup and recovery for a working environment, all the volumes you initially select are backed up using the policies that you select. You can select separate policies for Snapshot copies, replicated volumes, and backup files. If you want to assign different backup policies to certain volumes that have different recovery point objectives (RPO), you can create additional policies for that cluster and assign those policies to the other volumes after BlueXP backup and recovery is activated.

You can choose a combination of hourly, daily, weekly, monthly, and yearly backups of all volumes. For backup to object you can also select one of the system-defined policies that provide backups and retention for 3 months, 1 year, and 7 years. Backup protection policies that you have created on the cluster using ONTAP System Manager or the ONTAP CLI will also appear as selections. This includes policies created using custom SnapMirror labels.



The Snapshot policy applied to the volume must have one of the labels that you're using in your replication policy and backup to object policy. If matching labels are not found, no backup files will be created. For example, if you want to create "weekly" replicated volumes and backup files, you must use a Snapshot policy that creates "weekly" Snapshot copies.

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 (and so obsolete backups don't continue to take up space).

See [Backup schedules](#) for more details about how the available schedule options.

Note that you can [create an on-demand backup of a volume](#) from the Backup Dashboard at any time, in addition to those backup files created from the scheduled backups.



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.

Backup file protection settings

If your cluster is using ONTAP 9.11.1 or greater, you can protect your backups in object storage from deletion and ransomware attacks. Each backup policy provides a section for *DataLock and Ransomware Protection* that can be applied to your backup files for a specific period of time - the *retention period*.

- *DataLock* protects your backup files from being modified or deleted.
- *Ransomware protection* scans your backup files to look for evidence of a ransomware attack when a backup file is created, and when data from a backup file is being restored.

Scheduled ransomware protection scans are enabled by default. The default setting for the scan frequency is for 7 days. The scan occurs only on the latest Snapshot copy. The scheduled scans can be disabled to reduce your costs. You can enable or disable scheduled ransomware scans on the latest Snapshot copy by using the option on the Advanced Settings page. If you enable it, scans are performed weekly by default. You can change that schedule to days or weeks or disable it, saving costs.

The backup retention period is the same as the backup schedule retention period, plus a maximum 31-day buffer. For example, *weekly* backups with 5 copies retained will lock each backup file for 5 weeks. *Monthly* backups with 6 copies retained will lock each backup file for 6 months.

Support is currently available when your backup destination is Amazon S3, Azure Blob, or NetApp StorageGRID. Other storage provider destinations will be added in future releases.

For more details, refer to this information:

- [How DataLock and Ransomware protection work.](#)
- [How to update Ransomware protection options in the Advanced Settings page.](#)



DataLock can't be enabled if you are tiering backups to archival storage.

Archival storage for older backup files

When using certain cloud storage you can move older backup files to a less expensive storage class/access tier after a certain number of days. You can also choose to send your backup files to archival storage immediately without being written to standard cloud storage. Note that archival storage can't be used if you have enabled DataLock.

- In AWS, backups start in the *Standard* storage class and transition to the *Standard-Infrequent Access* storage class after 30 days.

If your cluster is using ONTAP 9.10.1 or greater, you can choose to tier older backups to either *S3 Glacier* or *S3 Glacier Deep Archive* storage in the BlueXP backup and recovery UI after a certain number of days for further cost optimization. [Learn more about AWS archival storage.](#)

- In Azure, backups are associated with the *Cool* access tier.

If your cluster is using ONTAP 9.10.1 or greater, you can choose to tier older backups to *Azure Archive* storage in the BlueXP backup and recovery UI after a certain number of days for further cost optimization. [Learn more about Azure archival storage.](#)

- In GCP, backups are associated with the *Standard* storage class.

If your cluster is using ONTAP 9.12.1 or greater, you can choose to tier older backups to *Archive* storage in the BlueXP backup and recovery UI after a certain number of days for further cost optimization. [Learn more about Google archival storage.](#)

- In StorageGRID, backups are associated with the *Standard* storage class.

If your on-prem cluster is using ONTAP 9.12.1 or greater, and your StorageGRID system is using 11.4 or greater, you can archive older backup files to public cloud archival storage after a certain number of days. Current support is for AWS S3 Glacier/S3 Glacier Deep Archive or Azure Archive storage tiers. [Learn more about archiving backup files from StorageGRID.](#)

See [Archival storage settings](#) for more details about archiving older backup files.

FabricPool tiering policy considerations

There are certain things you need to be aware of when the volume you are backing up resides on a FabricPool aggregate and it has an assigned tiering policy other than `none`:

- The first backup of a FabricPool-tiered volume requires reading all local and all tiered data (from the object store). A backup operation does not "reheat" the cold data tiered in object storage.

This operation could cause a one-time increase in cost to read the data from your cloud provider.

- Subsequent backups are incremental and do not have this effect.

- If the tiering policy is assigned to the volume when it is initially created you will not see this issue.
- Consider the impact of backups before assigning the all tiering policy to volumes. Because data is tiered immediately, BlueXP backup and recovery will read data from the cloud tier rather than from the local tier. Because concurrent backup operations share the network link to the cloud object store, performance degradation might occur if network resources become saturated. In this case, you may want to proactively configure multiple network interfaces (LIFs) to decrease this type of network saturation.

Plan your protection journey

The BlueXP backup and recovery service enables you to create up to three copies of your source volumes to protect your data. There are many options that you can select when enabling this service on your volumes, so you should review your choices so you're prepared.

We'll go over the following options:

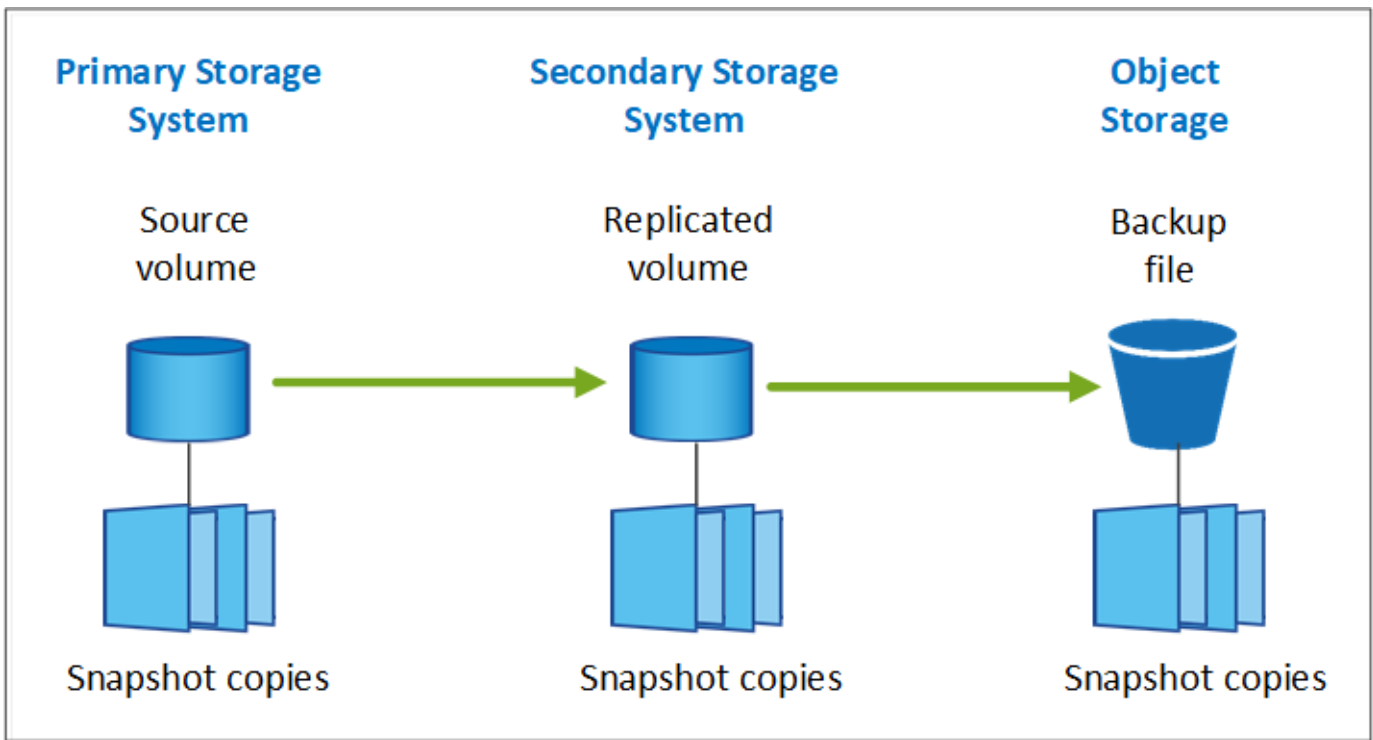
- Which protection features will you use: Snapshot copies, replicated volumes, and/or backup to cloud
- Which backup architecture will you use: a cascade or fan-out backup of your volumes
- Will you use the default backup policies, or do you need to create custom policies
- Do you want the service to create the cloud buckets for you, or do you want to make your object storage containers before you begin
- Which BlueXP Connector deployment mode are you using (standard, restricted, or private mode)

Which protection features will you use

Before you select the features you'll use, here's a quick explanation of what each features does, and what type of protection it provides.

Backup type	Description
Snapshot	Creates a read-only, point-in-time image of a volume within the source volume as a Snapshot copy. You can use the Snapshot copy to recover individual files, or to restore the entire contents of a volume.
Replication	Creates a secondary copy of your data on another ONTAP storage system and continually updates the secondary data. Your data is kept current and remains available whenever you need it.
Cloud backup	Creates backups of your data to the cloud for protection and for long-term archival purposes. If necessary, you can restore a volume, folder, or individual files from the backup to the same, or different, working environment.

Snapshots are the basis of all the backup methods, and they are required to use the backup and recovery service. A Snapshot copy is a read-only, point-in-time image of a volume. The image consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last Snapshot copy was made. The Snapshot copy that is created on your volume is used to keep the replicated volume and backup file synchronized with changes made to the source volume - as shown in the figure.



You can choose to create both replicated volumes on another ONTAP storage system and backup files in the cloud. Or you can choose just to create replicated volumes or backup files - it's your choice.

To summarize, these are the valid protection flows you can create for volumes in your ONTAP working environment:

- Source volume → Snapshot copy → Replicated volume → Backup file
- Source volume → Snapshot copy → Backup file
- Source volume → Snapshot copy → Replicated volume



The initial creation of a replicated volume or backup file includes a full copy of the source data — this is called a *baseline transfer*. Subsequent transfers contain only differential copies of the source data (the Snapshot).

Comparison of the different backup methods

The following table shows a generalized comparison of the three backup methods. While object storage space is typically less expensive than your on-premises disk storage, if you think you might restore data from the cloud frequently, then the egress fees from cloud providers can reduce some of your savings. You'll need to identify how often you need to restore data from the backup files in the cloud.

In addition to this criteria, cloud storage offers additional security options if you use the DataLock and Ransomware Protection feature, and additional cost savings by selecting archival storage classes for older backup files. [Learn more about DataLock and Ransomware protection](#) and [archival storage settings](#).

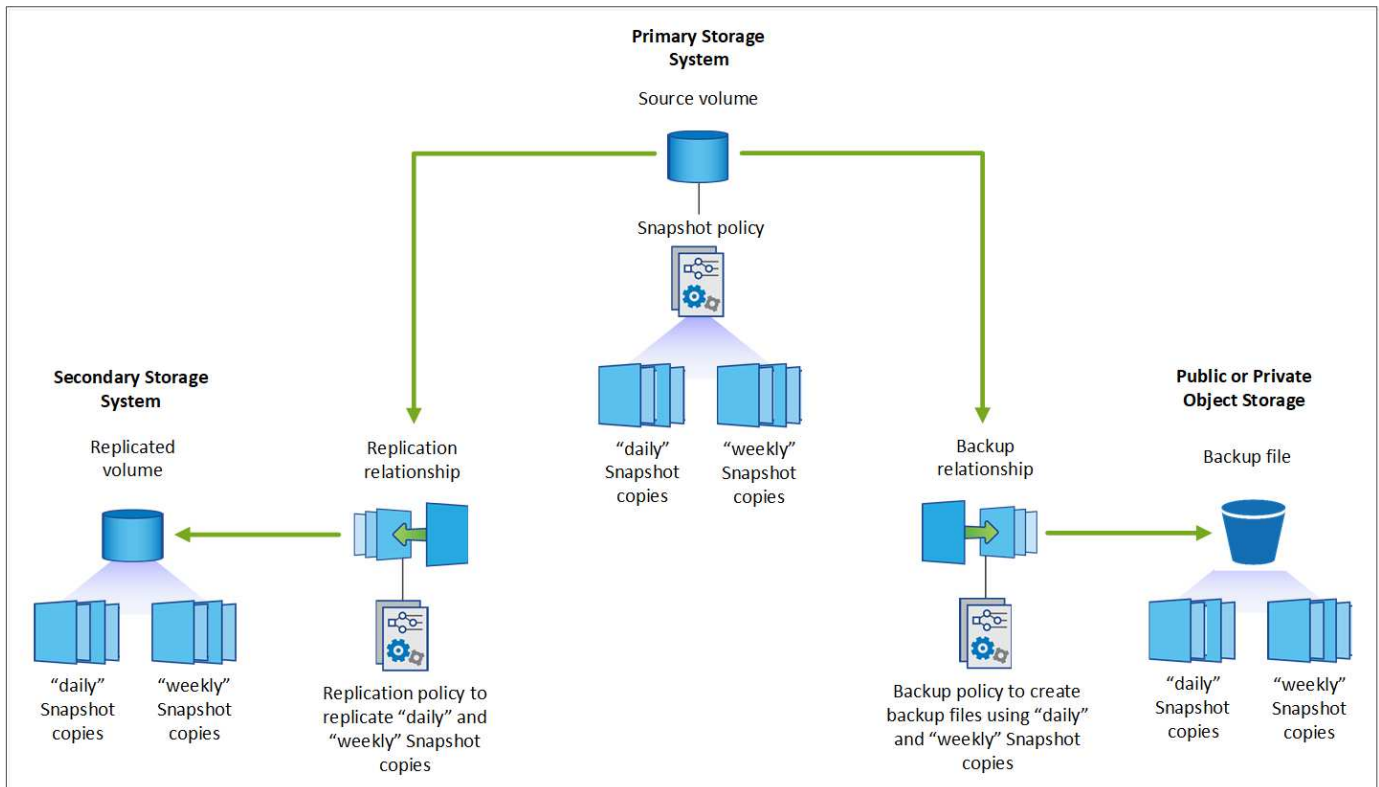
Backup type	Backup speed	Backup cost	Restore speed	Restore cost
Snapshot	High	Low (disk space)	High	Low
Replication	Medium	Medium (disk space)	Medium	Medium (network)

Backup type	Backup speed	Backup cost	Restore speed	Restore cost
Cloud backup	Low	Low (object space)	Low	High (provider fees)


Which backup architecture will you use

When creating both replicated volumes and backup files, you can choose a fan-out or cascade architecture to back up your volumes.

A **fan-out** architecture transfers the Snapshot copy independently to both the destination storage system and the backup object in the cloud.



A **cascade** architecture transfers the Snapshot copy to the destination storage system first, and then that system transfers the copy to the backup object in the cloud.



Comparison of the different architecture choices

This table provides a comparison of the fan-out and cascade architectures.

Fan-out	Cascade
Small performance impact on the source system because it is sending Snapshot copies to 2 distinct systems	Less effect on the performance of the source storage system because it sends the Snapshot copy only once
Easier to set up because all policies, networking, and ONTAP configurations are done on the source system	Requires some networking and ONTAP configuration to be done from the secondary system as well.

Will you use the default policies for Snapshot copies, replications, and backups

You can use the default policies provided by NetApp to create your backups, or you can create custom policies. When you use the activation wizard to enable the backup and recovery service for your volumes, you can select from the default policies and any other policies that already exist in the working environment (Cloud Volumes ONTAP or on-premises ONTAP system). If you want to use a policy different than those existing policies, you can create the policy before starting or while using the activation wizard.

- The default Snapshot policy creates hourly, daily, and weekly Snapshot copies, retaining 6 hourly, 2 daily, and 2 weekly Snapshot copies.
- The default replication policy replicates daily and weekly Snapshot copies, retaining 7 daily and 52 weekly Snapshot copies.
- The default backup policy replicates daily and weekly Snapshot copies, retaining 7 daily and 52 weekly Snapshot copies.

If you create custom policies for replication or backup, the policy labels (for example, "daily" or "weekly") must match the labels that exist in your Snapshot policies or replicated volumes and backup files won't be created.

You can create Snapshot, replication, and backup to object storage policies in the BlueXP backup and recovery UI. See the section for [adding a new backup policy](#) for details.

In addition to using BlueXP backup recovery to create custom policies, you can use System Manager or the ONTAP Command Line Interface (CLI).

[Create a Snapshot policy using System Manager](#)

[Create a Snapshot policy using the ONTAP CLI](#)

[Create a replication policy using System Manager](#)

[Create a replication policy using the ONTAP CLI](#)

[Create a backup policy using System Manager](#)

[Create a backup policy using the ONTAP CLI](#)

Note: When using System Manager, select **Asynchronous** as the policy type for replication policies, and select **Asynchronous** and **Back up to cloud** for backup to object policies.

Here are a few sample ONTAP CLI commands that may be helpful if you are creating custom policies. Note that you must use the *admin* vserver (storage VM) as the <vserver_name> in these commands.

Policy Description	Command
Simple Snapshot policy	<code>snapshot policy create -policy WeeklySnapshotPolicy -enabled true -schedule1 weekly -count1 10 -vserver ClusterA -snapmirror-label1 weekly</code>
Simple backup to cloud	<code>snapmirror policy create -policy <policy_name> -transfer -priority normal -vserver <vserver_name> -create -snapshot-on-source false -type vault snapmirror policy add-rule -policy <policy_name> -vserver <vserver_name> -snapmirror-label <snapmirror_label> -keep</code>
Backup to cloud with DataLock and Ransomware protection	<code>snapmirror policy create -policy CloudBackupService-Enterprise -snapshot-lock-mode enterprise -vserver <vserver_name> snapmirror policy add-rule -policy CloudBackupService-Enterprise -retention-period 30days</code>
Backup to cloud with archival storage class	<code>snapmirror policy create -vserver <vserver_name> -policy <policy_name> -archive-after-days <days> -create -snapshot-on-source false -type vault snapmirror policy add-rule -policy <policy_name> -vserver <vserver_name> -snapmirror-label <snapmirror_label> -keep</code>
Simple replication to another storage system	<code>snapmirror policy create -policy <policy_name> -type async-mirror -vserver <vserver_name> snapmirror policy add-rule -policy <policy_name> -vserver <vserver_name> -snapmirror-label <snapmirror_label> -keep</code>



Only vault policies can be used for backup to cloud relationships.

Where do my policies reside?

Backup policies reside in different locations depending on the backup architecture you plan to use: Fan-out or Cascading. Replication policies and Backup policies are not designed the same way because replications pair two ONTAP storage systems and backup to object uses a storage provider as the destination.

- Snapshot policies always reside on the primary storage system.
- Replication policies always reside on the secondary storage system.
- Backup-to-object policies are created on the system where the source volume resides - this is the primary cluster for fan-out configurations, and the secondary cluster for cascading configurations.

These differences are shown in the table.

Architecture	Snapshot policy	Replication policy	Backup policy
Fan-out	Primary	Secondary	Primary
Cascade	Primary	Secondary	Secondary

So if you're planning to create custom policies when using the cascading architecture, you'll need to create the

replication and backup to object policies on the secondary system where the replicated volumes will be created. If you're planning to create custom policies when using the fan-out architecture, you'll need to create the replication policies on the secondary system where the replicated volumes will be created and backup to object policies on the primary system.

If you're using the default policies that exist on all ONTAP systems, then you're all set.

Do you want to create your own object storage container

When you create backup files in object storage for a working environment, by default, the backup and recovery service creates the container (bucket or storage account) for the backup files in the object storage account that you have configured. The AWS or GCP bucket is named "netapp-backup-<uuid>" by default. The Azure Blob storage account is named "netappbackup<uuid>".

You can create the container yourself in the object provider account if you want to use a certain prefix or assign special properties. If you want to create your own container, you must create it before starting the activation wizard. The container must be used exclusively for storing ONTAP volume backup files - it cannot be used for any other purpose. The backup activation wizard will automatically discover your provisioned containers for the selected Account and credentials so that you can select the one you want to use.

You can create the bucket from BlueXP, or from your cloud provider.

- [Create Amazon S3 buckets from BlueXP](#)
- [Create Azure Blob storage accounts from BlueXP](#)
- [Create Google Cloud Storage buckets from BlueXP](#)

Note: At this time you cannot use your own S3 buckets when creating backups in StorageGRID systems or to ONTAP S3.

If you plan to use a different bucket prefix than "netapp-backup-xxxxxx", then you'll need to modify the S3 permissions for the Connector IAM Role. For details, refer to [how to create backups to AWS S3](#).

Advanced bucket settings

If you plan to move older backup files to archival storage, or if you plan to enable DataLock and Ransomware protection to lock your backup files and scan them for possible ransomware, you'll need to create the container with certain configuration settings:

- Archival storage on your own buckets is supported in AWS S3 storage at this time when using ONTAP 9.10.1 or greater software on your clusters. By default, backups start in the S3 *Standard* storage class. Ensure that you create the bucket with the appropriate lifecycle rules:
 - Move the objects in the entire scope of the bucket to S3 *Standard-IA* after 30 days.
 - Move the objects with the tag "smc_push_to_archive: true" to *Glacier Flexible Retrieval* (formerly S3 Glacier)
- DataLock and Ransomware protection is supported in AWS storage when using ONTAP 9.11.1 or greater software on your clusters, and Azure storage when using ONTAP 9.12.1 or greater software.
 - For AWS, you must enable Object Locking on the bucket using a 30-day retention period.
 - For Azure, you need to create the Storage Class with version-level immutability support.

Which BlueXP Connector deployment mode are you using

If you're already using BlueXP to manage your storage, then a BlueXP Connector has already been installed. If you plan to use the same Connector with BlueXP backup and recovery, then you're all set. If you need to use a different Connector, you'll need to install it before starting your backup and recovery implementation.

BlueXP offers multiple deployment modes that enable you to use BlueXP in a way that meets your business and security requirements. *Standard mode* leverages the BlueXP SaaS layer to provide full functionality, while *restricted mode* and *private mode* are available for organizations that have connectivity restrictions.

[Learn more about BlueXP deployment modes.](#)

Support for sites with full internet connectivity

When BlueXP backup and recovery is used in a site with full internet connectivity (also known as *standard mode* or *SaaS mode*), you can create replicated volumes on any on-premises ONTAP or Cloud Volumes ONTAP systems managed by BlueXP, and you can create backup files on object storage in any of the supported cloud providers. [See the full list of supported backup destinations.](#)

For a list of valid Connector locations, refer to one of the following backup procedures for the cloud provider where you plan to create backup files. There are some restrictions where the Connector must be installed manually on a Linux machine or deployed in a specific cloud provider.

- [Back up Cloud Volumes ONTAP data to Amazon S3](#)
- [Back up on-premises ONTAP data to Amazon S3](#)
- [Back up Cloud Volumes ONTAP data to Azure Blob](#)
- [Back up on-premises ONTAP data to Azure Blob](#)
- [Back up Cloud Volumes ONTAP data to Google Cloud](#)
- [Back up on-premises ONTAP data to Google Cloud](#)
- [Back up on-premises ONTAP data to StorageGRID](#)
- [Back up on-premises ONTAP to ONTAP S3](#)

Support for sites with limited internet connectivity

BlueXP backup and recovery can be used in a site with limited internet connectivity (also known as *restricted mode*) to back up volume data. In this case, you'll need to deploy the BlueXP Connector in the restricted region.

- You can back up data from Cloud Volumes ONTAP systems installed in AWS commercial regions to Amazon S3. [Back up Cloud Volumes ONTAP data to Amazon S3.](#)
- You can back up data from Cloud Volumes ONTAP systems installed in Azure commercial regions to Azure Blob. [Back up Cloud Volumes ONTAP data to Azure Blob.](#)

Support for sites with no internet connectivity

BlueXP backup and recovery can be used in a site with no internet connectivity (also known as *private mode* or *dark sites*) to back up volume data. In this case, you'll need to deploy the BlueXP Connector on a Linux host in the same site.

- You can back up data from local on-premises ONTAP systems to local NetApp StorageGRID systems. [Back up on-premises ONTAP data to StorageGRID.](#)

- You can back up data from local on-premises ONTAP systems to local on-premises ONTAP systems or Cloud Volumes ONTAP systems configured for S3 object storage. [Back up on-premises ONTAP data to ONTAP S3](#).

Manage backup policies for ONTAP volumes

You can use the default backup policies provided by NetApp to create your backups, or you can create custom policies. Policies govern the backup frequency, the time the backup is taken, and the number of backup files that are retained.

When you use the activation wizard to enable the backup and recovery service for your volumes, you can select from the default policies and any other policies that already exist in the working environment (Cloud Volumes ONTAP or on-premises ONTAP system). If you want to use a policy different than those existing policies, you can create the policy before or while you use the activation wizard.

To learn about the default backup policies provided, refer to [Plan your protection journey](#).

BlueXP backup and recovery provides three types of backups of ONTAP data: Snapshots, replications, and backups to object storage. Their policies reside in different locations based on the architecture that you use and the type of backup:

Architecture	Snapshot policy storage location	Replication policy storage location	Backup to object policy storage location
Fan-out	Primary	Secondary	Primary
Cascade	Primary	Secondary	Secondary

Create backup policies using the following tools depending on your environment, your preferences, and the protection type:

- BlueXP UI
- System Manager UI
- ONTAP CLI



When using System Manager, select **Asynchronous** as the policy type for replication policies, and select **Asynchronous** and **Back up to cloud** for backup to object policies.

View policies for a working environment

1. In the BlueXP UI, select **Volumes > Backup settings**.
2. From the Backup Settings page, select the working environment, select the **Actions** icon, and select **Policies management**.

The Policies management page appears.

Backup and recovery **Volumes** Restore Applications Virtual Machines Kubernetes Job Monitoring Reports

Volumes > Backup Settings > Policies Management


Working Environment: PrimaryClusterA

31
Total Policies

4
Snapshot Policies

20
Replication Policies

7
Backup Policies

Snapshot Policies (4) Replication Policies (20) Backup Policies (7) 

Snapshot policy name	Schedule name	Associated Volumes
hourly	Hourly Daily Weekly	1
default	Hourly Daily Weekly	1
default-1weekly	Hourly Daily Weekly	0

Snapshot policies are displayed by default.

- To view other policies that exist in the working environment, select either **Replication Policies** or **Backup Policies**. If the existing policies can be used for your backup plans, you're all set. If you need to have a policy with different characteristics, you can create new policies from this page.

Create policies

You can create policies that govern your Snapshot copies, replications and backups to object storage:


- [Create a Snapshot policy before initiating the Snapshot](#)
- [Create a replication policy before initiating the replication](#)
- [Create a backup-to-object-storage policy before initiating the backup](#)

Create a Snapshot policy before initiating the Snapshot

Part of your 3-2-1 strategy involves creating a Snapshot copy of the volume on the **primary** storage system.

Part of the policy creation process involves identifying Snapshot and SnapMirror labels that denote the schedule and retention. You can use predefined labels or create your own.

Steps

- In the BlueXP UI, select **Volumes > Backup settings**.
- From the Backup Settings page, select the working environment, select the **Actions**  icon, and select **Policies management**.

The Policies management page appears.

- In the Policies page, select **Create policy > Create Snapshot policy**.
- Specify the policy name.
- Select the Snapshot schedule or schedules. You can have a maximum of 5 labels. Or, create a schedule.

6. If you choose to create a schedule:
 - a. Select the frequency of hourly, daily, weekly, monthly, or yearly.
 - b. Specify the Snapshot labels denoting the schedule and retention.
 - c. Enter when and how often the Snapshot will be taken.
 - d. Retention: Enter the number of Snapshots to keep.
7. Select **Create**.

Snapshot policy example using cascading architecture

This example creates a Snapshot policy with two clusters:

1. Cluster 1:
 - a. Select Cluster 1 on the policy page.
 - b. Ignore the Replication and Backup to Object policy sections.
 - c. Create the Snapshot policy.
2. Cluster 2:
 - a. Select Cluster 2 on the Policy page.
 - b. Ignore the Snapshot policy section.
 - c. Configure the Replication and Backup to object policies.

Create a replication policy before initiating the replication

Your 3-2-1 strategy might include replicating a volume on a different storage system. The replication policy resides on the **secondary** storage system.

Steps

1. In the Policies page, select **Create policy > Create replication policy**.
2. In the Policy Details section, specify the policy name.
3. Specify the SnapMirror labels (maximum of 5) denoting the retention for each label.
4. Specify the transfer schedule.
5. Select **Create**.

Create a backup-to-object-storage policy before initiating the backup

Your 3-2-1 strategy might include backing up a volume to object storage.

This storage policy resides in different storage system locations depending on the backup architecture:

- Fan-out: Primary storage system
- Cascading: Secondary storage system

Steps

1. In the Policy management page, select **Create policy > Create backup policy**.
2. In the Policy Details section, specify the policy name.
3. Specify the SnapMirror labels (maximum of 5) denoting the retention for each label.

4. Specify the settings, including the transfer schedule and when to archive backups.
5. (Optional) To move older backup files to a less expensive storage class or access tier after a certain number of days, select the **Archive** option and indicate the number of days that should elapse before the data is archived. Enter **0** as the "Archive After Days" to send your backup file directly to archival storage.

[Learn more about archival storage settings.](#)

6. (Optional) To protect your backups from being modified or deleted, select the **DataLock & Ransomware protection** option.

If your cluster is using ONTAP 9.11.1 or greater, you can choose to protect your backups from deletion by configuring *DataLock* and *Ransomware protection*.

[Learn more about the available DataLock settings.](#)

7. Select **Create**.

Edit a policy

You can edit a custom Snapshot, replication, or backup policy.

Changing the backup policy affects all volumes that are using that policy.

Steps

1. In the Policies management page, select the policy, select the **Actions**  icon, and select **Edit policy**.



The process is the same for replication and backup policies.


2. In the Edit Policy page, make the changes.
3. Select **Save**.

Delete a policy

You can delete policies that are not associated with any volumes.

If a policy is associated with a volume and you want to delete the policy, you must remove the policy from the volume first.

Steps

1. In the Policies management page, select the policy, select the **Actions**  icon, and select **Delete Snapshot policy**.
2. Select **Delete**.

Find more information

For instructions on creating policies using System Manager or ONTAP CLI, see the following:

[Create a Snapshot policy using System Manager](#)

[Create a Snapshot policy using the ONTAP CLI](#)

[Create a replication policy using System Manager](#)

[Create a replication policy using the ONTAP CLI](#)

[Create a backup to object storage policy using System Manager](#)

Backup-to-object policy options

BlueXP backup and recovery enables you to create backup policies with a variety of settings for your on-prem ONTAP and Cloud Volumes ONTAP systems.



These policy settings are relevant for backup-to-object storage only. None of these settings affect your Snapshot or replication policies. Similar policy settings for Snapshots and replications will be added in the future.

Backup schedule options

BlueXP backup and recovery enables you to create multiple backup policies with unique schedules for each working environment (cluster). You can assign different backup policies to volumes that have different recovery point objectives (RPO).

Each backup policy provides a section for *Labels & Retention* that you can apply to your backup files. Note that the Snapshot policy applied to the volume must be one of the policies recognized by BlueXP backup and recovery or backup files will not be created.

The screenshot shows a configuration interface for a backup policy. At the top, there is a header bar with 'Name' and 'Default_Policy_Name' fields. Below this is a section titled 'Labels & Retention' which is highlighted with an orange border. Inside this section, on the left, is a list of 12 labels with checkboxes for 'Hourly', 'Daily', 'Weekly', 'Monthly', and 'Yearly'. 'Hourly' and 'Daily' are checked. On the right, under 'Selected Labels (2)', there is a table showing the selected labels and their retention values. The table has two rows: 'Hourly' with 'Number of Backups to Retain' set to 12, and 'Daily' with 'Number of Backups to Retain' set to 30. Below the 'Labels & Retention' section, there are two more sections: 'DataLock & Ransomware Protection' set to 'None' and 'Archival Policy' set to 'Disabled'.

Label	Number of Backups to Retain
Hourly	12
Daily	30

There are two parts of the schedule; the Label and the Retention value:

- The **label** defines how often a backup file is created (or updated) from the volume. You can select among the following types of labels:
 - You can choose one, or a combination of, **hourly**, **daily**, **weekly**, **monthly**, and **yearly** timeframes.
 - You can select one of the system-defined policies that provide backup and retention for 3 months, 1 year, or 7 years.
 - If you have created custom backup protection policies on the cluster using ONTAP System Manager or the ONTAP CLI, you can select one of those policies.

- The **retention** value defines how many backup files for each label (timeframe) are retained. Once the maximum number of backups have been reached in a category, or interval, older backups are removed so you always have the most current backups. This also saves you storage costs because obsolete backups don't continue to take up space in the cloud.

For example, say you create a backup policy that creates 7 **weekly** and 12 **monthly** backups:

- each week and each month a backup file is created for the volume
- at the 8th week, the first weekly backup is removed, and the new weekly backup for the 8th week is added (keeping a maximum of 7 weekly backups)
- at the 13th month, the first monthly backup is removed, and the new monthly backup for the 13th month is added (keeping a maximum of 12 monthly backups)

Note that Yearly backups will be deleted automatically from the source system after being transferred to object storage. This default behavior can be changed [in the Advanced Settings page](#) for the Working Environment.

DataLock and Ransomware protection options

BlueXP backup and recovery provides support for DataLock and Ransomware protection for your volume backups. These features enable you to lock your backup files and scan them to detect possible ransomware on the backup files. This is an optional setting that you can define in your backup policies when you want extra protection for your volume backups for a cluster.

Both of these features protect your backup files so that you'll always have a valid backup file to recover data from in case of a ransomware attack attempt on your backups. It's also helpful to meet certain regulatory requirements where backups need to be locked and retained for a certain period of time. When the DataLock and Ransomware Protection option is enabled, the cloud bucket that is provisioned as a part of BlueXP backup and recovery activation will have object locking and object versioning enabled.

[See the DataLock and Ransomware protection blog for more details.](#)

This feature does not provide protection for your source volumes; only for the backups of those source volumes. Use NetApp [Cloud Insights and Cloud Secure](#), or some of the [anti-ransomware protections provided from ONTAP](#) to protect your source volumes.



- If you plan to use DataLock and Ransomware protection, you can enable it when creating your first backup policy and activating BlueXP backup and recovery for that cluster. You can later enable it using BlueXP backup and recovery Advanced Settings.
- DataLock and Ransomware protection can be disabled for a cluster after it has been configured to save costs.
- When BlueXP scans a backup file for ransomware when restoring volume data, you'll incur extra egress costs from your cloud provider to access the contents of the backup file.

What is DataLock

DataLock protects your backup files from being modified or deleted for a certain period of time - also called *immutable storage*. This functionality uses technology from the object storage provider for "object locking." The period of time that the backup file is locked (and retained) is called the DataLock Retention Period. It is based on the backup policy schedule and retention setting that you defined, plus a maximum 31-day buffer. Any DataLock retention policy that is less than 31 days is rounded up to 31 days minimum.

Be aware that old backups are deleted after the DataLock Retention Period expires, not after the backup policy

retention period expires.

Let's look at some examples of how this works:

- If you create a Monthly backup schedule with 12 retentions, each backup is locked for 12 months (plus a maximum 31-day buffer) before it is deleted.
- If you create a backup policy that creates 30 daily, 7 weekly, 12 monthly backups there will be three locked retention periods. The "30 daily" backups would be retained for 44 days (30 days plus a maximum 31-day buffer), the "7 weekly" backups would be retained for 9 weeks (7 weeks plus a maximum 31-day buffer), and the "12 monthly" backups would be retained for 12 months (plus a maximum 31-day buffer).
- If you create an Hourly backup schedule with 24 retentions, you might think that backups are locked for 24 hours. However, since that is less than the minimum of 30 days, each backup will be locked and retained for 44 days (30 days plus a maximum 31-day buffer).

You can see in this last case that if each backup file is locked for 30 days (plus a maximum 31-day buffer), you'll end up with many more backup files than would typically be retained with an hourly/24 retentions policy. Usually, when BlueXP backup and recovery creates the 25th backup file it would delete the oldest backup to keep the maximum retentions at 24 (based on the policy). The DataLock retention setting overrides the policy retention setting from your backup policy in this case. This could affect your storage costs as your backup files will be saved in the object store for a longer period of time.

What is Ransomware protection

Ransomware protection scans your backup files to look for evidence of a ransomware attack. The detection of ransomware attacks is performed using a checksum comparison. If potential ransomware is identified in a new backup file versus the previous backup file, that newer backup file is replaced by the most recent backup file that does not show any signs of a ransomware attack. (The file that was identified as having a ransomware attack is deleted 1 day after it has been replaced.)

Ransomware scans happen at the following points in the backup and restore process:

- When a backup file is created.

You can optionally enable or disable ransomware scans.

The scan is not performed on the backup file when it is first written to cloud storage, but when the **next** backup file is written. For example, if you have a weekly backup schedule set for Tuesday, on Tuesday the 14th a backup is created. Then on Tuesday the 21st another backup is created. The ransomware scan is run on the backup file from the 14th at this time.

- When you attempt to restore data from a backup file

You can choose to run a scan before restoring data from a backup file, or skip this scan.

- Manually

You can run an on-demand ransomware protection scan at any time to verify the health of a specific backup file. This can be useful if you've had a ransomware issue on a particular volume and you want to verify that the backups for that volume are not affected.

DataLock and Ransomware Protection options

Each backup policy provides a section for *DataLock and Ransomware Protection* that you can apply to your

backup files.

AWS	Azure
<div><p>DataLock & Ransomware Protection</p><p>Backup copies are protected from being modified or deleted, and they are scanned for ransomware threats.</p><p><input checked="" type="radio"/> None</p><p><input type="radio"/> Governance Users with specific permissions can overwrite or delete protected backup files during the retention period</p><p><input type="radio"/> Compliance No users can overwrite or delete protected backup files during the retention period</p></div>	<div><p>DataLock & Ransomware Protection</p><p>Backup copies are protected from being modified or deleted, and they are scanned for ransomware threats.</p><p><input checked="" type="radio"/> None</p><p><input type="radio"/> Unlocked Backup files are protected during the retention period. The retention period can be increased or decreased. Typically used for 24 hours just to test the system.</p><p><input type="radio"/> Locked Backup files are protected during the retention period. The retention period can be increased, but it can't be decreased. Satisfies full regulatory compliance.</p></div>
<div><p>StorageGRID</p><div><p>DataLock & Ransomware Protection</p><p>Backup copies are protected from being modified or deleted, and they are scanned for ransomware threats.</p><p><input checked="" type="radio"/> None</p><p><input type="radio"/> Compliance No users can overwrite or delete protected backup files during the retention period</p></div></div>	

Ransomware protection scans are enabled by default. The default setting for the scan frequency is for 7 days. The scan occurs only on the latest Snapshot copy. You can enable or disable ransomware scans on the latest Snapshot copy by using the option on the Advanced Settings page. If you enable it, scans are performed every 7 days by default.

You can change that schedule to days or weeks or disable it, saving costs.

Refer to [How to update Ransomware protection options in the Advanced Settings page](#).

You can choose from the following settings for each backup policy:

AWS

- **None** (Default)

DataLock protection and ransomware protection are disabled.

- **Governance**

DataLock is set to *Governance* mode where users with `s3:BypassGovernanceRetention` permission ([see below](#)) can overwrite or delete backup files during the retention period. Ransomware protection is enabled.

- **Compliance**

DataLock is set to *Compliance* mode where no users can overwrite or delete backup files during the retention period. Ransomware protection is enabled.

Azure

- **None** (Default)

DataLock protection and ransomware protection are disabled.

- **Unlocked**

Backup files are protected during the retention period. The retention period can be increased or decreased. Typically used for 24 hours to test the system. Ransomware protection is enabled.

- **Locked**

Backup files are protected during the retention period. The retention period can be increased, but it can't be decreased. Satisfies full regulatory compliance. Ransomware protection is enabled.

StorageGRID

- **None** (Default)

DataLock protection and ransomware protection are disabled.

- **Compliance**

DataLock is set to *Compliance* mode where no users can overwrite or delete backup files during the retention period. Ransomware protection is enabled.

Supported working environments and object storage providers

You can enable DataLock and Ransomware protection on ONTAP volumes from the following working environments when using object storage in the following public and private cloud providers. Additional cloud providers will be added in future releases.

Source Working Environment	Backup File Destination
Cloud Volumes ONTAP in AWS	Amazon S3
Cloud Volumes ONTAP in Azure	Azure Blob

Source Working Environment	Backup File Destination
On-premises ONTAP system	Amazon S3 Azure Blob NetApp StorageGRID

Requirements

- For AWS:
 - Your clusters must running ONTAP 9.11.1 or greater
 - The Connector can be deployed in the cloud or on your premises
 - The following S3 permissions must be part of the IAM role that provides the Connector with permissions. They reside in the "backupS3Policy" section for the resource "arn:aws:s3:::netapp-backup-*".

AWS S3 permissions

- s3:GetObjectVersionTagging
- s3:GetBucketObjectLockConfiguration
- s3:GetObjectVersionAcl
- s3:PutObjectTagging
- s3:DeleteObject
- s3:DeleteObjectTagging
- s3:GetObjectRetention
- s3:DeleteObjectVersionTagging
- s3:PutObject
- s3:GetObject
- s3:PutBucketObjectLockConfiguration
- s3:GetLifecycleConfiguration
- s3:GetBucketTagging
- s3:DeleteObjectVersion
- s3:ListBucketVersions
- s3:ListBucket
- s3:PutBucketTagging
- s3:GetObjectTagging
- s3:PutBucketVersioning
- s3:PutObjectVersionTagging
- s3:GetBucketVersioning
- s3:GetBucketAcl
- s3:BypassGovernanceRetention
- s3:PutObjectRetention
- s3:GetBucketLocation
- s3:GetObjectVersion

[View the full JSON format for the policy where you can copy and paste required permissions.](#)

- For Azure:
 - Your clusters must running ONTAP 9.12.1 or greater
 - The Connector can be deployed in the cloud or on your premises
- For StorageGRID:
 - Your clusters must running ONTAP 9.11.1 or greater
 - Your StorageGRID systems must be running 11.6.0.3 or greater
 - The Connector must be deployed on your premises (it can be installed in a site with or without internet access)

- The following S3 permissions must be part of the IAM role that provides the Connector with permissions:

StorageGRID S3 permissions

- s3:GetObjectVersionTagging
- s3:GetBucketObjectLockConfiguration
- s3:GetObjectVersionAcl
- s3:PutObjectTagging
- s3:DeleteObject
- s3:DeleteObjectTagging
- s3:GetObjectRetention
- s3:DeleteObjectVersionTagging
- s3:PutObject
- s3:GetObject
- s3:PutBucketObjectLockConfiguration
- s3:GetLifecycleConfiguration
- s3:GetBucketTagging
- s3:DeleteObjectVersion
- s3:ListBucketVersions
- s3:ListBucket
- s3:PutBucketTagging
- s3:GetObjectTagging
- s3:PutBucketVersioning
- s3:PutObjectVersionTagging
- s3:GetBucketVersioning
- s3:GetBucketAcl
- s3:PutObjectRetention
- s3:GetBucketLocation
- s3:GetObjectVersion

Restrictions

- The DataLock and Ransomware protection feature is not available if you have configured archival storage in the backup policy.
- The DataLock option you select when activating BlueXP backup and recovery must be used for all backup policies for that cluster.
- You cannot use multiple DataLock modes on a single cluster.
- If you enable DataLock, all volume backups will be locked. You can't mix locked and non-locked volume backups for a single cluster.

- DataLock and Ransomware protection is applicable for new volume backups using a backup policy with DataLock and Ransomware protection enabled. You can later enable or disable these features using the Advanced Settings option.
- FlexGroup volumes can use DataLock and Ransomware protection only when using ONTAP 9.13.1 or greater.

Tips on how to mitigate DataLock costs

You can enable or disable the Ransomware Scan feature while keeping the DataLock feature active. To avoid extra charges, you can disable scheduled ransomware scans. This lets you customize your security settings and avoid incurring costs from the cloud provider.

Even if scheduled ransomware scans are disabled, you can still perform on-demand scans when needed.

You can choose different levels of protection:

- **DataLock *without* ransomware scans:** Provides protection for backup data in the destination storage that can be either in Governance or Compliance mode.
 - **Governance mode:** Offers flexibility to administrators to overwrite or delete protected data.
 - **Compliance mode:** Provides complete indelibility until the retention period expires. This helps meet the most stringent data security requirements of highly regulated environments. The data cannot be overwritten or modified during its lifecycle, providing the strongest level of protection for your backup copies.



Microsoft Azure uses a Lock and Unlock mode instead.

- **DataLock *with* ransomware scans:** Provides an additional layer of security for your data. This feature helps detect any attempts to change backup copies. If any attempt is made, a new version of the data is created discreetly. The scan frequency can be changed to 1, 2, 3, 4, 5, 6, or 7 days. If scans are set to every 7 days, the costs decrease significantly.

For more tips to mitigate DataLock costs, refer to

<https://community.netapp.com/t5/Tech-ONTAP-Blogs/Understanding-BlueXP-Backup-and-Recovery-DataLock-and-Ransomware-Feature-TCO/ba-p/453475>

Additionally, you can get estimates for the cost associated with DataLock by visiting the [BlueXP backup and recovery Total Cost of Ownership \(TCO\) calculator](#).

Archival storage options

When using AWS, Azure, or Google cloud storage, you can move older backup files to a less expensive archival storage class or access tier after a certain number of days. You can also choose to send your backup files to archival storage immediately without being written to standard cloud storage. Just enter **0** as the "Archive After Days" to send your backup file directly to archival storage. This can be especially helpful for users who rarely need to access data from cloud backups or users who are replacing a backup to tape solution.

Data in archival tiers can't be accessed immediately when needed, and will require a higher retrieval cost, so you'll need to consider how often you may need to restore data from backup files before deciding to archive your backup files.



- Even if you select “0” to send all data blocks to archival cloud storage, metadata blocks are always written to standard cloud storage.
- Archival storage can’t be used if you have enabled DataLock.
- You can’t change the archival policy after selecting **0** days (archive immediately).

Each backup policy provides a section for *Archival Policy* that you can apply to your backup files.

Name	Default_Policy_Name	⌵
Labels & Retention	30 Daily	⌵
DataLock & Ransomware Protection	None	⌵
Archival Policy	<p>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.</p> <p><input checked="" type="checkbox"/> Tier Backups to Archive</p> <p>Archive After (Days) <input type="text" value="30"/> Storage Class <input type="text" value="S3 Glacier"/></p>	

- In AWS, backups start in the *Standard* storage class and transition to the *Standard-Infrequent Access* storage class after 30 days.

If your cluster is using ONTAP 9.10.1 or greater, you can tier older backups to either *S3 Glacier* or *S3 Glacier Deep Archive* storage. [Learn more about AWS archival storage.](#)

- If you select no archive tier in your first backup policy when activating BlueXP backup and recovery, then *S3 Glacier* will be your only archive option for future policies.
- If you select *S3 Glacier* in your first backup policy, then you can change to the *S3 Glacier Deep Archive* tier for future backup policies for that cluster.
- If you select *S3 Glacier Deep Archive* in your first backup policy, then that tier will be the only archive tier available for future backup policies for that cluster.
- In Azure, backups are associated with the *Cool* access tier.

If your cluster is using ONTAP 9.10.1 or greater, you can tier older backups to *Azure Archive* storage. [Learn more about Azure archival storage.](#)

- In GCP, backups are associated with the *Standard* storage class.

If your on-prem cluster is using ONTAP 9.12.1 or greater, you can choose to tier older backups to *Archive* storage in the BlueXP backup and recovery UI after a certain number of days for further cost optimization. [Learn more about Google archival storage.](#)

- In StorageGRID, backups are associated with the *Standard* storage class.

If your on-prem cluster is using ONTAP 9.12.1 or greater, and your StorageGRID system is using 11.4 or greater, you can archive older backup files to public cloud archival storage.

- For AWS, you can tier backups to AWS *S3 Glacier* or *S3 Glacier Deep Archive* storage. [Learn more](#)

[about AWS archival storage.](#)

- For Azure, you can tier older backups to *Azure Archive* storage. [Learn more about Azure archival storage.](#)

[Learn more about archiving backup files from StorageGRID.](#)

Manage backup-to-object storage options in the Advanced Settings page

You can change cluster-level, backup-to-object storage settings that you set when activating BlueXP backup and recovery for each ONTAP system by using the Advanced Settings page. You can also modify some settings that are applied as "default" backup settings. This includes changing the transfer rate of backups to object storage, whether historical Snapshot copies are exported as backup files, and enabling or disabling ransomware scans for a working environment.



These settings are available for backup-to-object storage only. None of these settings affect your Snapshot or replication settings. Similar cluster-level replications settings for Snapshots and replications will be added in the future.

You can change the following options in the Advanced Settings page:


- Changing the network bandwidth allocated to upload backups to object storage using the Max Transfer Rate option
- Changing whether historical Snapshot copies are exported as backup files and included in your initial baseline backup files for future volumes
- Changing whether "yearly" snapshots are removed from the source system
- Enabling or disabling ransomware scans for a working environment, including scheduled scans

View cluster-level backup settings

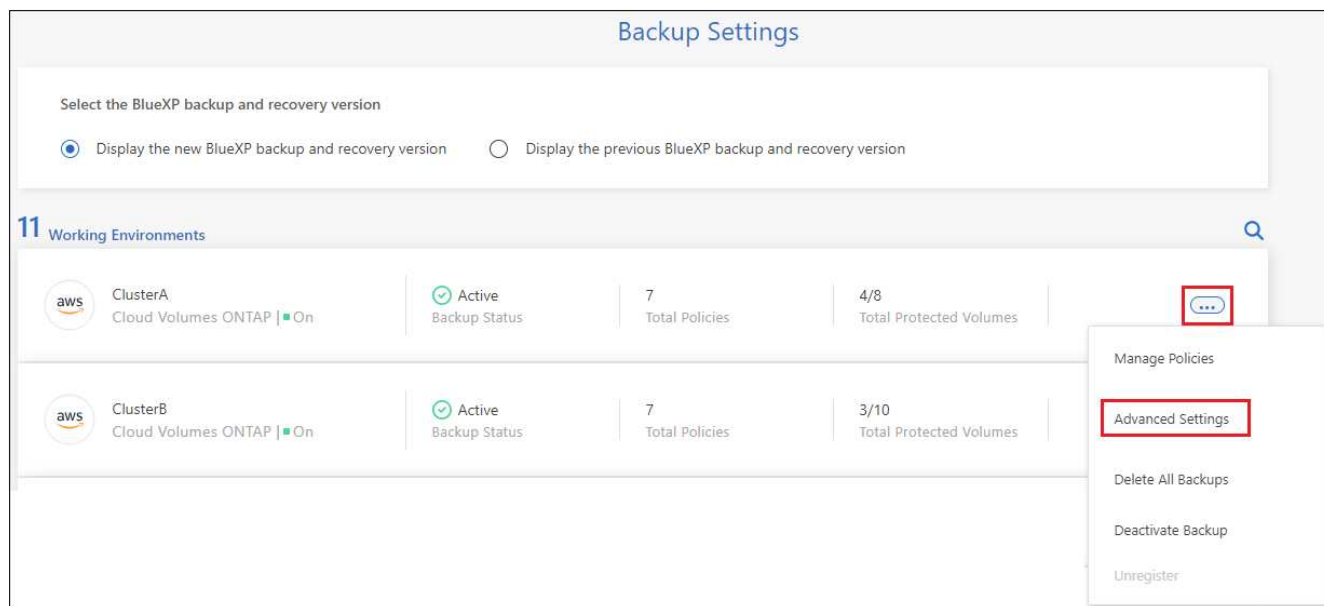
You can view the cluster-level backup settings for each working environment.

Steps

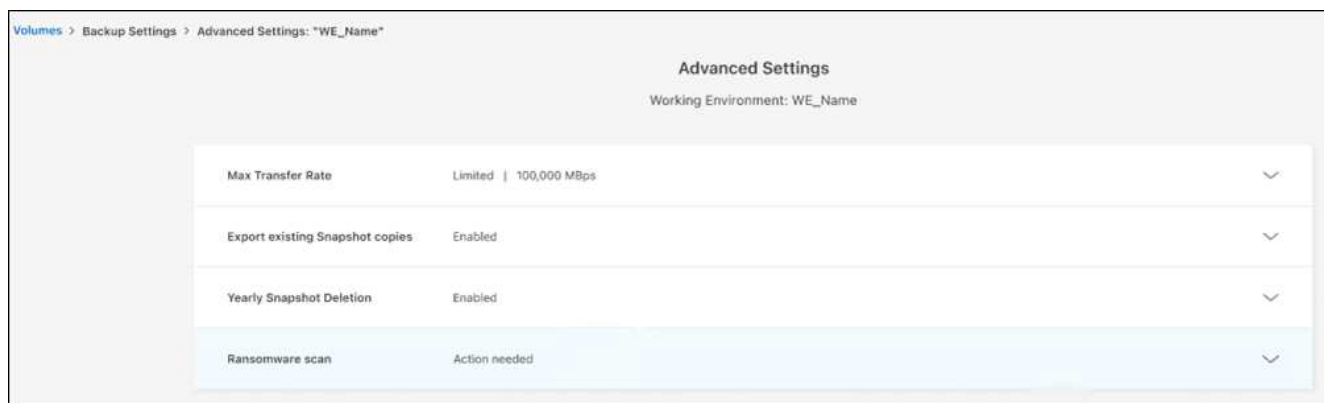
1. From the BlueXP menu, select **Protection > Backup and recovery**.
2. From the **Volumes** tab, select **Backup Settings**.



3. From the *Backup Settings* page, click ... for the working environment and select **Advanced Settings**.



The *Advanced Settings* page displays the current settings for that working environment.



4. Expand the option and make the change.

All backup operations after the change will use the new values.

Note that some options are unavailable based on the version of ONTAP on the source cluster, and based on the cloud provider destination where the backups reside.

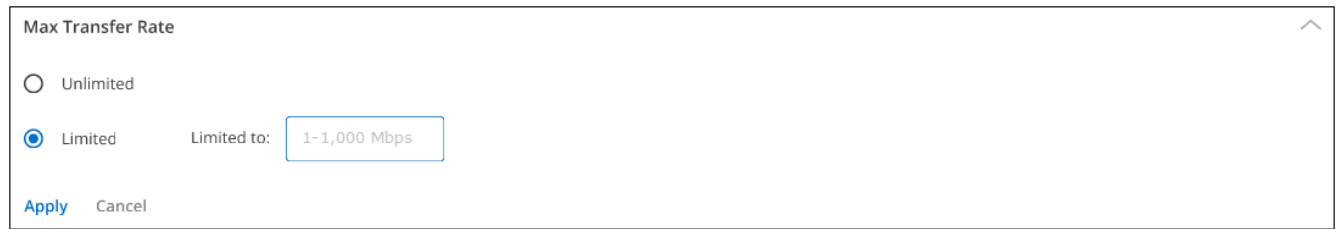
Change the network bandwidth available to upload backups to object storage

When you activate BlueXP backup and recovery for a working environment, by default, ONTAP can use an unlimited amount of bandwidth to transfer the backup data from volumes in the working environment to object storage. If you notice that backup traffic is affecting normal user workloads, you can throttle the amount of network bandwidth that is used during the transfer using the Max Transfer Rate option in the Advanced Settings page.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings* page, click ... for the working environment and select **Advanced Settings**.

3. In the Advanced Settings page, expand the **Max Transfer Rate** section.

A dialog box titled "Max Transfer Rate" with a close button in the top right corner. It contains two radio buttons: "Unlimited" and "Limited". The "Limited" radio button is selected. To the right of the "Limited" radio button is a text input field labeled "Limited to:" containing the value "1-1,000 Mbps". At the bottom left are two buttons: "Apply" and "Cancel".

Max Transfer Rate

☐ Unlimited

☒ Limited Limited to:

Apply Cancel

4. Choose a value between 1 and 1,000 Mbps as the maximum transfer rate.
5. Select the **Limited** radio button and enter the maximum bandwidth that can be used, or select **Unlimited** to indicate that there is no limit.
6. Select **Apply**.

This setting does not affect the bandwidth allocated to any other replication relationships that may be configured for volumes in the working environment.

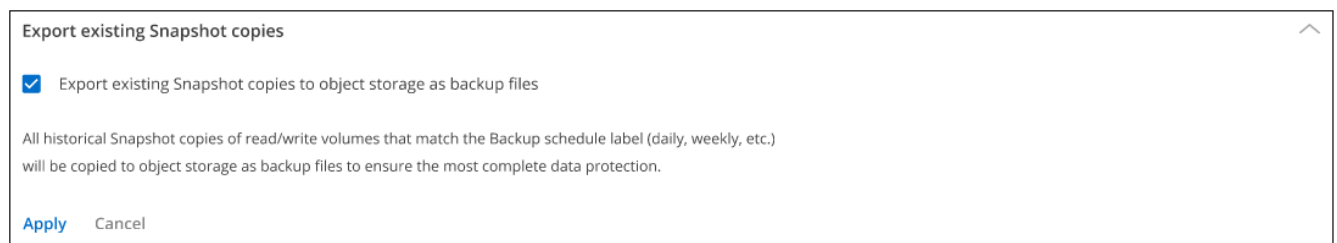
Change whether historical Snapshot copies are exported as backup files

If there are any local Snapshot copies for volumes that match the backup schedule label you're using in this working environment (for example, daily, weekly, etc.), you can export those historic snapshots to object storage as backup files. This enables you to initialize your backups in the cloud by moving older snapshot copies into the baseline backup copy.

Note that this option only applies to new backup files for new read/write volumes, and it is not supported with data protection (DP) volumes.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings* page, click **...** for the working environment and select **Advanced Settings**.
3. In the Advanced Settings page, expand the **Export existing Snapshot copies** section.

A dialog box titled "Export existing Snapshot copies" with a close button in the top right corner. It contains a checked checkbox labeled "Export existing Snapshot copies to object storage as backup files". Below the checkbox is a paragraph of text: "All historical Snapshot copies of read/write volumes that match the Backup schedule label (daily, weekly, etc.) will be copied to object storage as backup files to ensure the most complete data protection." At the bottom left are two buttons: "Apply" and "Cancel".

Export existing Snapshot copies

☒ Export existing Snapshot copies to object storage as backup files

All historical Snapshot copies of read/write volumes that match the Backup schedule label (daily, weekly, etc.) will be copied to object storage as backup files to ensure the most complete data protection.

Apply Cancel

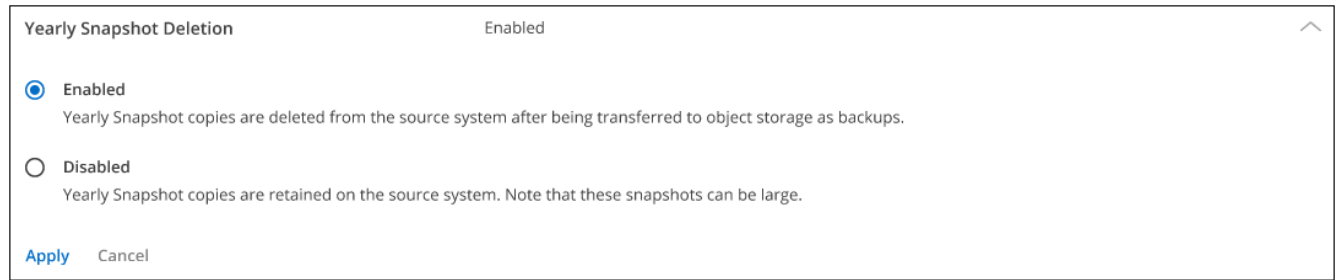
4. Select whether you want existing Snapshot copies to be exported.
5. Select **Apply**.

Change whether "yearly" snapshots are removed from the source system

When you select the "yearly" backup label for a backup policy for any of your volumes, the Snapshot copy that is created is very large. By default, these yearly snapshots are deleted automatically from the source system after being transferred to object storage. You can change this default behavior from the Yearly Snapshot Deletion section.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings page*, click **...** for the working environment and select **Advanced Settings**.
3. In the Advanced Settings page, expand the **Yearly Snapshot Deletion** section.



4. Select **Disabled** to retain the yearly snapshots on the source system.
5. Select **Apply**.

Enable or disable ransomware scans

Ransomware protection scans are enabled by default. The default setting for the scan frequency is for 7 days. The scan occurs only on the latest Snapshot copy. You can enable or disable ransomware scans on the latest Snapshot copy by using the option on the Advanced Settings page. If you enable it, scans are performed every 7 days by default.

You can change that schedule to days or weeks or disable it, saving costs.



Enabling ransomware scans will incur extra charges depending on the cloud provider.

Scheduled ransomware scans run only on the latest Snapshot copy.

If the scheduled ransomware scans are disabled, you can still perform on-demand scans and the scan during a restore operation will still occur.

Refer to [Manage policies](#) for details about managing policies that implement ransomware detection.

Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings page*, click **...** for the working environment and select **Advanced Settings**.
3. In the Advanced Settings page, expand the **Ransomware scan** section.
4. Enable or disable **Ransomware Scan**.
5. Select **Scheduled ransomware scan**.
6. Optionally, change the every week default scan to days or weeks.
7. Set the how often in days or weeks that the scan should run.
8. Select **Apply**.

Back up Cloud Volumes ONTAP data to Amazon S3

Complete a few steps to get started backing up volume data from your Cloud Volumes ONTAP systems 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 9.8 or later in AWS (ONTAP 9.8P13 and later is recommended).
- You have a valid cloud provider subscription for the storage space where your backups will be located.
- You have subscribed to the [BlueXP Marketplace Backup offering](#), an [AWS annual contract](#), or you have purchased [and activated](#) a BlueXP backup and recovery BYOL license from NetApp.
- You have a Connector installed in AWS:
 - The Connector can be installed in a site with full internet access ("standard mode") or with limited internet connectivity ("restricted mode").
 - The IAM role that provides the BlueXP Connector with permissions includes S3 permissions from the latest [BlueXP policy](#).

2

Prepare your BlueXP Connector

If you already have a Connector deployed in an AWS region, then you're all set. If not, then you'll need to install a BlueXP Connector in AWS to back up Cloud Volumes ONTAP data to AWS. The Connector can be installed in a site with full internet access ("standard mode") or with limited internet connectivity ("restricted mode").

3

Verify license requirements

You'll need to check license requirements for both AWS and BlueXP.

4

Verify ONTAP networking requirements for replicating volumes

Ensure that the primary and secondary storage systems meet ONTAP version and networking requirements.

5

Enable BlueXP backup and recovery

Select the working environment and click **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.

6

Activate backups on your ONTAP volumes

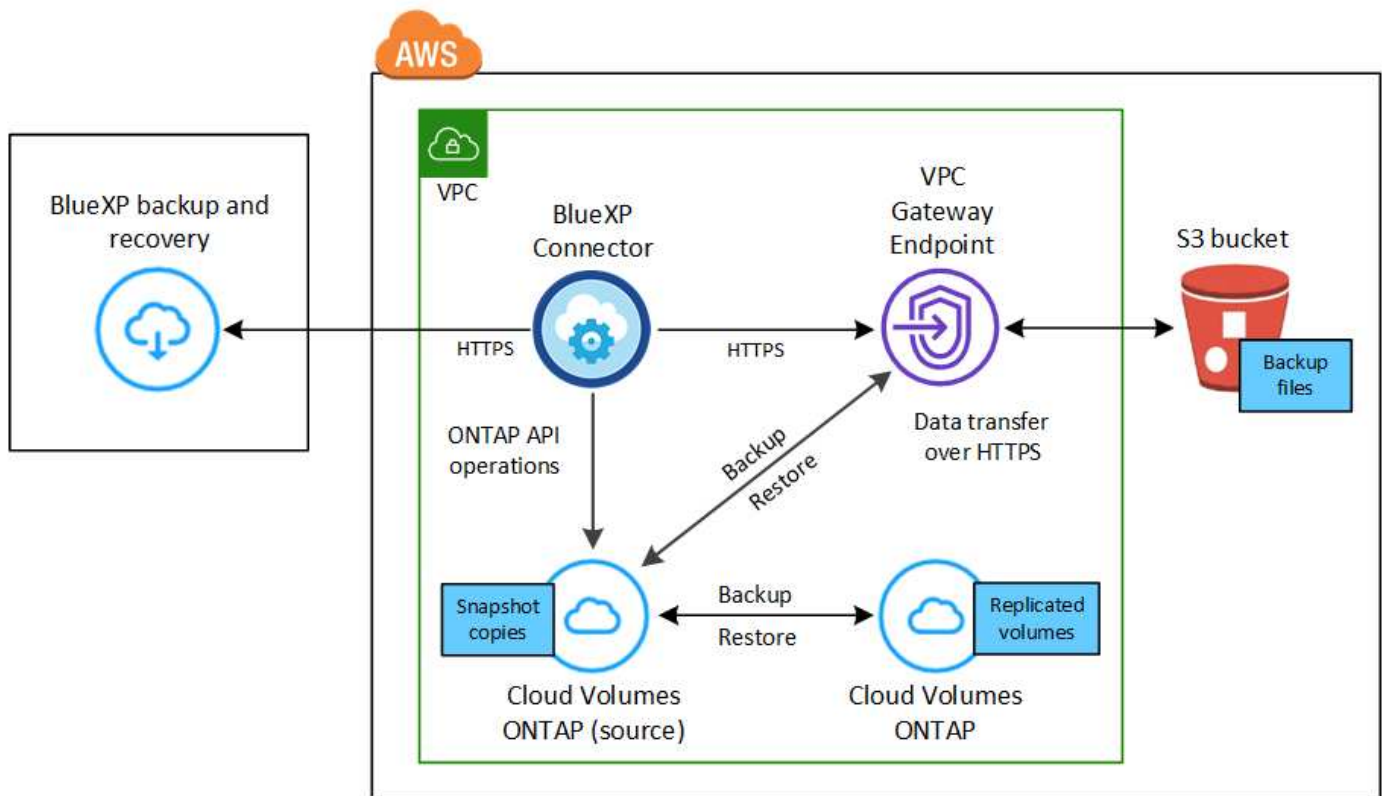
Follow the setup wizard to select the replication and backup policies that you'll use and the volumes you want to back up.

Verify support for your configuration

Read the following requirements to make sure that you have a supported configuration before you start backing up volumes to S3.

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

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.



The VPC gateway endpoint must exist in your VPC already. [Learn more about gateway endpoints.](#)

Supported ONTAP versions

Minimum of ONTAP 9.8; ONTAP 9.8P13 and later is recommended.

Required information for using customer-managed keys for data encryption

You can choose your own customer-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.](#)

Verify license requirements

For BlueXP backup and recovery PAYGO licensing, a BlueXP subscription is available in the AWS Marketplace that enables deployments of Cloud Volumes ONTAP and BlueXP backup and recovery. You need to [subscribe to this BlueXP subscription](#) before you enable BlueXP backup and recovery. Billing for BlueXP backup and recovery is done through this subscription.

For an annual contract that enables you to back up both Cloud Volumes ONTAP data and on-premises ONTAP data, you need to subscribe from the [AWS Marketplace page](#) and then [associate the subscription with your AWS credentials](#).

For an annual contract that enables you to bundle Cloud Volumes ONTAP and BlueXP backup and recovery, you must set up the annual contract when you create a Cloud Volumes ONTAP working environment. This option doesn't enable you to back up on-prem data.

For BlueXP backup and recovery BYOL licensing, 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 must use a BYOL license when the Connector and Cloud Volumes ONTAP system are deployed in a dark site.

And you need to have an AWS account for the storage space where your backups will be located.

Prepare your BlueXP Connector

The Connector must be installed in an AWS region with full or limited internet access ("standard" or "restricted" mode). [See BlueXP deployment modes for details](#).

- [Learn about Connectors](#)
- [Deploy a Connector in AWS in standard mode \(full internet access\)](#)
- [Install the Connector in restricted mode \(limited outbound access\)](#)

Verify or add permissions to the Connector

The IAM role that provides BlueXP with permissions must include S3 permissions from the latest [BlueXP policy](#). If the policy does not contain all of these permissions, see the [AWS Documentation: Editing IAM policies](#).

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:DeleteObject",
    "s3:PutObject",
    "s3:ListBucket",
    "s3:ListAllMyBuckets",
    "s3:GetBucketTagging",
    "s3:GetBucketLocation",
    "s3:GetBucketPolicyStatus",
    "s3:GetBucketPublicAccessBlock",
    "s3:GetBucketAcl",
    "s3:GetBucketPolicy",
    "s3:PutBucketPolicy",
    "s3:PutBucketOwnershipControls",
    "s3:PutBucketPublicAccessBlock",
    "s3:PutEncryptionConfiguration",
    "s3:GetObjectVersionTagging",
    "s3:GetBucketObjectLockConfiguration",
    "s3:GetObjectVersionAcl",
    "s3:PutObjectTagging",
    "s3:DeleteObjectTagging",
    "s3:GetObjectRetention",
    "s3:DeleteObjectVersionTagging",
    "s3:PutBucketObjectLockConfiguration",
    "s3:DeleteObjectVersion",
    "s3:GetObjectTagging",
    "s3:PutBucketVersioning",
    "s3:PutObjectVersionTagging",
    "s3:GetBucketVersioning",
    "s3:BypassGovernanceRetention",
    "s3:PutObjectRetention",
    "s3:GetObjectVersion",
    "athena:StartQueryExecution",
    "athena:GetQueryResults",
    "athena:GetQueryExecution",
    "glue:GetDatabase",
    "glue:GetTable",

```



```

        "glue:CreateTable",
        "glue:CreateDatabase",
        "glue:GetPartitions",
        "glue:BatchCreatePartition",
        "glue:BatchDeletePartition"
    ],
    "Resource": [
        "arn:aws:s3:::netapp-backup-*"
    ]
}

```



When creating backups in AWS China regions, you need to change the AWS Resource Name "arn" under all *Resource* sections in the IAM policies from "aws" to "aws-cn"; for example `arn:aws-cn:s3:::netapp-backup-*`.

Required AWS Cloud Volumes ONTAP permissions

When your Cloud Volumes ONTAP system is running ONTAP 9.12.1 or greater software, the IAM role that provides that working environment with permissions must include a new set of S3 permissions specifically for BlueXP backup and recovery from the latest [Cloud Volumes ONTAP policy](#).

If you created the Cloud Volumes ONTAP working environment using BlueXP version 3.9.23 or greater, these permissions should be part of the IAM role already. Otherwise you'll need to add the missing permissions.

Supported AWS regions

BlueXP backup and recovery is supported in all AWS regions [where Cloud Volumes ONTAP is supported](#), including AWS GovCloud regions.

Required setup for creating backups in a different AWS account

By default, backups are created using the same account as the one used for your Cloud Volumes ONTAP system. If you want to use a different AWS account for your backups, you must:

- Verify that the permissions "s3:PutBucketPolicy" and "s3:PutBucketOwnershipControls" are part of the IAM role that provides the BlueXP Connector with permissions.
- Add the destination AWS account credentials in BlueXP. [See how to do this](#).
- Add the following permissions in the user credentials in the second account:

```
"athena:StartQueryExecution",
"athena:GetQueryResults",
"athena:GetQueryExecution",
"glue:GetDatabase",
"glue:GetTable",
"glue:CreateTable",
"glue:CreateDatabase",
"glue:GetPartitions",
"glue:BatchCreatePartition",
"glue:BatchDeletePartition"
```

Create your own buckets

By default, the service creates buckets for you. If you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets.](#)

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in task-backup-to-s3.adoc - include::.../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

- To replicate data between two Cloud Volumes ONTAP systems in different subnets, the subnets must be routed together (this is the default setting).

Enable BlueXP backup and recovery on Cloud Volumes ONTAP

Enabling BlueXP backup and recovery is easy. The steps differ slightly depending on whether you have an existing Cloud Volumes ONTAP system or a new one.

Enable BlueXP backup and recovery on a new system

BlueXP backup and recovery is enabled by default in the working environment wizard. Be sure to keep the option enabled.

See [Launching Cloud Volumes ONTAP in AWS](#) for requirements and details for creating your Cloud Volumes ONTAP system.

Steps

1. From the BlueXP Canvas, select **Add Working Environment**, choose the cloud provider, and select **Add New**. Select **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 select **Continue**.



5. Complete the pages in the wizard to deploy the system.

Result

BlueXP backup and recovery is enabled on the system. After you've created volumes on these Cloud Volumes ONTAP systems, launch BlueXP backup and recovery and [activate backup on each volume that you want to protect](#).

Enable BlueXP backup and recovery on an existing system

Enable BlueXP backup and recovery on an existing system at any time directly from the working environment.

Steps

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

If the Amazon S3 destination for your backups exists as a working environment on the Canvas, you can drag the cluster onto the Amazon S3 working environment to initiate the setup wizard.



To modify backup settings or add replication, refer to [Manage ONTAP backups](#).

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises working environment.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)


You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future working environments.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.



If the AWS destination for your backups exists as a working environment on the Canvas, you can drag the ONTAP cluster onto the AWS object storage.

- Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions** ... icon option and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a BlueXP Connector, you're all set. Just select **Next**.
- If you don't already have a BlueXP Connector, the **Add a Connector** option appears. Refer to [Prepare your BlueXP Connector](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a working environment. See how to [activate backup for additional volumes in the working environment](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row. (☒ Volume Name).

- To back up individual volumes, check the box for each volume (☒ Volume_1).

2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture
- Local Snapshot policy
- Replication target and policy



If the volumes you choose have different Snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots**: If you are performing replication or back up to object storage, local Snapshots must be created.
 - **Replication**: Creates replicated volumes on another ONTAP storage system.
 - **Backup**: Backs up volumes to object storage.
2. **Architecture**: If you chose replication and backup, choose one of the following flows of information:
 - **Cascading**: Information flows from the primary storage system to the secondary, and from secondary to object storage.
 - **Fan out**: Information flows from the primary storage system to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot**: Choose an existing Snapshot policy or create a new one.



To create a custom policy before activating the Snapshot, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

4. **Replication**: Set the following options:

- **Replication target**: Select the destination working environment and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy**: Choose an existing replication policy or create one.



To create a custom policy, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **Amazon Web Services**.
- **Provider settings:** Enter the provider details and region where the backups will be stored.

Enter the AWS Account used to store the backups. This can be a different account than where the Cloud Volumes ONTAP system resides.

If you want to use a different AWS account for your backups, you must add the destination AWS account credentials in BlueXP, and add the permissions "s3:PutBucketPolicy" and "s3:PutBucketOwnershipControls" to the IAM role that provides BlueXP with permissions.

Select the region where the backups will be stored. This can be a different region than where the Cloud Volumes ONTAP system resides.

Either create a new bucket or select an existing one.

- **Encryption key:** If you created a new bucket, enter encryption key information given to you from the provider. Choose whether you'll use the default AWS 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 encryption keys](#)).

If you choose to use your own customer-managed keys, enter the key vault and key information.



If you chose an existing bucket, encryption information is already available, so you don't need to enter it now.

- **Backup policy:** Select an existing backup-to-object storage policy or create one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- For backup-to-object policies, set the DataLock and Ransomware Protection settings. For details on DataLock and Ransomware Protection, refer to [Backup-to-object policy settings](#).

- Select **Create**.

- **Export existing Snapshot copies to object storage as backup copies:** If there are any local Snapshot copies for volumes in this working environment that match the backup schedule label you just selected for this working environment (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to

ensure the most complete protection for your volumes.

1. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

BlueXP backup and recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary storage system data contained in Snapshot copies.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage volume.

An S3 bucket is created 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.

You can also monitor the status of backup and restore jobs using the [Job Monitoring panel](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future working environments.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.
- You can [manage cluster-level backup settings](#). This includes changing the storage keys ONTAP uses to access cloud storage, changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to a Cloud Volumes ONTAP system in AWS, or to an on-premises ONTAP system.

Back up Cloud Volumes ONTAP data to Azure Blob storage

Complete a few steps to get started backing up volume data from your Cloud Volumes ONTAP systems 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 9.8 or later in Azure (ONTAP 9.8P13 and later is recommended).
- You have a valid cloud provider subscription for the storage space where your backups will be located.
- You have subscribed to the [BlueXP Marketplace Backup offering](#), or you have purchased [and activated](#) a BlueXP backup and recovery BYOL license from NetApp.

2

Prepare your BlueXP Connector

If you already have a Connector deployed in an Azure region, then you're all set. If not, then you'll need to install a BlueXP Connector in Azure to back up Cloud Volumes ONTAP data to Azure Blob storage. The Connector can be installed in a site with full internet access ("standard mode") or with limited internet connectivity ("restricted mode").

[Prepare your BlueXP Connector](#)

3

Verify license requirements

You'll need to check license requirements for both Azure and BlueXP.

Refer to [Verify license requirements](#).

4

Verify ONTAP networking requirements for replicating volumes

Ensure that the source and destination systems meet ONTAP version and networking requirements.

[Verify ONTAP networking requirements for replicating volumes](#).

5

Enable BlueXP backup and recovery

Select the working environment and click **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.

[Enable BlueXP backup and recovery on Cloud Volumes ONTAP](#).

6

Activate backups on your ONTAP volumes

Follow the setup wizard to select the replication and backup policies that you'll use and the volumes you want

to back up.


[Activate backups on your ONTAP volumes.](#)

Verify support for your configuration

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

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

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.



Supported ONTAP versions

Minimum of ONTAP 9.8; ONTAP 9.8P13 and later is recommended.

Supported Azure regions

BlueXP backup and recovery is supported in all Azure regions [where Cloud Volumes ONTAP is supported](#); including Azure Government regions.

By default, BlueXP backup and recovery provisions the Blob container with Local redundancy (LRS) for cost optimization. You can change this setting to Zone redundancy (ZRS) after BlueXP backup and recovery has been activated if you want to make sure your data is replicated between different zones. See the Microsoft instructions for [changing how your storage account is replicated](#).

Required setup for creating backups in a different Azure subscription

By default, backups are created using the same subscription as the one used for your Cloud Volumes ONTAP system. If you want to use a different Azure subscription for your backups, you must [log in to the Azure portal and link the two subscriptions](#).

Verify license requirements

For BlueXP backup and recovery PAYGO licensing, a subscription through the Azure Marketplace is required before you enable BlueXP backup and recovery. Billing for BlueXP backup and recovery is done through this subscription. [You can subscribe from the Details & Credentials page of the working environment wizard.](#)

For BlueXP backup and recovery BYOL licensing, 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 must use a BYOL license when the Connector and Cloud Volumes ONTAP system are deployed in a dark site ("private mode").

And you need to have a Microsoft Azure subscription for the storage space where your backups will be located.

Prepare your BlueXP Connector

The Connector can be installed in an Azure region with full or limited internet access ("standard" or "restricted" mode). [See BlueXP deployment modes for details.](#)

- [Learn about Connectors](#)
- [Deploy a Connector in Azure in standard mode \(full internet access\)](#)
- [Install the Connector in restricted mode \(limited outbound access\)](#)

Verify or add permissions to the Connector

To use the BlueXP backup and recovery Search & Restore functionality, you need to have specific permissions in the role for the Connector so that it can access the Azure Synapse Workspace and Data Lake Storage Account. See the permissions below, and follow the steps if you need to modify the policy.

Before you start

- You must register the Azure Synapse Analytics Resource Provider (called "Microsoft.Synapse") with your Subscription. [See how to register this resource provider for your subscription.](#) You must be the Subscription **Owner** or **Contributor** to register the resource provider.
- Port 1433 must be open for communication between the Connector and the Azure Synapse SQL services.

Steps

1. Identify the role assigned to the Connector virtual machine:
 - a. In the Azure portal, open the virtual machines service.
 - b. Select the Connector virtual machine.
 - c. Under Settings, select **Identity**.
 - d. Select **Azure role assignments**.
 - e. Make note of the custom role assigned to the Connector virtual machine.
2. Update the custom role:
 - a. In the Azure portal, open your Azure subscription.
 - b. Select **Access control (IAM) > Roles**.
 - c. Select the ellipsis (...) for the custom role and then select **Edit**.
 - d. Select **JSON** and add the following permissions:

```

"Microsoft.Storage/storageAccounts/listkeys/action",
"Microsoft.Storage/storageAccounts/read",
"Microsoft.Storage/storageAccounts/write",
"Microsoft.Storage/storageAccounts/blobServices/containers/read",
"Microsoft.Storage/storageAccounts/listAccountSas/action",
"Microsoft.KeyVault/vaults/read",
"Microsoft.KeyVault/vaults/accessPolicies/write",
"Microsoft.Network/networkInterfaces/read",
"Microsoft.Resources/subscriptions/locations/read",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/virtualNetworks/subnets/read",
"Microsoft.Resources/subscriptions/resourceGroups/read",
"Microsoft.Resources/subscriptions/resourcegroups/resources/read",
"Microsoft.Resources/subscriptions/resourceGroups/write",
"Microsoft.Authorization/locks/*",
"Microsoft.Network/privateEndpoints/write",
"Microsoft.Network/privateEndpoints/read",
"Microsoft.Network/privateDnsZones/virtualNetworkLinks/write",
"Microsoft.Network/virtualNetworks/join/action",
"Microsoft.Network/privateDnsZones/A/write",
"Microsoft.Network/privateDnsZones/read",
"Microsoft.Network/privateDnsZones/virtualNetworkLinks/read",
"Microsoft.Network/networkInterfaces/delete",
"Microsoft.Network/networkSecurityGroups/delete",
"Microsoft.Resources/deployments/delete",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
"Microsoft.Synapse/workspaces/write",
"Microsoft.Synapse/workspaces/read",
"Microsoft.Synapse/workspaces/delete",
"Microsoft.Synapse/register/action",
"Microsoft.Synapse/checkNameAvailability/action",
"Microsoft.Synapse/workspaces/operationStatuses/read",
"Microsoft.Synapse/workspaces/firewallRules/read",
"Microsoft.Synapse/workspaces/replaceAllIpFirewallRules/action",
"Microsoft.Synapse/workspaces/operationResults/read",
"Microsoft.Synapse/workspaces/privateEndpointConnectionsApproval/
action"

```

[View the full JSON format for the policy](#)

e. Click **Review + update** and then click **Update**.

Required information for using customer-managed keys for data encryption

You can use your own customer-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.](#)

BlueXP backup and recovery supports *Azure access policies*, the *Azure role-based access control* (Azure RBAC) permission model and the *Managed Hardware Security Model* (HSM) (refer to [What is Azure Key Vault Managed HSM?](#)).

Create your Azure Blob storage account

By default, the service creates storage accounts for you. If you want to use your own storage accounts, you can create them before you start the backup activation wizard and then select those storage accounts in the wizard.

[Learn more about creating your own storage accounts.](#)

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in task-backup-to-azure.adoc - include::.../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

- To replicate data between two Cloud Volumes ONTAP systems in different subnets, the subnets must be routed together (this is the default setting).

Enable BlueXP backup and recovery on Cloud Volumes ONTAP

Enabling BlueXP backup and recovery is easy. The steps differ slightly depending on whether you have an existing Cloud Volumes ONTAP system or a new one.

Enable BlueXP backup and recovery on a new system

BlueXP backup and recovery is enabled by default in the working environment wizard. Be sure to keep the option enabled.

See [Launching Cloud Volumes ONTAP in Azure](#) for requirements and details for creating your Cloud Volumes ONTAP system.



If you want to pick the name of the resource group, **disable** BlueXP backup and recovery when deploying Cloud Volumes ONTAP. Follow the steps for [enabling BlueXP backup and recovery on an existing system](#) to enable BlueXP backup and recovery and choose the resource group.

Steps

1. From the BlueXP Canvas, select **Add Working Environment**, choose the cloud provider, and select **Add New**. Select **Create Cloud Volumes ONTAP**.
2. Select **Microsoft Azure** as the cloud provider and then choose a single node or HA system.
3. In the Define Azure Credentials page, enter the credentials name, client ID, client secret, and directory ID, and click **Continue**.
4. Fill out the Details & Credentials page and be sure that an Azure Marketplace subscription is in place, and click **Continue**.
5. On the Services page, leave the service enabled and click **Continue**.



6. Complete the pages in the wizard to deploy the system.

Result

BlueXP backup and recovery is enabled on the system. After you've created volumes on these Cloud Volumes ONTAP systems, launch BlueXP backup and recovery and [activate backup on each volume that you want to protect](#).


Enable BlueXP backup and recovery on an existing system

Enable BlueXP backup and recovery at any time directly from the working environment.

Steps

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

If the Azure Blob destination for your backups exists as a working environment on the Canvas, you can drag the cluster onto the Azure Blob working environment to initiate the setup wizard.



2. Complete the pages in the wizard to deploy BlueXP backup and recovery.
3. When you want to initiate backups, continue with [Activate backups on your ONTAP volumes](#).

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises working environment.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)


You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future working environments.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.



If the Azure destination for your backups exists as a working environment on the Canvas, you can drag the ONTAP cluster onto the Azure Blob object storage.

- Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions** ... icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a BlueXP Connector, you're all set. Just select **Next**.
- If you don't already have a BlueXP Connector, the **Add a Connector** option appears. Refer to [Prepare your BlueXP Connector](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup-to-object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a working environment. See how to [activate backup for additional volumes in the working environment](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes. (FlexGroup volumes can be selected one at a time only.) To back up all existing FlexVol volumes, check one volume first and then check the box in the title row. (☒ Volume Name).

- To back up individual volumes, check the box for each volume (☒ Volume_1).

2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture
- Local Snapshot policy
- Replication target and policy



If the volumes you choose have different Snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots**: If you are performing replication or back up to object storage, local Snapshots must be created.
 - **Replication**: Creates replicated volumes on another ONTAP storage system.
 - **Backup**: Backs up volumes to object storage.
2. **Architecture**: If you chose replication and backup, choose one of the following flows of information:
 - **Cascading**: Information flows from the primary storage system to the secondary, and from secondary to object storage.
 - **Fan out**: Information flows from the primary storage system to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot**: Choose an existing Snapshot policy or create one.



To create a custom policy before activating the Snapshot, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

4. **Replication**: Set the following options:

- **Replication target**: Select the destination working environment and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy**: Choose an existing replication policy or create one.



To create a custom policy before activating the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **Microsoft Azure**.
- **Provider settings:** Enter the provider details.

Enter the region where the backups will be stored. This can be a different region than where the Cloud Volumes ONTAP system resides.

Either create a new storage account or select an existing one.

Enter the Azure subscription used to store the backups. This can be a different subscription than where the Cloud Volumes ONTAP system resides. If you want to use a different Azure subscription for your backups, you must [log in to the Azure portal and link the two subscriptions](#).

Either create your own resource group that manages the Blob container or select the resource group type and group.



If you want to protect your backup files from being modified or deleted, ensure that the storage account was created with immutable storage enabled using a 30-day retention period.



If you want to tier older backup files to Azure Archive Storage for further cost optimization, ensure that the storage account has the appropriate Lifecycle rule.

- **Encryption key:** If you created a new Azure storage account, enter encryption key information given to you from the provider. Choose whether you'll use the default Azure encryption keys, or choose your own customer-managed keys from your Azure account, to manage encryption of your data.

If you choose to use your own customer-managed keys, enter the key vault and key information. [Learn how to use your own keys](#).



If you chose an existing Microsoft storage account, encryption information is already available, so you don't need to enter it now.

- **Networking:** Choose the IPspace, and whether you'll be using a Private Endpoint. Private Endpoint is disabled by default.
 - 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 Azure private endpoint that you have previously configured. [Learn about using an Azure private endpoint](#).
- **Backup policy:** Select an existing backup-to-object storage policy.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- For backup-to-object policies, set the DataLock and Ransomware Protection settings. For details on DataLock and Ransomware Protection, refer to [Backup-to-object policy settings](#).
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.
 - **Export existing Snapshot copies to object storage as backup copies:** If there are any local Snapshot copies for volumes in this working environment that match the backup schedule label you just selected for this working environment (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.
- 1. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

BlueXP backup and recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary storage data contained in Snapshot copies.

A replicated volume is created in the destination cluster that will be synchronized with the primary volume.

A Blob storage container is created in the resource group you entered, and the backup files are stored there.

By default, BlueXP backup and recovery provisions the Blob container with Local redundancy (LRS) for cost optimization. You can change this setting to Zone redundancy (ZRS) if you want to make sure your data is replicated between different zones. See the Microsoft instructions for [changing how your storage account is replicated](#).

The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring panel](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future working environments.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.
- You can [manage cluster-level backup settings](#). This includes changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to a Cloud Volumes ONTAP system in Azure, or to an on-premises ONTAP system.

Back up Cloud Volumes ONTAP data to Google Cloud Storage

Complete a few steps to get started backing up volume data from your Cloud Volumes ONTAP systems to Google Cloud 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 9.8 or later in GCP (ONTAP 9.8P13 and later is recommended).
- You have a valid GCP subscription for the storage space where your backups will be located.
- You have a service account in your Google Cloud Project that has a custom role with a reduced set of permissions.



The Storage Admin role is no longer required for the service account that enables BlueXP backup and recovery to access Google Cloud Storage buckets.

- You have subscribed to the [BlueXP Marketplace Backup offering](#), or you have purchased [and activated](#) a BlueXP backup and recovery BYOL license from NetApp.

2

Prepare your BlueXP Connector

If you already have a Connector deployed in a GCP region, then you're all set. If not, then you'll need to install a BlueXP Connector in GCP to back up Cloud Volumes ONTAP data to Google Cloud Storage. The Connector can be installed in a site with full internet access ("standard mode") or with limited internet connectivity ("restricted mode").

3

Verify license requirements

You'll need to check license requirements for both Google Cloud and BlueXP.

4

Verify ONTAP networking requirements for replicating volumes

Ensure that the source and destination systems meet ONTAP version and networking requirements.

5

Enable BlueXP backup and recovery

Select the working environment and click **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.

6

Prepare Google Cloud as your backup target

Set up permissions for the Connector to create and manage the Google Cloud bucket.

Optionally, you can set up your own custom-managed keys for data encryption instead of using the default Google Cloud encryption keys. [Learn how to get your Google Cloud environment ready to receive ONTAP backups.](#)

7

Activate backups on your ONTAP volumes

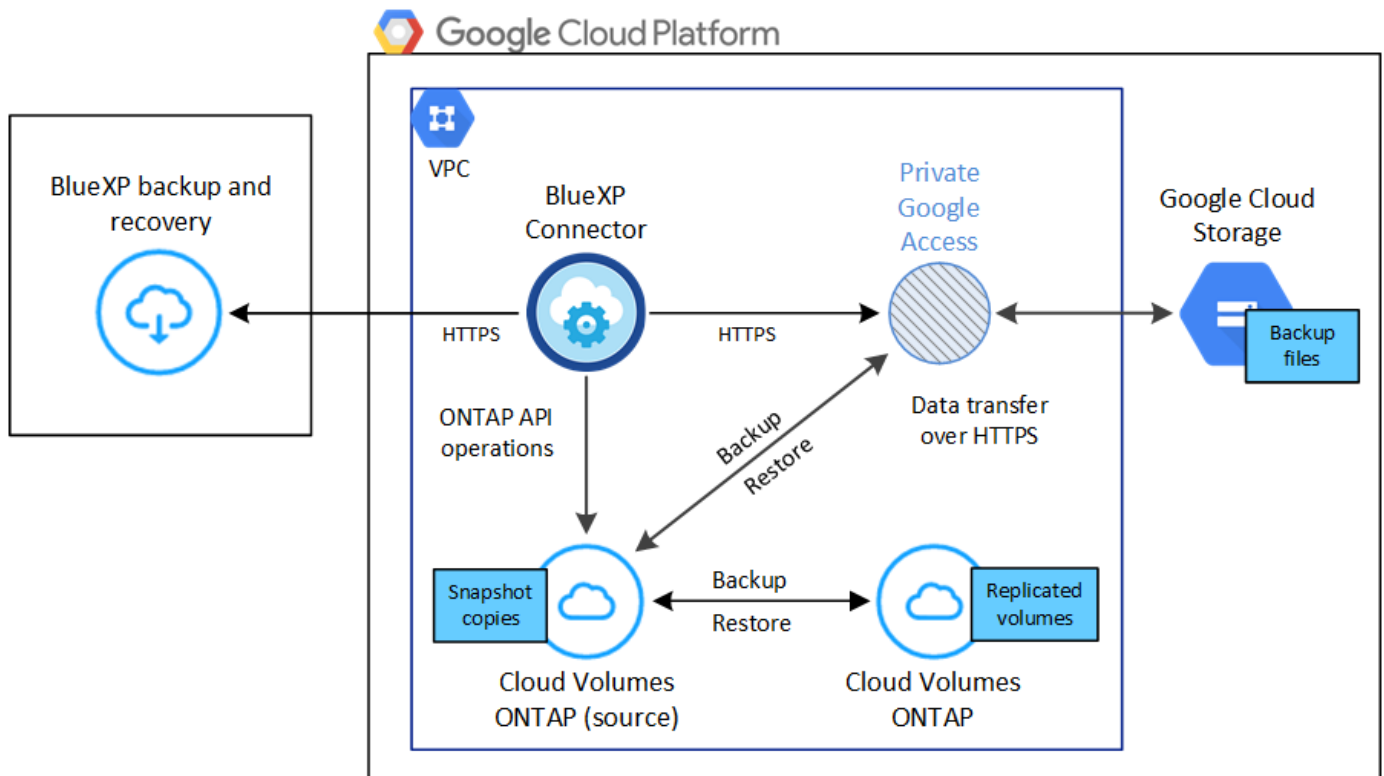
Follow the setup wizard to select the replication and backup policies that you'll use and the volumes you want to back up.

Verify support for your configuration

Read the following requirements to make sure that you have a supported configuration before you start backing up volumes to Google Cloud Storage.

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

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.



Supported ONTAP versions

Minimum of ONTAP 9.8; ONTAP 9.8P13 and later is recommended.

Supported GCP regions

BlueXP backup and recovery is supported in all GCP regions [where Cloud Volumes ONTAP is supported](#).

GCP Service Account

You need to have a service account in your Google Cloud Project that has the custom role. [Learn how to create a service account](#).



The Storage Admin role is no longer required for the service account that enables BlueXP backup and recovery to access Google Cloud Storage buckets.

Verify license requirements

For BlueXP backup and recovery PAYGO licensing, a BlueXP subscription is available in the Google Marketplace that enables deployments of Cloud Volumes ONTAP and BlueXP backup and recovery. You need to [subscribe to this BlueXP subscription](#) before you enable BlueXP backup and recovery. Billing for BlueXP backup and recovery is done through this subscription. [You can subscribe from the Details & Credentials page of the working environment wizard](#).

For BlueXP backup and recovery BYOL licensing, 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](#).

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

Prepare your BlueXP Connector

The Connector must be installed in a Google region with internet access.

- [Learn about Connectors](#)
- [Deploy a Connector in Google Cloud](#)

Verify or add permissions to the Connector

To use the BlueXP backup and recovery "Search & Restore" functionality, you need to have specific permissions in the role for the Connector so that it can access the Google Cloud BigQuery service. See the permissions below, and follow the steps if you need to modify the policy.

Steps

1. In the [Google Cloud Console](#), go to the **Roles** page.
2. Using the drop-down list at the top of the page, select the project or organization that contains the role that you want to edit.
3. Select a custom role.
4. Select **Edit Role** to update the role's permissions.
5. Select **Add Permissions** to add the following new permissions to the role.

```
bigquery.jobs.get  
bigquery.jobs.list  
bigquery.jobs.listAll  
bigquery.datasets.create  
bigquery.datasets.get  
bigquery.jobs.create  
bigquery.tables.get  
bigquery.tables.getData  
bigquery.tables.list  
bigquery.tables.create
```

6. Select **Update** to save the edited role.

Required information for using customer-managed encryption keys (CMEK)

You can use your own customer-managed keys for data encryption instead of using the default Google-managed encryption keys. Both cross-region and cross-project keys are supported, so you can choose a project for a bucket that is different than the project of the CMEK key. If you're planning to use your own customer-managed keys:

- You'll need to have the Key Ring and the Key Name so you can add this information in the activation wizard. [Learn more about customer-managed encryption keys](#).
- You'll need to verify that these required permissions are included in the role for the Connector:

```
cloudkms.cryptoKeys.get
cloudkms.cryptoKeys.getIamPolicy
cloudkms.cryptoKeys.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.keyRings.get
cloudkms.keyRings.getIamPolicy
cloudkms.keyRings.list
cloudkms.keyRings.setIamPolicy
```

- You'll need to verify that the Google "Cloud Key Management Service (KMS)" API is enabled in your project. See the [Google Cloud documentation: Enabling APIs](#) for details.

CMEK considerations:

- Both HSM (hardware-backed) and software-generated keys are supported.
- Both newly created or imported Cloud KMS keys are supported.
- Only regional keys are supported; global keys are not supported.
- Currently, only the "Symmetric encrypt/decrypt" purpose is supported.
- The service agent associated with the Storage Account is assigned the "CryptoKey Encrypter/Decrypter (roles/cloudkms.cryptoKeyEncrypterDecrypter)" IAM role by BlueXP backup and recovery.

Create your own buckets

By default, the service creates buckets for you. If you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets.](#)

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in task-backup-to-gcp.adoc - include::.../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

- To replicate data between two Cloud Volumes ONTAP systems in different subnets, the subnets must be routed together (this is the default setting).

Enable BlueXP backup and recovery on Cloud Volumes ONTAP

Enabling BlueXP backup and recovery is easy. The steps differ slightly depending on whether you have an existing Cloud Volumes ONTAP system or a new one.

Enable BlueXP backup and recovery on a new system

BlueXP backup and recovery can be enabled when you complete the working environment wizard to create a new Cloud Volumes ONTAP system.

You must have a Service Account already configured. If you don't 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.

See [Launching Cloud Volumes ONTAP in GCP](#) for requirements and details for creating your Cloud Volumes ONTAP system.

Steps

1. From the BlueXP Canvas, select **Add Working Environment**, choose the cloud provider, and select **Add New**. Select **Create Cloud Volumes ONTAP**.
2. **Choose a Location**: Select **Google Cloud Platform**.
3. **Choose Type**: Select **Cloud Volumes ONTAP** (either single-node or high-availability).
4. **Details & Credentials**: Enter the following information:
 - a. Click **Edit Project** and select a new project if the one you want to use is different than the default Project (where the Connector resides).
 - b. Specify the cluster name.
 - c. Enable the **Service Account** switch and select the Service Account that has the predefined Storage Admin role. This is required to enable backups and tiering.
 - d. Specify the credentials.

Make sure that a GCP Marketplace subscription is in place.

The screenshot shows the 'Details & Credentials' configuration page. At the top, there are two tabs: 'Project1' (selected) and 'MPAWSSubscription1222'. Below the tabs, it shows 'Google Cloud Project' and 'Marketplace Subscription'. An 'Edit Project' button is in the top right. The main content is divided into two columns: 'Details' and 'Credentials'. In the 'Details' column, there is a 'Working Environment Name (Cluster Name)' field with the value 'TamiVSA'. Below it is a 'Service Account' section with a toggle switch turned on, a 'Service Account Name' dropdown menu showing 'ServiceAccount1', and an 'Add Labels' button with the text 'Optional Field | Up to four labels'. In the 'Credentials' column, there is a 'User Name' field with the value 'admin', a 'Password' field with masked characters '*****', and a 'Confirm Password' field with masked characters '*****'.

5. **Services**: Leave the BlueXP backup and recovery service enabled and click **Continue**.

The screenshot shows the 'Services' configuration page. It features a single service entry, 'Backup to Cloud', which includes a cloud icon, the service name, a toggle switch that is turned on, and a dropdown arrow on the right.

6. Complete the pages in the wizard to deploy the system as described in [Launching Cloud Volumes ONTAP in GCP](#).



To modify backup settings or add replication, refer to [Manage ONTAP backups](#).

Result

BlueXP backup and recovery is enabled on the system. After you've created volumes on these Cloud Volumes ONTAP systems, launch BlueXP backup and recovery and [activate backup on each volume that you want to protect](#).

Enable BlueXP backup and recovery on an existing system

You can enable BlueXP backup and recovery at any time directly from the working environment.

Steps

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

If the Google Cloud Storage destination for your backups exists as a working environment on the Canvas, you can drag the cluster onto the Google Cloud Storage working environment to initiate the setup wizard.



To modify backup settings or add replication, refer to [Manage ONTAP backups](#).

Prepare Google Cloud Storage as your backup target

Preparing Google Cloud Storage as your backup target involves the following steps:

- Set up permissions.
- (Optional) Create your own buckets. (The service will create buckets for you if you want.)
- (Optional) Set up customer-managed keys for data encryption

Set up permissions

You need to provide storage access keys for a service account that has specific permissions using a custom role. A service account enables BlueXP backup and recovery 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. In the [Google Cloud Console](#), go to the **Roles** page.
2. [Create a new role](#) with the following permissions:


```
storage.buckets.create
storage.buckets.delete
storage.buckets.get
storage.buckets.list
storage.buckets.update
storage.buckets.getIamPolicy
storage.multipartUploads.create
storage.objects.create
storage.objects.delete
storage.objects.get
storage.objects.list
storage.objects.update
```

3. In the Google Cloud console, [go to the Service accounts page](#).
4. Select your Cloud project.
5. Select **Create service account** and provide the required information:
 - a. **Service account details**: Enter a name and description.
 - b. **Grant this service account access to project**: Select the custom role that you just created.
 - c. Select **Done**.
6. Go to [GCP Storage Settings](#) and create access keys for the service account:
 - a. Select a project, and select **Interoperability**. If you haven't already done so, select **Enable interoperability access**.
 - b. Under **Access keys for service accounts**, select **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 BlueXP backup and recovery later when you configure the backup service.

Create your own buckets

By default, the service creates buckets for you. Or, if you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets](#).

Set up customer-managed encryption keys (CMEK) for data encryption

You can use your own customer-managed keys for data encryption instead of using the default Google-managed encryption keys. Both cross-region and cross-project keys are supported, so you can choose a project for a bucket that is different than the project of the CMEK key.

If you're planning to use your own customer-managed keys:

- You'll need to have the Key Ring and the Key Name so you can add this information in the activation wizard. [Learn more about customer-managed encryption keys](#).
- You'll need to verify that these required permissions are included in the role for the Connector:

```
cloudkms.cryptoKeys.get
cloudkms.cryptoKeys.getIamPolicy
cloudkms.cryptoKeys.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.keyRings.get
cloudkms.keyRings.getIamPolicy
cloudkms.keyRings.list
cloudkms.keyRings.setIamPolicy
```

- You'll need to verify that the Google "Cloud Key Management Service (KMS)" API is enabled in your project. See the [Google Cloud documentation: Enabling APIs](#) for details.

CMEK considerations:

- Both HSM (Hardware-backed) and Software-generated keys are supported.
- Both newly created or imported Cloud KMS keys are supported.
- Only regional keys are supported, global keys are not supported.
- Currently, only the "Symmetric encrypt/decrypt" purpose is supported.
- The service agent associated with the Storage Account is assigned the "CryptoKey Encrypter/Decrypter (roles/cloudkms.cryptoKeyEncrypterDecrypter)" IAM role by BlueXP backup and recovery.

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises working environment.

A wizard takes you through the following major steps:

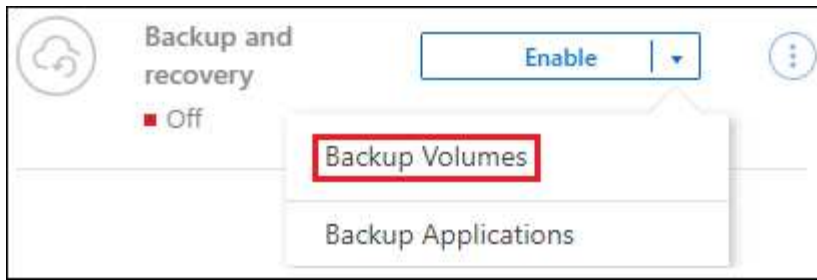
- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future working environments.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:
 - From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.



If the GCP destination for your backups exists as a working environment on the Canvas, you can drag the ONTAP cluster onto the GCP object storage.

- Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions** ... icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a BlueXP Connector, you're all set. Just select **Next**.
- If you don't already have a BlueXP Connector, the **Add a Connector** option appears. Refer to [Prepare your BlueXP Connector](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a working environment. See how to [activate backup for additional volumes in the working environment](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row. (☒ Volume Name).
 - To back up individual volumes, check the box for each volume (☒ Volume_1).
2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture
- Local Snapshot policy
- Replication target and policy



If the volumes you choose have different Snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots:** If you are performing replication or back up to object storage, local Snapshots must be created.
 - **Replication:** Creates replicated volumes on another ONTAP storage system.
 - **Backup:** Backs up volumes to object storage.
2. **Architecture:** If you chose replication and backup, choose one of the following flows of information:
 - **Cascading:** Information flows from the primary storage system to the secondary, and from secondary to object storage.
 - **Fan out:** Information flows from the primary storage system to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing Snapshot policy or create one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

4. **Replication:** Set the following options:

- **Replication target:** Select the destination working environment and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy:** Choose an existing replication policy or create one.



To create a custom policy before activating the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **Google Cloud**.
- **Provider settings:** Enter the provider details and region where the backups will be stored.

Either create a new bucket or select an existing one.

- **Encryption key:** If you created a new Google bucket, enter encryption key information given to you from the provider. Choose whether you'll use the default Google Cloud encryption keys, or choose your own customer-managed keys from your Google account, to manage encryption of your data.

If you choose to use your own customer-managed keys, enter the key vault and key information.



If you chose an existing Google Cloud bucket, encryption information is already available, so you don't need to enter it now.

- **Backup policy:** Select an existing backup-to-object storage policy or create one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.
- **Export existing Snapshot copies to object storage as backup copies:** If there are any local Snapshot copies for volumes in this working environment that match the backup schedule label you just selected for this working environment (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

BlueXP backup and recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent

transfers contain differential copies of the primary storage system data contained in Snapshot copies.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage system volume.

A Google Cloud Storage bucket is created in the service account indicated by the Google access key and secret key you entered, and the backup files are stored there.

Backups are associated with the *Standard* storage class by default. You can use the lower cost *Nearline*, *Coldline*, or *Archive* storage classes. However, you configure the storage class through Google, not through the BlueXP backup and recovery UI. See the Google topic [Changing the default storage class of a bucket](#) for details.

The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring panel](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future working environments.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.
- You can [manage cluster-level backup settings](#). This includes changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to a Cloud Volumes ONTAP system in Google, or to an on-premises ONTAP system.

Back up on-premises ONTAP data to Amazon S3

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



"On-premises ONTAP systems" include FAS, AFF, and ONTAP Select systems.

Quick start

Get started quickly by following these steps. Details for each step are provided in the following sections in this topic.



Identify the connection method you'll use

Choose whether you'll connect your on-premises ONTAP cluster directly to AWS S3 over the public internet, or

whether you'll use a VPN or AWS Direct Connect and route traffic through a private VPC Endpoint interface to AWS S3.

[Identify the connection method.](#)

2

Prepare your BlueXP Connector

If you already have a Connector deployed in your AWS VPC or on your premises, then you're all set. If not, then you'll need to create a BlueXP Connector to back up ONTAP data to AWS S3 storage. You'll also need to customize network settings for the Connector so that it can connect to AWS S3.

[Learn how to create a Connector and how to define required network settings.](#)

3

Verify license requirements

You'll need to check license requirements for both AWS and BlueXP.

Refer to [Verify license requirements](#).

4

Prepare your ONTAP clusters

Discover your ONTAP clusters in BlueXP, verify that the clusters meet minimum requirements, and customize network settings so the clusters can connect to AWS S3.

[Learn how to get your ONTAP clusters ready.](#)

5

Prepare Amazon S3 as your backup target

Set up permissions for the Connector to create and manage the S3 bucket. You'll also need to set up permissions for the on-premises ONTAP cluster so it can read and write data to the S3 bucket.

Optionally, you can set up your own custom-managed keys for data encryption instead of using the default Amazon S3 encryption keys. [Learn how to get your AWS S3 environment ready to receive ONTAP backups.](#)

6

Activate backups on your ONTAP volumes

Select the working environment and click **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel. Then follow the setup wizard to select the replication and backup policies that you'll use and the volumes you want to back up.

[Activate backups on your ONTAP volumes.](#)

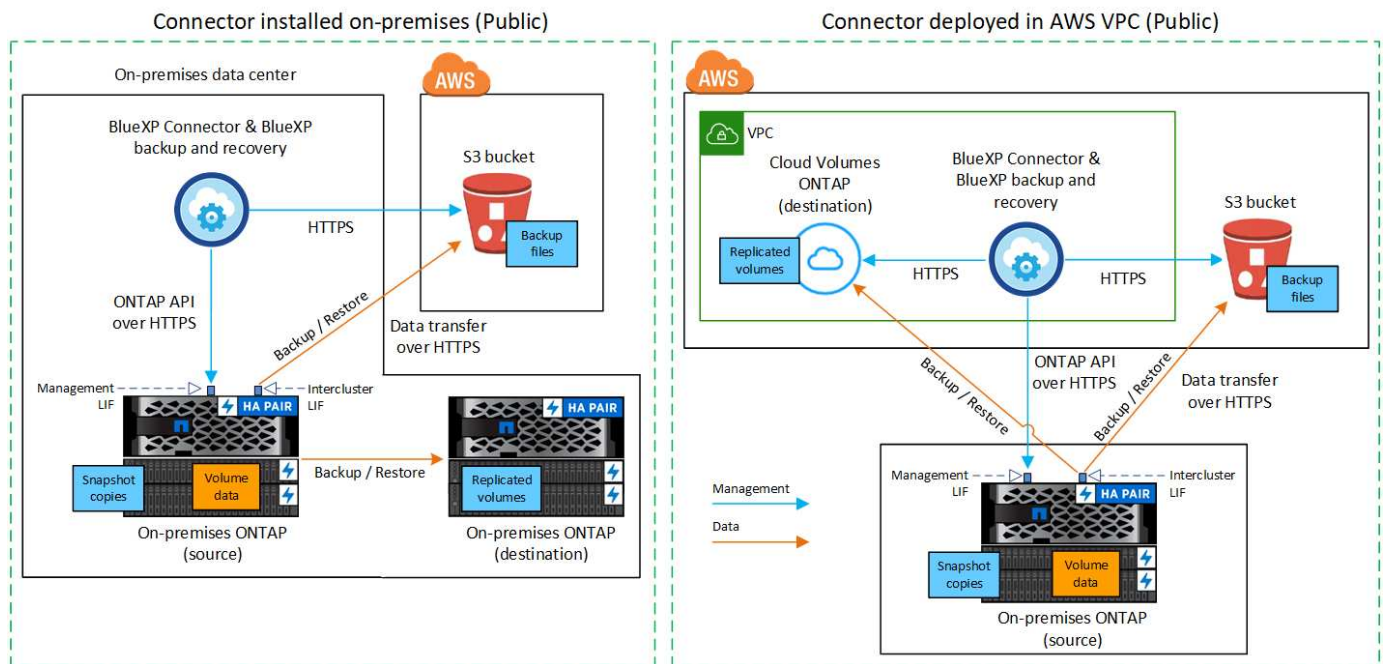
Identify the connection method

Choose which of the two connection methods you will use when configuring backups from on-premises ONTAP systems to AWS S3.

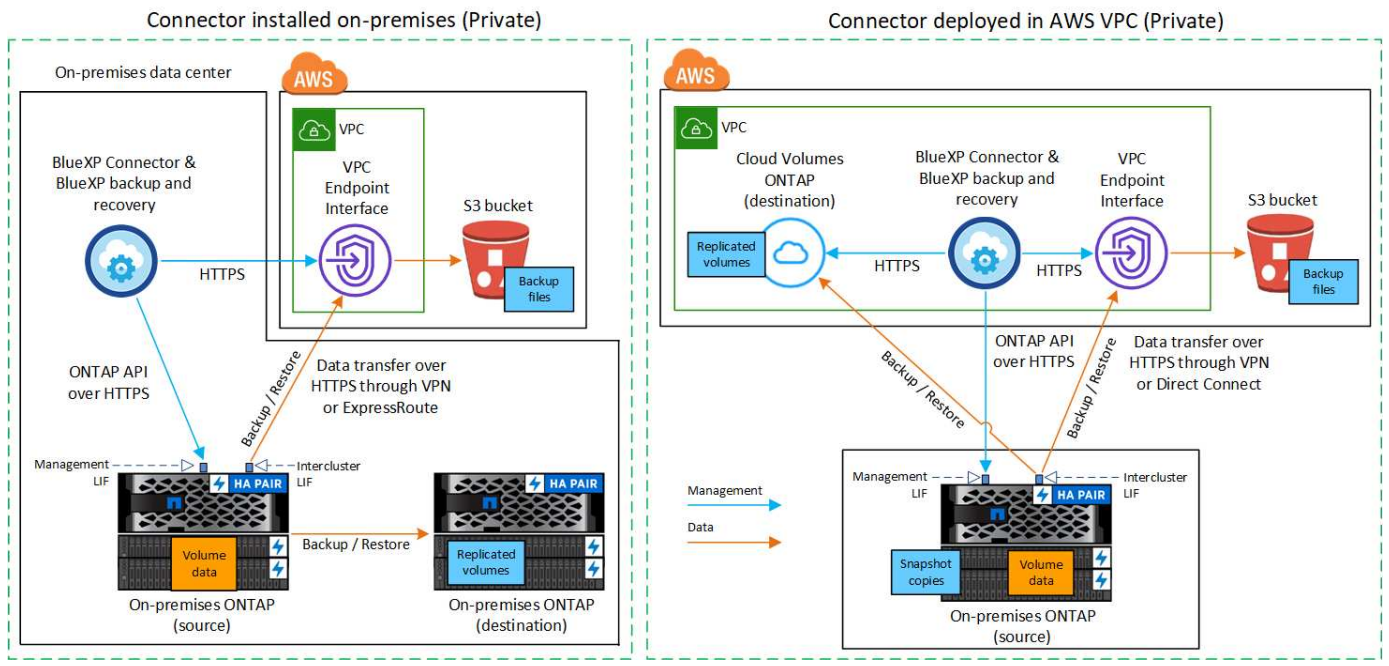
- **Public connection** - Directly connect the ONTAP system to AWS S3 using a public S3 endpoint.
- **Private connection** - Use a VPN or AWS Direct Connect and route traffic through a VPC Endpoint interface that uses a private IP address.

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.

The following diagram shows the **public connection** method and the connections that you need to prepare between the components. You can use a Connector that you've installed on your premises, or a Connector that you've deployed in the AWS VPC.



The following diagram shows the **private connection** method and the connections that you need to prepare between the components. You can use a Connector that you've installed on your premises, or a Connector that you've deployed in the AWS VPC.



Prepare your BlueXP Connector

The BlueXP Connector is the main software for BlueXP functionality. A Connector is required to back up and restore your ONTAP data.

Create or switch Connectors

If you already have a Connector deployed in your AWS VPC or on your premises, then you're all set.

If not, then you'll need to create a Connector in one of those locations to back up ONTAP data to AWS S3 storage. You can't use a Connector that's deployed in another cloud provider.

- [Learn about Connectors](#)
- [Install a Connector in AWS](#)
- [Install a Connector in your premises](#)
- [Install a Connector in an AWS GovCloud region](#)

BlueXP backup and recovery is supported in GovCloud regions when the Connector is deployed in the cloud - not when it's installed in your premises. Additionally, you must deploy the Connector from the AWS Marketplace. You can't deploy the Connector in a Government region from the BlueXP SaaS website.

Prepare Connector networking requirements

Ensure that the following networking requirements are met:

- Ensure that the network where the Connector is installed enables the following connections:
 - An HTTPS connection over port 443 to the BlueXP backup and recovery service and to your S3 object storage ([see the list of endpoints](#))
 - An HTTPS connection over port 443 to your ONTAP cluster management LIF
 - Additional inbound and outbound security group rules are required for AWS and AWS GovCloud deployments. See [Rules for the Connector in AWS](#) for details.
- [Ensure that the Connector has permissions to manage the S3 bucket.](#)
- 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 (a **private** connection), you'll need to enable a VPC Endpoint interface to S3. [See how to set up a VPC endpoint interface.](#)

Verify license requirements

You'll need to verify license requirements for both AWS and BlueXP:

- Before you can activate BlueXP backup and recovery for your cluster, you'll need to either subscribe to a pay-as-you-go (PAYGO) BlueXP Marketplace offering from AWS, or purchase and activate a BlueXP backup and recovery BYOL license from NetApp. These licenses are for your account and can be used across multiple systems.
 - For BlueXP backup and recovery PAYGO licensing, you'll need a subscription to the [NetApp BlueXP offering from the AWS Marketplace](#). Billing for BlueXP backup and recovery is done through this subscription.
 - For BlueXP backup and recovery BYOL licensing, 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.](#)

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

Supported regions

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

Prepare your ONTAP clusters

Unresolved directive in task-backup-onprem-to-aws.adoc - include::.../_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must configure the following requirements on the system that connects to object storage.

- For a fan-out backup architecture, configure the following settings on the *primary* system.
- For a cascaded backup architecture, configure the following settings on the *secondary* system.

The following ONTAP cluster networking requirements are needed:

- The cluster requires an inbound HTTPS connection from the Connector to the cluster management LIF.
- An intercluster LIF is required on each ONTAP node that hosts the volumes you want to back up. These intercluster LIFs must be able to access the object store.

The cluster initiates an outbound HTTPS connection over port 443 from the intercluster LIFs 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.

- The intercluster LIFs must be associated with the *IPspace* that ONTAP should use to connect to object storage. [Learn more about IPspaces.](#)

When you set up BlueXP backup and recovery, you are prompted for the IPspace to use. You should choose the IPspace that these LIFs are associated with. That might be the "Default" IPspace or a custom IPspace that you created.

If you are using a different IPspace than "Default", then you might need to create a static route to get access to the object storage.

All intercluster LIFs within the IPspace must have access to the object store. If you can't configure this for the current IPspace, then you'll need to create a dedicated IPspace where all intercluster LIFs have access to the object store.

- DNS servers must have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- Update firewall rules, if necessary, to allow BlueXP backup and recovery 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).
- If you are using a Private VPC Interface Endpoint in AWS for the S3 connection, then in order for HTTPS/443 to be used, you'll need to load the S3 endpoint certificate into the ONTAP cluster. [See how to](#)

set up a VPC endpoint interface and load the S3 certificate.

- Ensure that your ONTAP cluster has permissions to access the S3 bucket.

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in task-backup-onprem-to-aws.adoc - include::.../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

Prepare Amazon S3 as your backup target

Preparing Amazon S3 as your backup target involves the following steps:

- Set up S3 permissions.
- (Optional) Create your own S3 buckets. (The service will create buckets for you if you want.)
- (Optional) Set up customer-managed AWS keys for data encryption.
- (Optional) Configure your system for a private connection using a VPC endpoint interface.

Set up S3 permissions

You'll need to configure two sets of permissions:

- Permissions for the Connector to create and manage the S3 bucket.
- Permissions for the on-premises ONTAP cluster so it can read and write data to the S3 bucket.

Steps

1. Confirm that the following S3 permissions (from the latest [BlueXP policy](#)) are part of the IAM role that provides the Connector with permissions. If they are not, see the [AWS Documentation: Editing IAM policies](#).

```

{
  "Sid": "backupPolicy",
  "Effect": "Allow",
  "Action": [
    "s3:DeleteBucket",
    "s3:GetLifecycleConfiguration",
    "s3:PutLifecycleConfiguration",
    "s3:PutBucketTagging",
    "s3:ListBucketVersions",
    "s3:GetObject",
    "s3:DeleteObject",
    "s3:PutObject",
    "s3:ListBucket",
    "s3:ListAllMyBuckets",
    "s3:GetBucketTagging",
    "s3:GetBucketLocation",
    "s3:GetBucketPolicyStatus",
    "s3:GetBucketPublicAccessBlock",
    "s3:GetBucketAcl",
    "s3:GetBucketPolicy",
    "s3:PutBucketPolicy",
    "s3:PutBucketOwnershipControls",
    "s3:PutBucketPublicAccessBlock",
    "s3:PutEncryptionConfiguration",
    "s3:GetObjectVersionTagging",
    "s3:GetBucketObjectLockConfiguration",
    "s3:GetObjectVersionAcl",
    "s3:PutObjectTagging",
    "s3:DeleteObjectTagging",
    "s3:GetObjectRetention",
    "s3:DeleteObjectVersionTagging",
    "s3:PutBucketObjectLockConfiguration",
    "s3:DeleteObjectVersion",
    "s3:GetObjectTagging",
    "s3:PutBucketVersioning",
    "s3:PutObjectVersionTagging",
    "s3:GetBucketVersioning",
    "s3:BypassGovernanceRetention",
    "s3:PutObjectRetention",
    "s3:GetObjectVersion",
    "athena:StartQueryExecution",
    "athena:GetQueryResults",
    "athena:GetQueryExecution",
    "glue:GetDatabase",
    "glue:GetTable",

```

```
        "glue:CreateTable",
        "glue:CreateDatabase",
        "glue:GetPartitions",
        "glue:BatchCreatePartition",
        "glue:BatchDeletePartition"
    ],
    "Resource": [
        "arn:aws:s3:::netapp-backup-*"
    ]
}
```



When creating backups in AWS China regions, you need to change the AWS Resource Name "arn" under all *Resource* sections in the IAM policies from "aws" to "aws-cn"; for example `arn:aws-cn:s3:::netapp-backup-*`.

2. When you activate the service, the Backup wizard will prompt you to enter an access key and secret key. These credentials are passed to the ONTAP cluster so that ONTAP can back up and restore data to the S3 bucket. For that, you'll need to create an IAM user with the following permissions.

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

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject",
        "s3:ListBucket",
        "s3:ListAllMyBuckets",
        "s3:GetBucketLocation",
        "s3:PutEncryptionConfiguration"
      ],
      "Resource": "arn:aws:s3:::netapp-backup-*",
      "Effect": "Allow",
      "Sid": "backupPolicy"
    },
    {
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": "arn:aws:s3:::netapp-backup*",
      "Effect": "Allow"
    },
    {
      "Action": [
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject",
        "s3:ListAllMyBuckets",
        "s3:PutObjectTagging",
        "s3:GetObjectTagging",
        "s3:RestoreObject",
        "s3:GetBucketObjectLockConfiguration",
        "s3:GetObjectRetention",
        "s3:PutBucketObjectLockConfiguration",
        "s3:PutObjectRetention"
      ],
      "Resource": "arn:aws:s3:::netapp-backup*/**",
      "Effect": "Allow"
    }
  ]
}

```

Create your own buckets

By default, the service creates buckets for you. Or, if you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets.](#)

If you create your own buckets, you should use a bucket name of “netapp-backup”. If you need to use a custom name, edit the `ontapcloud-instance-policy-netapp-backup` IAMRole for the existing CVOs and add the following list to the S3 permissions. You need to include `“Resource”: “arn:aws:s3:::”` and assign all the necessary permissions that need to be associated with the bucket.

```
"Action": [  
  "S3:ListBucket"  
  "S3:GetBucketLocation"  
]  
"Resource": "arn:aws:s3:::",  
"Effect": "Allow"  
},  
{  
  "Action": [  
    "S3:GetObject",  
    "S3:PutObject",  
    "S3:DeleteObject",  
    "S3:ListAllMyBuckets",  
    "S3:PutObjectTagging",  
    "S3:GetObjectTagging",  
    "S3:RestoreObject",  
    "S3:GetBucketObjectLockConfiguration",  
    "S3:GetObjectRetention",  
    "S3:PutBucketObjectLockConfiguration",  
    "S3:PutObjectRetention"  
  ]  
  "Resource": "arn:aws:s3:::",
```

Set up customer-managed AWS keys for data encryption

If you want to use the default Amazon S3 encryption keys to encrypt the data passed between your on-prem cluster and the S3 bucket, then you are all set because the default installation uses that type of encryption.

If instead you want to use your own customer-managed keys for data encryption rather than using the default keys, then you’ll need to have the encryption managed keys already set up before you start the BlueXP backup and recovery wizard. [Refer to how to use your own keys.](#)

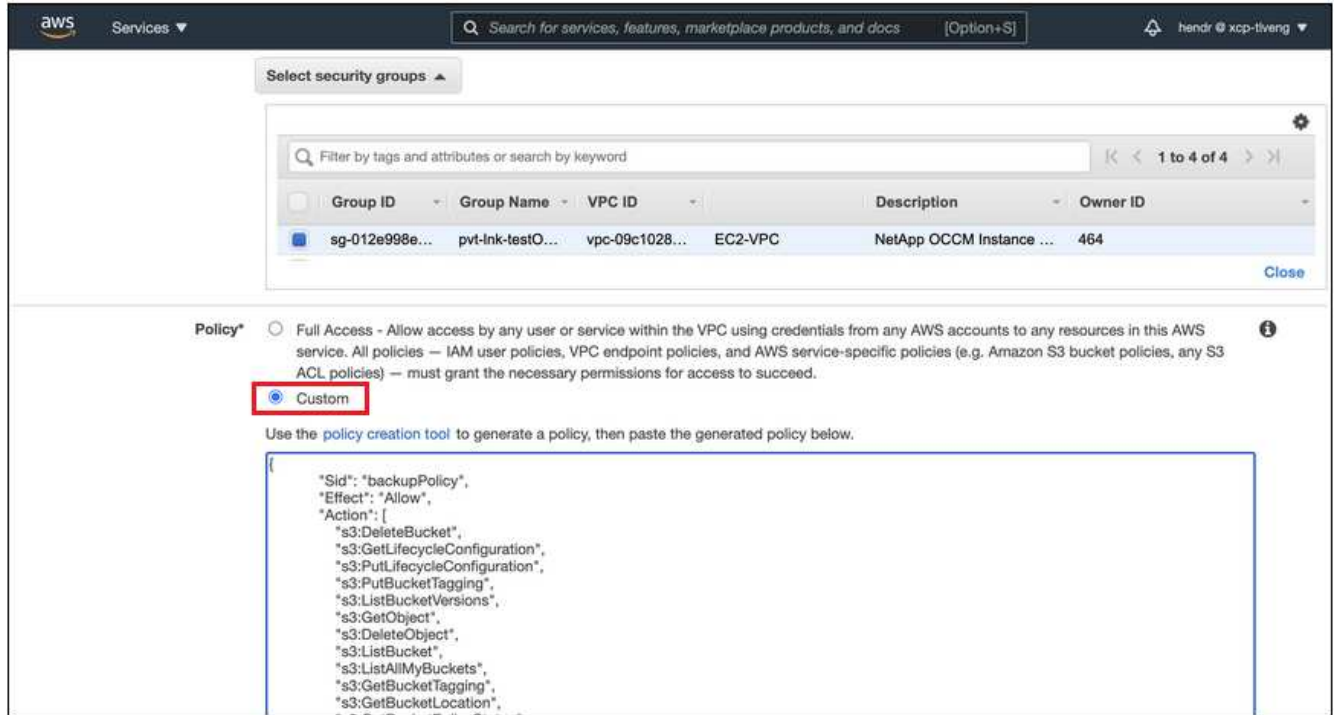
Configure your system for a private connection using a VPC endpoint interface

If you want to use a standard public internet connection, then all the permissions are set by the Connector and there is nothing else you need to do. This type of connection is shown in the [first diagram](#).

If you want to have a more secure connection over the internet from your on-prem data center to the VPC, there’s an option to select an AWS PrivateLink connection in the Backup activation wizard. It’s required if you plan to use a VPN or AWS Direct Connect to connect your on-premises system through a VPC Endpoint interface that uses a private IP address. This type of connection is shown in the [second diagram](#).

Steps


1. Create an Interface endpoint configuration using the Amazon VPC console or the command line. [Refer to details about using AWS PrivateLink for Amazon S3.](#)
2. Modify the security group configuration that's associated with the BlueXP Connector. You must change the policy to "Custom" (from "Full Access"), and you must [add the S3 permissions from the backup policy](#) as shown earlier.



If you're using port 80 (HTTP) for communication to the private endpoint, you're all set. You can enable BlueXP backup and recovery on the cluster now.

If you're using port 443 (HTTPS) for communication to the private endpoint, you must copy the certificate from the VPC S3 endpoint and add it to your ONTAP cluster, as shown in the next 4 steps.

3. Obtain the DNS name of the endpoint from the AWS Console.



- Obtain the certificate from the VPC S3 endpoint. You do this by [logging into the VM that hosts the BlueXP Connector](#) and running the following command. When entering the DNS name of the endpoint, add “bucket” to the beginning, replacing the “*”:

```
[ec2-user@ip-10-160-4-68 ~]$ openssl s_client -connect bucket.vpce-0ff5c15df7e00fbab-yxs7lt8v.s3.us-west-2.vpce.amazonaws.com:443 -showcerts
```

- From the output of this command, copy the data for the S3 certificate (all data between, and including, the BEGIN / END CERTIFICATE tags):

```
Certificate chain
0 s:/CN=s3.us-west-2.amazonaws.com`
  i:/C=US/O=Amazon/OU=Server CA 1B/CN=Amazon
-----BEGIN CERTIFICATE-----
MIIM6zCCC9OgAwIBAgIQA7MGJ4FaD8uL0KR3oltTANBgkqhkiG9w0BAQsFADBG
...
...
GqvboZ/oO2NWLLFCqI+xmKLCmiPrZy+/6Af+HH2mLCM4EsI2b+IpBmPkriWnnxo=
-----END CERTIFICATE-----
```

- Log into the ONTAP cluster CLI and apply the certificate you copied using the following command (substitute your own storage VM name):

```
cluster1::> security certificate install -vserver cluster1 -type server-  
ca  
Please enter Certificate: Press <Enter> when done
```

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises working environment.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future working environments.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:
 - From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.

If the Amazon S3 destination for your backups exists as a working environment on the Canvas, you can drag the ONTAP cluster onto the Amazon S3 object storage.

 - Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions**  icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.
2. Continue with the following options:
 - If you already have a BlueXP Connector, you're all set. Just select **Next**.
 - If you don't already have a BlueXP Connector, the **Add a Connector** option appears. Refer to [Prepare your BlueXP Connector](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a working environment. See how to [activate backup for additional volumes in the working environment](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row. (☒ Volume Name).
 - To back up individual volumes, check the box for each volume (☒ Volume_1).
2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture
- Local Snapshot policy
- Replication target and policy



If the volumes you choose have different Snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots:** If you are performing replication or back up to object storage, local Snapshots must be created.
 - **Replication:** Creates replicated volumes on another ONTAP storage system.
 - **Backup:** Backs up volumes to object storage.
2. **Architecture:** If you chose replication and backup, choose one of the following flows of information:
 - **Cascading:** Information flows from the primary to the secondary to object storage and from the secondary to object storage.
 - **Fan out:** Information flows from the primary to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing Snapshot policy or create a policy.



To create a custom policy before activating the Snapshot, refer to [Create a policy](#).

4. To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
 - For backup-to-object policies, set the DataLock and Ransomware Protection settings. For details on DataLock and Ransomware Protection, refer to [Backup-to-object policy settings](#).
- Select **Create**.

5. **Replication**: Set the following options:

- **Replication target**: Select the destination working environment and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy**: Choose an existing replication policy or create a policy.



To create a custom policy before activating the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

6. **Back up to Object**: If you selected **Backup**, set the following options:

- **Provider**: Select **Amazon Web Services**.
- **Provider settings**: Enter the provider details and AWS region where the backups will be stored.

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

- **Bucket**: Either choose an existing S3 bucket or create a new one. Refer to [Add S3 buckets](#).
- **Encryption key**: If you created a new S3 bucket, enter encryption key information given to you from the provider. Choose 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.



If you chose an existing bucket, encryption information is already available, so you don't need to enter it now.

- **Networking**: Choose the IPspace, and whether you'll be using a Private Endpoint. Private Endpoint is disabled by default.
 - 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 AWS PrivateLink for Amazon S3](#).
- **Backup policy**: Select an existing backup policy or create a policy.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.
- **Export existing Snapshot copies to object storage as backup copies:** If there are any local Snapshot copies for volumes in this working environment that match the backup schedule label you just selected for this working environment (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

7. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

BlueXP backup and recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary data contained in Snapshot copies.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage volume.

The S3 bucket is created 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.

You can also monitor the status of backup and restore jobs using the [Job Monitoring panel](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future working environments.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.

- You can [manage cluster-level backup settings](#). This includes changing the storage keys ONTAP uses to access cloud storage, changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to a Cloud Volumes ONTAP system in AWS, or to an on-premises ONTAP system.

Back up on-premises ONTAP data to Azure Blob storage

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



"On-premises ONTAP systems" include FAS, AFF, and ONTAP Select systems.

Quick start

Get started quickly by following these steps. Details for each step are provided in the following sections in this topic.

1

Identify the connection method you'll use

Choose whether you'll connect your on-premises ONTAP cluster directly to Azure over the public internet, or whether you'll use a VPN or Azure ExpressRoute and route traffic through a private VPC Endpoint interface to Azure.

[Identify the connection method.](#)

2

Prepare your BlueXP Connector

If you already have a Connector deployed in your Azure VNet or on your premises, then you're all set. If not, then you'll need to create a BlueXP Connector to back up ONTAP data to Azure Blob storage. You'll also need to customize network settings for the Connector so that it can connect to Azure.

[Learn how to create a Connector and how to define required network settings.](#)

3

Verify license requirements

You'll need to check license requirements for both Azure and BlueXP.

Refer to [Verify license requirements](#).

4

Prepare your ONTAP clusters

Discover your ONTAP clusters in BlueXP, verify that the clusters meet minimum requirements, and customize network settings so the clusters can connect to Azure.

[Learn how to get your ONTAP clusters ready.](#)

5

Prepare Azure Blob as your backup target

Set up permissions for the Connector to create and manage the Azure bucket. You'll also need to set up permissions for the on-premises ONTAP cluster so it can read and write data to the Azure bucket.

Optionally, you can set up your own custom-managed keys for data encryption instead of using the default Azure encryption keys. [Learn how to get your Azure environment ready to receive ONTAP backups.](#)

6

Activate backups on your ONTAP volumes

Select the working environment and click **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel. Then follow the setup wizard to select the replication and backup policies that you'll use and the volumes you want to back up.

[Activate backups on your ONTAP volumes.](#)

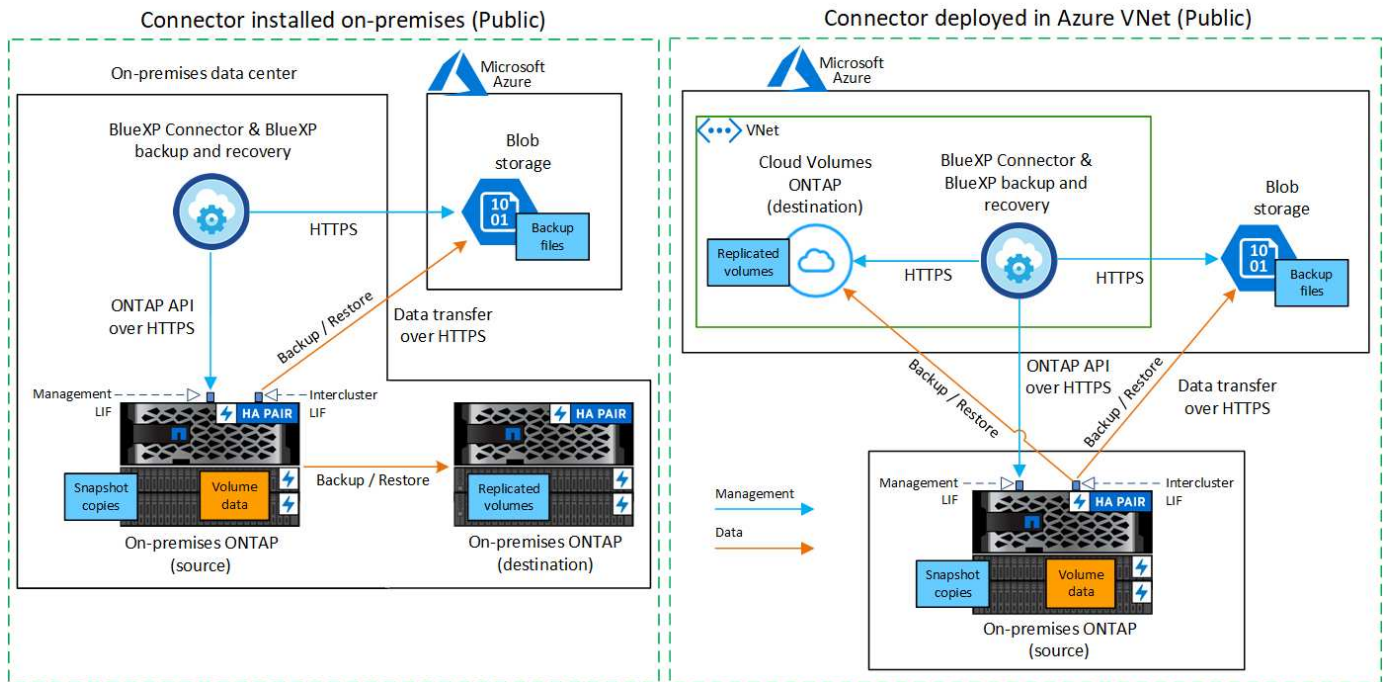
Identify the connection method

Choose which of the two connection methods you will use when configuring backups from on-premises ONTAP systems to Azure Blob.


- **Public connection** - Directly connect the ONTAP system to Azure Blob storage using a public Azure endpoint.
- **Private connection** - Use a VPN or ExpressRoute and route traffic through a VNet Private Endpoint that uses a private IP address.

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.

The following diagram shows the **public connection** method and the connections that you need to prepare between the components. You can use a Connector that you've installed on your premises, or a Connector that you've deployed in the Azure VNet.



The following diagram shows the **private connection** method and the connections that you need to prepare between the components. You can use a Connector that you've installed on your premises, or a Connector that you've deployed in the Azure VNet.



Prepare your BlueXP Connector

The BlueXP Connector is the main software for BlueXP functionality. A Connector is required to back up and restore your ONTAP data.

Create or switch Connectors

If you already have a Connector deployed in your Azure VNet or on your premises, then you're all set.

If not, then you'll need to create a Connector in one of those locations to back up ONTAP data to Azure Blob storage. You can't use a Connector that's deployed in another cloud provider.

- [Learn about Connectors](#)
- [Install a Connector in Azure](#)
- [Install a Connector in your premises](#)
- [Install a Connector in an Azure Government region](#)

BlueXP backup and recovery is supported in Azure Government regions when the Connector is deployed in the cloud - not when it's installed in your premises. Additionally, you must deploy the Connector from the Azure Marketplace. You can't deploy the Connector in a Government region from the BlueXP SaaS website.

Prepare 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 the BlueXP backup and recovery service and to your Blob object storage ([see the list of endpoints](#))
 - An HTTPS connection over port 443 to your ONTAP cluster management LIF
 - In order for the BlueXP backup and recovery Search & Restore functionality to work, port 1433 must be open for communication between the Connector and the Azure Synapse SQL services.
 - Additional inbound security group rules are required for Azure and Azure Government deployments. See [Rules for the Connector in Azure](#) for details.
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 (a **private** connection).

Verify or add permissions to the Connector

To use the BlueXP backup and recovery Search & Restore functionality, you need to have specific permissions in the role for the Connector so that it can access the Azure Synapse Workspace and Data Lake Storage Account. See the permissions below, and follow the steps if you need to modify the policy.

Before you start

You must register the Azure Synapse Analytics Resource Provider (called "Microsoft.Synapse") with your Subscription. [See how to register this resource provider for your subscription](#). You must be the Subscription **Owner** or **Contributor** to register the resource provider.

Steps

1. Identify the role assigned to the Connector virtual machine:
 - a. In the Azure portal, open the Virtual machines service.
 - b. Select the Connector virtual machine.
 - c. Under **Settings**, select **Identity**.
 - d. Select **Azure role assignments**.

- e. Make note of the custom role assigned to the Connector virtual machine.
2. Update the custom role:
- a. In the Azure portal, open your Azure subscription.
 - b. Select **Access control (IAM) > Roles**.
 - c. Select the ellipsis (...) for the custom role and then select **Edit**.
 - d. Select **JSON** and add the following permissions:

```

"Microsoft.Storage/storageAccounts/listkeys/action",
"Microsoft.Storage/storageAccounts/read",
"Microsoft.Storage/storageAccounts/write",
"Microsoft.Storage/storageAccounts/blobServices/containers/read",
"Microsoft.Storage/storageAccounts/listAccountSas/action",
"Microsoft.KeyVault/vaults/read",
"Microsoft.KeyVault/vaults/accessPolicies/write",
"Microsoft.Network/networkInterfaces/read",
"Microsoft.Resources/subscriptions/locations/read",
"Microsoft.Network/virtualNetworks/read",
"Microsoft.Network/virtualNetworks/subnets/read",
"Microsoft.Resources/subscriptions/resourceGroups/read",
"Microsoft.Resources/subscriptions/resourcegroups/resources/read",
"Microsoft.Resources/subscriptions/resourceGroups/write",
"Microsoft.Authorization/locks/*",
"Microsoft.Network/privateEndpoints/write",
"Microsoft.Network/privateEndpoints/read",
"Microsoft.Network/privateDnsZones/virtualNetworkLinks/write",
"Microsoft.Network/virtualNetworks/join/action",
"Microsoft.Network/privateDnsZones/A/write",
"Microsoft.Network/privateDnsZones/read",
"Microsoft.Network/privateDnsZones/virtualNetworkLinks/read",
"Microsoft.Network/networkInterfaces/delete",
"Microsoft.Network/networkSecurityGroups/delete",
"Microsoft.Resources/deployments/delete",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
"Microsoft.Synapse/workspaces/write",
"Microsoft.Synapse/workspaces/read",
"Microsoft.Synapse/workspaces/delete",
"Microsoft.Synapse/register/action",
"Microsoft.Synapse/checkNameAvailability/action",
"Microsoft.Synapse/workspaces/operationStatuses/read",
"Microsoft.Synapse/workspaces/firewallRules/read",
"Microsoft.Synapse/workspaces/replaceAllIpFirewallRules/action",
"Microsoft.Synapse/workspaces/operationResults/read",
"Microsoft.Synapse/workspaces/privateEndpointConnectionsApproval/
action"

```

[View the full JSON format for the policy](#)

e. Select **Review + update** and then select **Update**.

Verify license requirements

You'll need to verify license requirements for both Azure and BlueXP:

- Before you can activate BlueXP backup and recovery for your cluster, you'll need to either subscribe to a pay-as-you-go (PAYGO) BlueXP Marketplace offering from Azure, or purchase and activate a BlueXP backup and recovery BYOL license from NetApp. These licenses are for your account and can be used across multiple systems.
 - For BlueXP backup and recovery PAYGO licensing, you'll need a subscription to the [NetApp BlueXP offering from the Azure Marketplace](#). Billing for BlueXP backup and recovery is done through this subscription.
 - For BlueXP backup and recovery BYOL licensing, 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](#).
- You need to have an Azure subscription for the object storage space where your backups will be located.

Supported regions

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

Prepare your ONTAP clusters

Unresolved directive in task-backup-onprem-to-azure.adoc - include::.../_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must configure the following requirements on the system that connects to object storage.

- For a fan-out backup architecture, configure the following settings on the *primary* system.
- For a cascaded backup architecture, configure the following settings on the *secondary* system.

The following ONTAP cluster networking requirements are needed:

- 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 BlueXP backup and recovery, 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](#).
- 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 BlueXP backup and recovery 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).

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in task-backup-onprem-to-azure.adoc - include::.../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

Prepare Azure Blob as your backup target

1. You can use 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. [Learn how to use your own keys](#).

Note that Backup and recovery supports *Azure access policies* as the permission model. The *Azure role-based access control* (Azure RBAC) permission model is not currently supported.

2. 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. [Refer to details about using a Private Endpoint](#).

Create your Azure Blob storage account

By default, the service creates storage accounts for you. If you want to use your own storage accounts, you can create them before you start the backup activation wizard and then select those storage accounts in the wizard.

[Learn more about creating your own storage accounts](#).

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises working environment.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)


You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future working environments.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.



If the Azure destination for your backups exists as a working environment on the Canvas, you can drag the ONTAP cluster onto the Azure Blob object storage.

- Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions** ... icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a BlueXP Connector, you're all set. Just select **Next**.
- If you don't already have a BlueXP Connector, the **Add a Connector** option appears. Refer to [Prepare your BlueXP Connector](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a working environment. See how to [activate backup for additional volumes in the working environment](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.

- Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
- After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then

check the box in the title row. ( Volume Name).

- To back up individual volumes, check the box for each volume ( Volume_1).

2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture
- Local Snapshot policy
- Replication target and policy



If the volumes you choose have different Snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots**: If you are performing replication or back up to object storage, local Snapshots must be created.
 - **Replication**: Creates replicated volumes on another ONTAP storage system.
 - **Backup**: Backs up volumes to object storage.
2. **Architecture**: If you chose replication and backup, choose one of the following flows of information:
 - **Cascading**: Information flows from the primary to the secondary, and from secondary to object storage.
 - **Fan out**: Information flows from the primary to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot**: Choose an existing Snapshot policy or create a new one.



To create a custom policy before activating the Snapshot, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

4. **Replication**: Set the following options:

- **Replication target**: Select the destination working environment and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy**: Choose an existing replication policy or create a new one.



To create a custom policy before activating the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **Microsoft Azure**.
- **Provider settings:** Enter the provider details and region where the backups will be stored.

Either create a new storage account or select an existing one.

Either create your own resource group that manages the Blob container or select the resource group type and group.



If you want to protect your backup files from being modified or deleted, ensure that the storage account was created with immutable storage enabled using a 30-day retention period.



If you want to tier older backup files to Azure Archive Storage for further cost optimization, ensure that the storage account has the appropriate Lifecycle rule.

- **Encryption key:** If you created a new Azure storage account, enter encryption key information given to you from the provider. Choose whether you'll use the default Azure encryption keys, or choose your own customer-managed keys from your Azure account, to manage encryption of your data.

If you choose to use your own customer-managed keys, enter the key vault and key information.



If you chose an existing Microsoft storage account, encryption information is already available, so you don't need to enter it now.

- **Networking:** Choose the IPspace, and whether you'll be using a Private Endpoint. Private Endpoint is disabled by default.
 - 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 Azure private endpoint that you have previously configured. [Learn about using an Azure private endpoint](#).
- **Backup policy:** Select an existing Backup to object storage policy or create a new one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- For backup-to-object policies, set the DataLock and Ransomware Protection settings. For details

on DataLock and Ransomware Protection, refer to [Backup-to-object policy settings](#).

- Select **Create**.

- **Export existing Snapshot copies to object storage as backup copies:** If there are any local Snapshot copies for volumes in this working environment that match the backup schedule label you just selected for this working environment (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

BlueXP backup and recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary storage system data contained in Snapshot copies.

A replicated volume is created in the destination cluster that will be synchronized with the primary volume.

A Blob storage account is created in the resource group you entered, and the backup files are stored there. The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

You can also monitor the status of backup and restore jobs using the [Job Monitoring panel](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future working environments.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.
- You can [manage cluster-level backup settings](#). This includes changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to a Cloud Volumes ONTAP system in Azure, or to an on-premises ONTAP system.

Back up on-premises ONTAP data to Google Cloud Storage

Complete a few steps to get started backing up volume data from your on-premises primary ONTAP systems to a secondary storage system and to Google Cloud Storage.



"On-premises ONTAP systems" include FAS, AFF, and ONTAP Select systems.

Quick start

Get started quickly by following these steps. Details for each step are provided in the following sections in this topic.

1

Identify the connection method you'll use

Choose whether you'll connect your on-premises ONTAP cluster directly to Google Cloud Storage over the public internet, or whether you'll use a VPN or Google Cloud Interconnect and route traffic through a private Google Access interface that uses a private IP address.

2

Prepare your BlueXP Connector

If you already have a Connector deployed in your Google Cloud Platform VPC, then you're all set. If not, then you'll need to create a BlueXP Connector to back up ONTAP data to Google Cloud storage. You'll also need to customize network settings for the Connector so that it can connect to Google Cloud.

3

Prepare networking for the Connector

Ensure that the Connector has the required networking connections.

4

Verify license requirements

You'll need to check license requirements for both Google Cloud and BlueXP.

5

Prepare your ONTAP clusters

Discover your ONTAP clusters in BlueXP, verify that the clusters meet minimum requirements, and customize network settings so the clusters can connect to Google Cloud.

6

Prepare Google Cloud as your backup target

Set up permissions for the Connector to create and manage the Google Cloud bucket. You'll also need to set up permissions for the on-premises ONTAP cluster so it can read and write data to the Google Cloud bucket.

Optionally, you can set up your own custom-managed keys for data encryption instead of using the default Google Cloud encryption keys.

7

Activate backups on your ONTAP volumes

Select the working environment and click **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel. Then follow the setup wizard to select the replication and backup policies that you'll use and the volumes you want to back up.

Identify the connection method


Choose which of the two connection methods you will use when configuring backups from on-premises ONTAP systems to Google Cloud Storage.

- **Public connection** - Directly connect the ONTAP system to Google Cloud Storage using a public Google endpoint.
- **Private connection** - Use a VPN or Google Cloud Interconnect and route traffic through a Private Google Access interface that uses a private IP address.

Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.

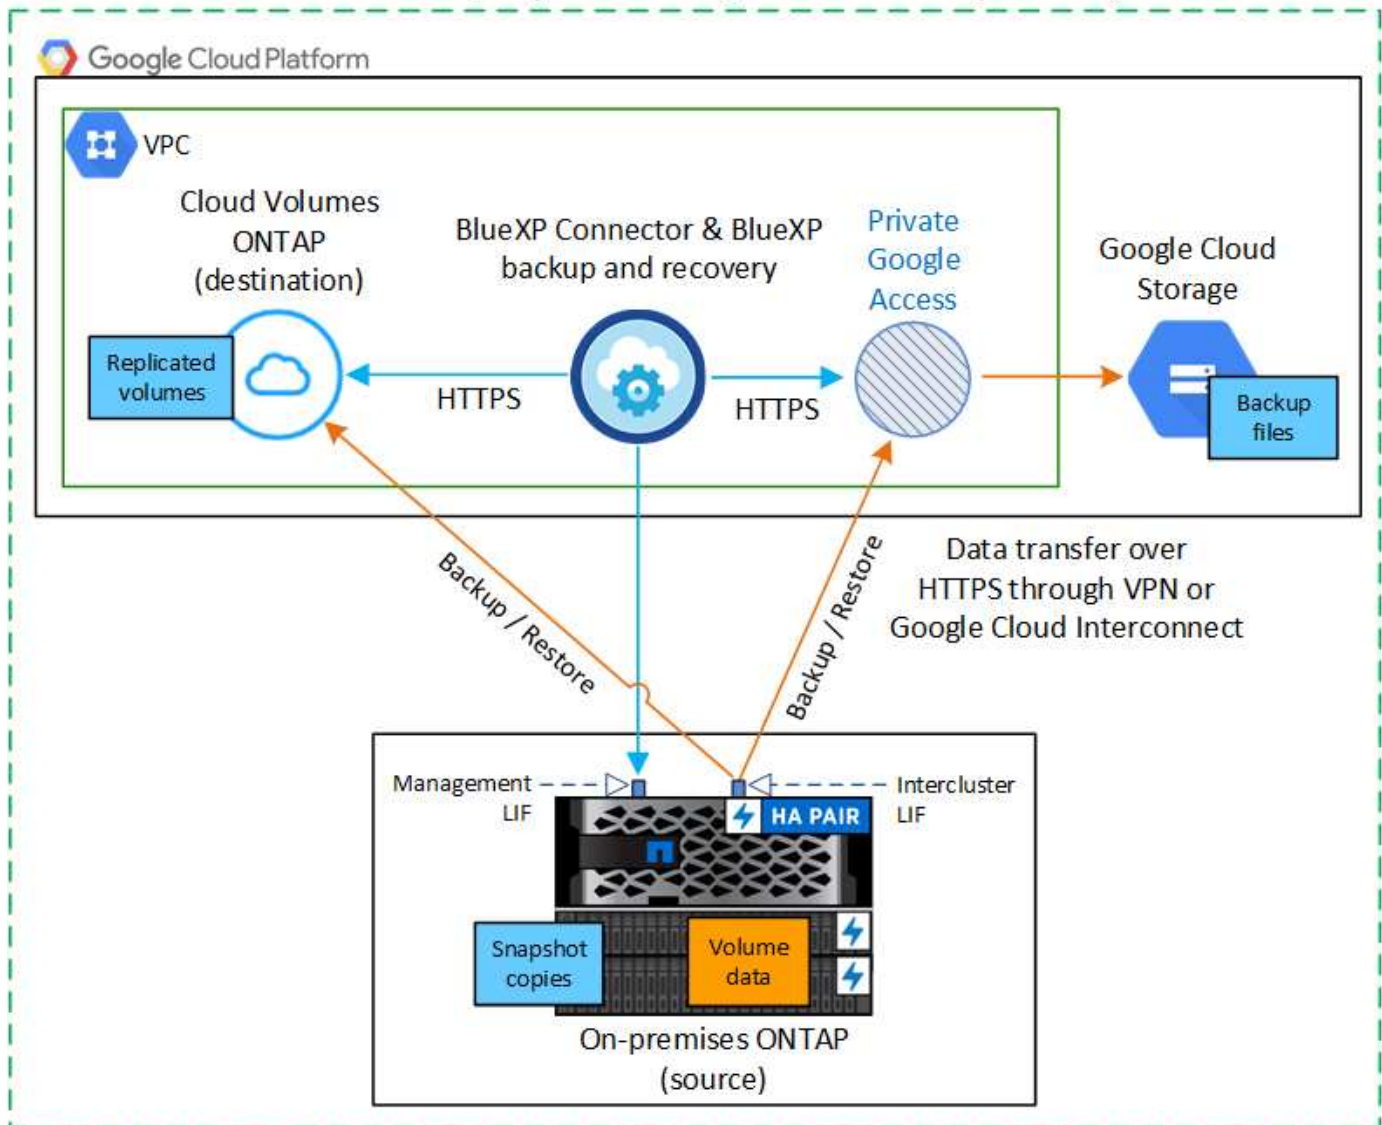
The following diagram shows the **public connection** method and the connections that you need to prepare between the components. The Connector must be deployed in the Google Cloud Platform VPC.

Connector deployed in Google Cloud VPC (Public)



The following diagram shows the **private connection** method and the connections that you need to prepare between the components. The Connector must be deployed in the Google Cloud Platform VPC.

Connector deployed in Google Cloud VPC (Private)



Prepare your BlueXP Connector

The BlueXP Connector is the main software for BlueXP functionality. A Connector is required to back up and restore your ONTAP data.

Create or switch Connectors

If you already have a Connector deployed in your Google Cloud Platform VPC, then you're all set.

If not, then you'll need to create a Connector in that location to back up ONTAP data to Google Cloud Storage. You can't use a Connector that's deployed in another cloud provider, or on-premises.

- [Learn about Connectors](#)
- [Install a Connector in GCP](#)

Prepare 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 the BlueXP backup and recovery service and to your Google Cloud storage ([see the list of endpoints](#))
 - An HTTPS connection over port 443 to your ONTAP cluster management LIF
2. Enable Private Google Access (or Private Service Connect) on the subnet where you plan to deploy the Connector. [Private Google Access](#) or [Private Service Connect](#) are 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 (a **private** connection).

Follow the Google instructions for setting up these Private access options. Make sure your DNS servers have been configured to point `www.googleapis.com` and `storage.googleapis.com` to the correct internal (private) IP addresses.

Verify or add permissions to the Connector

To use the BlueXP backup and recovery "Search & Restore" functionality, you need to have specific permissions in the role for the Connector so that it can access the Google Cloud BigQuery service. Review the permissions below, and follow the steps if you need to modify the policy.

Steps

1. In the [Google Cloud Console](#), go to the **Roles** page.
2. Using the drop-down list at the top of the page, select the project or organization that contains the role that you want to edit.
3. Select a custom role.
4. Select **Edit Role** to update the role's permissions.
5. Select **Add Permissions** to add the following new permissions to the role.

```
bigquery.jobs.get  
bigquery.jobs.list  
bigquery.jobs.listAll  
bigquery.datasets.create  
bigquery.datasets.get  
bigquery.jobs.create  
bigquery.tables.get  
bigquery.tables.getData  
bigquery.tables.list  
bigquery.tables.create
```

6. Select **Update** to save the edited role.

Verify license requirements

- Before you can activate BlueXP backup and recovery for your cluster, you'll need to either subscribe to a pay-as-you-go (PAYGO) BlueXP Marketplace offering from Google, or purchase and activate a BlueXP backup and recovery BYOL license from NetApp. These licenses are for your account and can be used across multiple systems.
 - For BlueXP backup and recovery PAYGO licensing, you'll need a subscription to the [NetApp BlueXP offering from the Google Marketplace](#). Billing for BlueXP backup and recovery is done through this subscription.
 - For BlueXP backup and recovery BYOL licensing, 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](#).
- You need to have a Google subscription for the object storage space where your backups will be located.

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 backups will be stored when you set up the service.

Prepare your ONTAP clusters

Unresolved directive in task-backup-onprem-to-gcp.adoc - include::.../_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must configure the following requirements on the system that connects to object storage.

- For a fan-out backup architecture, configure the following settings on the *primary* system.
- For a cascaded backup architecture, configure the following settings on the *secondary* system.

The following ONTAP cluster networking requirements are needed:

- 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 BlueXP backup and recovery, 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.](#)

If you're using Private Google Access or Private Service Connect, make sure your DNS servers have been configured to point `storage.googleapis.com` to the correct internal (private) IP address.

- 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 BlueXP backup and recovery 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).

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in task-backup-onprem-to-gcp.adoc - include::.../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

Prepare Google Cloud Storage as your backup target

Preparing Google Cloud Storage as your backup target involves the following steps:

- Set up permissions.
- (Optional) Create your own buckets. (The service will create buckets for you if you want.)
- (Optional) Set up customer-managed keys for data encryption

Set up permissions

You need to provide storage access keys for a service account that has specific permissions using a custom role. A service account enables BlueXP backup and recovery 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. In the [Google Cloud Console](#), go to the **Roles** page.
2. [Create a new role](#) with the following permissions:

```
storage.buckets.create
storage.buckets.delete
storage.buckets.get
storage.buckets.list
storage.buckets.update
storage.buckets.getIamPolicy
storage.multipartUploads.create
storage.objects.create
storage.objects.delete
storage.objects.get
storage.objects.list
storage.objects.update
```


3. In the Google Cloud console, [go to the Service accounts page](#).
4. Select your Cloud project.
5. Select **Create service account** and provide the required information:
 - a. **Service account details**: Enter a name and description.
 - b. **Grant this service account access to project**: Select the custom role that you just created.
 - c. Select **Done**.
6. Go to [GCP Storage Settings](#) and create access keys for the service account:
 - a. Select a project, and select **Interoperability**. If you haven't already done so, select **Enable interoperability access**.
 - b. Under **Access keys for service accounts**, select **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 BlueXP backup and recovery later when you configure the backup service.

Create your own buckets

By default, the service creates buckets for you. Or, if you want to use your own buckets, you can create them before you start the backup activation wizard and then select those buckets in the wizard.

[Learn more about creating your own buckets](#).

Set up customer-managed encryption keys (CMEK) for data encryption

You can use your own customer-managed keys for data encryption instead of using the default Google-managed encryption keys. Both cross-region and cross-project keys are supported, so you can choose a project for a bucket that is different than the project of the CMEK key.

If you're planning to use your own customer-managed keys:

- You'll need to have the Key Ring and the Key Name so you can add this information in the activation wizard. [Learn more about customer-managed encryption keys](#).
- You'll need to verify that these required permissions are included in the role for the Connector:

```
cloudkms.cryptoKeys.get
cloudkms.cryptoKeys.getIamPolicy
cloudkms.cryptoKeys.list
cloudkms.cryptoKeys.setIamPolicy
cloudkms.keyRings.get
cloudkms.keyRings.getIamPolicy
cloudkms.keyRings.list
cloudkms.keyRings.setIamPolicy
```

- You'll need to verify that the Google "Cloud Key Management Service (KMS)" API is enabled in your project. See the [Google Cloud documentation: Enabling APIs](#) for details.

CMEK considerations:

- Both HSM (Hardware-backed) and Software-generated keys are supported.
- Both newly created or imported Cloud KMS keys are supported.
- Only regional keys are supported, global keys are not supported.
- Currently, only the "Symmetric encrypt/decrypt" purpose is supported.
- The service agent associated with the Storage Account is assigned the "CryptoKey Encrypter/Decrypter (roles/cloudkms.cryptoKeyEncrypterDecrypter)" IAM role by BlueXP backup and recovery.

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises working environment.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future working environments.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:
 - From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.



If the Google Cloud Storage destination for your backups exists as a working environment on the Canvas, you can drag the ONTAP cluster onto the Google Cloud object storage.

- Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions** ... icon and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:
 - If you already have a BlueXP Connector, you're all set. Just select **Next**.
 - If you don't already have a BlueXP Connector, the **Add a Connector** option appears. Refer to [Prepare](#)

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a working environment. See how to [activate backup for additional volumes in the working environment](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row. (☒ Volume Name).
 - To back up individual volumes, check the box for each volume (☒ Volume_1).
2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture
- Local Snapshot policy
- Replication target and policy



If the volumes you choose have different Snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots:** If you are performing replication or back up to object storage, local Snapshots must be created.
 - **Replication:** Creates replicated volumes on another ONTAP storage system.

- **Backup:** Backs up volumes to object storage.

2. **Architecture:** If you chose replication and backup, choose one of the following flows of information:

- **Cascading:** Information flows from the primary to the secondary and from the secondary to object storage.
- **Fan out:** Information flows from the primary to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing Snapshot policy or create a new one.



To create a custom policy before activating the Snapshot, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

4. **Replication:** Set the following options:

- **Replication target:** Select the destination working environment and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
- **Replication policy:** Choose an existing replication policy or create a new one.



To create a custom policy before you activate the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **Google Cloud**.
- **Provider settings:** Enter the provider details and region where the backups will be stored.

Either create a new bucket or select one that you've already created.



If you want to tier older backup files to Google Cloud Archive storage for further cost optimization, ensure that the bucket has the appropriate Lifecycle rule.

Enter the Google Cloud access key and secret key.

- **Encryption key:** If you created a new Google Cloud storage account, enter encryption key information given to you from the provider. Choose whether you'll use the default Google Cloud encryption keys, or choose your own customer-managed keys from your Google Cloud account, to manage encryption of your data.



If you chose an existing Google Cloud storage account, encryption information is already available, so you don't need to enter it now.

If you choose to use your own customer-managed keys, enter the key ring and key name. [Learn more about customer-managed encryption keys.](#)

- **Networking:** Choose the IPspace.

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.

- **Backup policy:** Select an existing Backup to object storage policy or create a new one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
 - Select up to 5 schedules, typically of different frequencies.
 - Select **Create**.
- **Export existing Snapshot copies to object storage as backup copies:** If there are any local Snapshot copies for volumes in this working environment that match the backup schedule label you just selected for this working environment (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

BlueXP backup and recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the primary storage system data. Subsequent transfers contain differential copies of the primary storage system data contained in Snapshot copies.

A replicated volume is created in the destination cluster that will be synchronized with the source volume.

A Google Cloud Storage bucket is created automatically in the service account indicated by the Google 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.

You can also monitor the status of backup and restore jobs using the [Job Monitoring panel](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future working environments.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.
- You can [manage cluster-level backup settings](#). This includes changing the storage keys ONTAP uses to access cloud storage, changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to a Cloud Volumes ONTAP system in Google, or to an on-premises ONTAP system.

Back up on-premises ONTAP data to ONTAP S3

Complete a few steps to get started backing up volume data from your primary on-premises ONTAP systems. You can send backups to a secondary ONTAP storage system (a replicated volume) or to a bucket on an ONTAP system configured as an S3 server (a backup file), or both.

The primary on-premises ONTAP system can be a FAS, AFF, or ONTAP Select system. The secondary ONTAP system can be an on-premises ONTAP or Cloud Volumes ONTAP system. The object storage can be on an on-premises ONTAP system or a Cloud Volumes ONTAP system on which you have enabled a Simple Storage Service (S3) object storage server.

Quick start

Get started quickly by following these steps. Details for each step are provided in the following sections in this topic.

1

Identify the connection method you'll use

Review how you'll connect your primary on-premises ONTAP cluster to the secondary ONTAP cluster for replication and to the ONTAP cluster configured as an S3 server for backup to object storage.

[Identify the connection method.](#)

2

Prepare your BlueXP Connector

If you've already deployed a BlueXP Connector, then you're all set. If not, then you'll need to create a BlueXP Connector to back up ONTAP data to ONTAP S3. You'll also need to customize network settings for the Connector so that it can connect to ONTAP S3.

[Learn how to create a Connector and how to define required network settings.](#)

3

Verify license requirements

You'll need to check license requirements for your ONTAP systems and for BlueXP backup and recovery.

[Verify license requirements.](#)

4

Prepare your ONTAP clusters

Discover your primary and secondary ONTAP clusters in BlueXP, verify that the clusters meet minimum requirements, and customize network settings so the clusters can connect to ONTAP S3 object storage.

[Learn how to get your ONTAP clusters ready.](#)

5

Prepare ONTAP S3 as your backup target

Set up permissions for the Connector so it can manage the ONTAP S3 bucket. You'll also need to set up permissions for the source on-premises ONTAP cluster so that it can read and write data to the ONTAP S3 bucket.

[Learn how to get your ONTAP S3 environment ready to receive ONTAP backups.](#)

6

Activate backups on your ONTAP volumes

Select the primary working environment and click **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel. Then follow the setup wizard to select the volumes you want to back up, and the Snapshot, replication, and backup to object policies that you'll use.


[Activate backups on your ONTAP volumes.](#)

Identify the connection method

There are many configurations in which you can create backups to an S3 bucket on an ONTAP system. Two scenarios are shown below.

The following image shows each component when backing up a primary on-premises ONTAP system to an on-premises ONTAP system configured for S3 and the connections that you need to prepare between them. It also shows a connection to a secondary ONTAP system in the same on-premises location to replicate volumes.


Connector installed on-premises (Public)



When the Connector and primary on-premises ONTAP system are installed in an on-premises location without internet access (a "private" mode deployment), the ONTAP S3 system must be located in the same on-premises data center.

The following image shows each component when backing up a primary on-premises ONTAP system to a Cloud Volumes ONTAP system configured for S3 and the connections that you need to prepare between them. It also shows a connection to a secondary Cloud Volumes ONTAP system in the same cloud provider environment to replicate volumes.

Connector deployed in cloud (Public)



In this scenario the Connector should be deployed in the same cloud provider environment in which the Cloud Volumes ONTAP systems are deployed.

Prepare your BlueXP Connector

The BlueXP Connector is the main software for BlueXP functionality. A Connector is required to back up and restore your ONTAP data.

Create or switch Connectors

When you back up data to ONTAP S3, a BlueXP Connector must be available on your premises or in the cloud. You'll either need to install a new Connector or make sure that the currently selected Connector resides in one of these locations. The on-premises Connector can be installed in a site with or without internet access.

- [Learn about Connectors](#)
- [Install the Connector in your cloud environment](#)

- [Installing the Connector on a Linux host with internet access](#)
- [Installing the Connector on a Linux host without internet access](#)
- [Switching between Connectors](#)

Prepare Connector networking requirements

Ensure that the network where the Connector is installed enables the following connections:

- An HTTPS connection over port 443 to the ONTAP S3 server
- An HTTPS connection over port 443 to your source ONTAP cluster management LIF
- An outbound internet connection over port 443 to BlueXP backup and recovery (not required when the Connector is installed in a "dark" site)

Private mode (dark site) considerations

BlueXP backup and recovery functionality is built into the BlueXP Connector. When it is installed in private mode, you'll need to update the Connector software periodically to get access to new features. Check the [BlueXP backup and recovery What's New](#) to see the new features in each BlueXP backup and recovery release. When you want to use the new features, follow the steps to [upgrade the Connector software](#).

When you use BlueXP backup and recovery in a standard SaaS environment, the BlueXP backup and recovery configuration data is backed up to the cloud. When you use BlueXP backup and recovery in a site with no internet access, the BlueXP backup and recovery configuration data is backed up to the ONTAP S3 bucket where your backups are being stored. If you ever have a Connector failure in your private mode site, you can [restore the BlueXP backup and recovery data to a new Connector](#).

Verify license requirements

Before you can activate BlueXP backup and recovery for your cluster, you'll need to purchase and activate a BlueXP backup and recovery BYOL license from NetApp. The license is for backup and restore to object storage - no license is needed to create Snapshot copies or replicated volumes. 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 supported when backing up files to ONTAP S3.

Prepare your ONTAP clusters

Unresolved directive in task-backup-onprem-to-ontap-s3.adoc - include::.../_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must ensure that the following requirements are met on the system that connects to object storage.



- When you use a fan-out backup architecture, the settings must be configured on the *primary* storage system.
- When you use a cascaded backup architecture, the settings must be configured on the *secondary* storage system.

[Learn more about the types of backup architecture.](#)

The following ONTAP cluster networking requirements are needed:

- The ONTAP cluster initiates an HTTPS connection over a user-specified port from the intercluster LIF to the ONTAP S3 server 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.
- 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 BlueXP backup and recovery, 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 (not required when the Connector is installed in a "dark" site).
- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- If you are using a different IPspace than Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow BlueXP backup and recovery 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).

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in task-backup-onprem-to-ontap-s3.adoc - include::.../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

Prepare ONTAP S3 as your backup target

You must enable a Simple Storage Service (S3) object storage server in the ONTAP cluster that you plan to use for object storage backups. See the [ONTAP S3 documentation](#) for details.

Note: You can discover this cluster to the BlueXP Canvas, but it is not identified as being an S3 object storage server, and you can't drag and drop a source working environment onto this S3 working environment to initiate backup activation.

This ONTAP system must meet the following requirements.

Supported ONTAP versions

ONTAP 9.8 and later is required for on-premises ONTAP systems.

ONTAP 9.9.1 and later is required for Cloud Volumes ONTAP systems.

S3 credentials

You must have created an S3 user to control access to your ONTAP S3 storage. [See the ONTAP S3 docs for details.](#)

When you set up backup to ONTAP S3, the backup wizard prompts you for an S3 access key and secret key for a user account. The user account enables BlueXP backup and recovery to authenticate and access the ONTAP S3 buckets used to store backups. The keys are required so that ONTAP S3 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"
```

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises working environment.

A wizard takes you through the following major steps:

- Select the volumes that you want to back up
- Define the backup strategy and policies
- Review your selections

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future working environments.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:
 - From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.
 - Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions (...)** option and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replications, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a BlueXP Connector, you're all set. Just select **Next**.
- If you don't have a BlueXP Connector, the **Add a Connector** option appears. Refer to [Prepare your BlueXP Connector](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a working environment. See how to [activate backup for additional volumes in the working environment](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row. (☒ Volume Name).
 - To back up individual volumes, check the box for each volume (☒ Volume_1).
2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves configuring the following options:

- Protection options: Whether you want to implement one or all of the backup options: local Snapshots, replication, and backup to object storage
- Architecture: Whether you want to use a fan-out or cascading backup architecture
- Local Snapshot policy
- Replication target and policy
- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define Backup Strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots**: Creates local Snapshot copies.
 - **Replication**: Creates replicated volumes on another ONTAP storage system.

- **Backup:** Backs up volumes to a bucket on an ONTAP system configured for S3.

2. **Architecture:** If you chose both replication and backup, choose one of the following flows of information:

- **Cascading:** Backup data flows from the primary to the secondary system, and then from the secondary to object storage.
- **Fan out:** Backup data flows from the primary to the secondary system *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing Snapshot policy or create a new one.



If you want to create a custom policy before activating the Snapshot, you can use System Manager or the ONTAP CLI `snapmirror policy create` command. Refer to [ONTAP CLI for snapmirror policy](#).



To create a custom policy using this service before activating the Snapshot, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

4. **Replication:** If you selected **Replication**, set the following options:

- **Replication target:** Select the destination working environment and SVM. Optionally, select the destination aggregate (or aggregates for FlexGroup volumes) and a prefix or suffix that will be added to the replicated volume name.
- **Replication policy:** Choose an existing replication policy or create a new one.

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **ONTAP S3**.
- **Provider settings:** Enter the S3 server FQDN details, port, and the users' access key and secret key.

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

- **Networking:** Choose the IPspace in the source ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access (not required when the Connector is installed in a "dark" site).



Selecting the correct IPspace ensures that BlueXP backup and recovery can set up a connection from ONTAP to your ONTAP S3 object storage.

- **Backup policy:** Select an existing backup policy or create a new one.



You can create a policy with System Manager or the ONTAP CLI. To create a custom policy using the ONTAP CLI `snapmirror policy create` command, refer to [ONTAP CLI for snapmirror policy](#).



To create a custom policy before activating the backup using the UI, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
 - Select up to 5 schedules, typically of different frequencies.
 - For backup-to-object policies, set the DataLock and Ransomware Protection settings. For details on DataLock and Ransomware Protection, refer to [Backup-to-object policy settings](#).
 - Select **Create**.
- **Export existing Snapshot copies to object storage as backup files:** If there are any local Snapshot copies for volumes in this working environment that match the backup schedule label you just selected (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies. If the policies don't match, backups will not be created.
3. Select **Activate Backup**.

Result

BlueXP backup and recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the source data. Subsequent transfers contain differential copies of the primary storage data contained in Snapshot copies.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage volume.

An S3 bucket is created 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.

You can also monitor the status of backup and restore jobs using the [Job Monitoring panel](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future working environments.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.
- You can [manage cluster-level backup settings](#). This includes changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to an on-premises ONTAP system.

Back up on-premises ONTAP data to StorageGRID

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



"On-premises ONTAP systems" include FAS, AFF, and ONTAP Select systems.

Quick start

Get started quickly by following these steps. Details for each step are provided in the following sections in this topic.

1

Identify the connection method you'll use

Review how you'll connect your on-premises ONTAP cluster directly to StorageGRID over the public internet, or whether you'll use a VPN and route traffic through a private VPC Endpoint interface to StorageGRID.

[Identify the connection method.](#)

2

Prepare your BlueXP Connector

If you already have a Connector deployed in your premises, then you're all set. If not, then you'll need to create a BlueXP Connector to back up ONTAP data to StorageGRID. You'll also need to customize network settings for the Connector so that it can connect to StorageGRID.

[Learn how to create a Connector and how to define required network settings.](#)

3

Verify license requirements

You'll need to check license requirements for both StorageGRID and BlueXP.

Refer to [Verify license requirements](#).

4

Prepare your ONTAP clusters

Discover your ONTAP clusters in BlueXP, verify that the clusters meet minimum requirements, and customize network settings so the clusters can connect to StorageGRID.

[Learn how to get your ONTAP clusters ready](#).

5

Prepare StorageGRID as your backup target

Set up permissions for the Connector to create and manage the StorageGRID bucket. You'll also need to set up permissions for the on-premises ONTAP cluster so it can read and write data to the bucket.

Optionally, you can set up your own custom-managed keys for data encryption instead of using the default StorageGRID encryption keys. [Learn how to get your StorageGRID environment ready to receive ONTAP backups](#).

6

Activate backups on your ONTAP volumes


Select the working environment and click **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel. Then follow the setup wizard to select the replication and backup policies that you'll use and the volumes you want to back up.

[Activate backups on your ONTAP volumes](#).

Identify the connection method

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

Optionally, you can connect to a secondary ONTAP system in the same on-premises location to replicate volumes.



When the Connector and on-premises ONTAP system are installed in an on-premises location without internet access (a "dark site"), the StorageGRID system must be located in the same on-premises data center. Archival of older backup files to public cloud is not supported in dark site configurations.

Prepare your BlueXP Connector

The BlueXP Connector is the main software for BlueXP functionality. A Connector is required to back up and restore your ONTAP data.

Create or switch Connectors

When you back up data to StorageGRID, a BlueXP 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-premises. The Connector can be installed in a site with or without internet access.

- [Learn about Connectors](#)
- [Installing the Connector on a Linux host with internet access](#)
- [Installing the Connector on a Linux host without internet access](#)
- [Switching between Connectors](#)

Prepare Connector networking requirements

Ensure that the network where the Connector is installed enables the following connections:

- An HTTPS connection over port 443 to the StorageGRID Gateway Node
- An HTTPS connection over port 443 to your ONTAP cluster management LIF
- An outbound internet connection over port 443 to BlueXP backup and recovery (not required when the Connector is installed in a "dark" site)

Private mode (dark site) considerations

- BlueXP backup and recovery functionality is built into the BlueXP Connector. When it is installed in private mode, you'll need to update the Connector software periodically to get access to new features. Check the [BlueXP backup and recovery What's New](#) to see the new features in each BlueXP backup and recovery release. When you want to use the new features, follow the steps to [upgrade the Connector software](#).

The new version of BlueXP backup and recovery that includes the ability to schedule and create Snapshot copies and replicated volumes, in addition to creating backups to object storage, requires that you are using version 3.9.31 or greater of the BlueXP Connector. So it is recommended that you get this newest release to manage all your backups.

- When you use BlueXP backup and recovery in a SaaS environment, the BlueXP backup and recovery configuration data is backed up to the cloud. When you use BlueXP backup and recovery in a site with no internet access, the BlueXP backup and recovery configuration data is backed up to the StorageGRID bucket where your backups are being stored. If you ever have a Connector failure in your private mode site, you can [restore the BlueXP backup and recovery data to a new Connector](#).

Verify license requirements

Before you can activate BlueXP backup and recovery for your cluster, you'll need to purchase and activate a BlueXP backup and recovery 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 supported when backing up files to StorageGRID.

Prepare your ONTAP clusters

Unresolved directive in task-backup-onprem-private-cloud.adoc - include::.../_include/backup-onprem-prepare-onprem-ONTAP-cluster.adoc[]

Verify ONTAP networking requirements for backing up data to object storage

You must configure the following requirements on the system that connects to object storage.

- When you use a fan-out backup architecture, the following settings must be configured on the *primary* storage system.
- When you use a cascaded backup architecture, the following settings must be configured on the *secondary* storage system.

The following ONTAP cluster networking requirements are needed:

- The ONTAP cluster initiates an HTTPS connection over a user-specified port from the intercluster LIF to the StorageGRID Gateway Node 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 BlueXP backup and recovery, 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 (not required when the Connector is installed in a "dark" site).
- DNS servers have been configured for the storage VM where the volumes are located. See how to [configure DNS services for the SVM](#).
- 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 BlueXP backup and recovery 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).

Verify ONTAP networking requirements for replicating volumes

Unresolved directive in task-backup-onprem-private-cloud.adoc - include::../_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

Prepare StorageGRID as your backup target

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

Supported StorageGRID versions

StorageGRID 10.3 and later is supported.

To use DataLock & Ransomware Protection for your backups, your StorageGRID systems must be running version 11.6.0.3 or greater.

To tier older backups to cloud archival storage, your StorageGRID systems must be running version 11.3 or greater. Additionally, your StorageGRID systems must be discovered to the BlueXP Canvas.

S3 credentials

You must have created an S3 tenant account to control access to your StorageGRID storage. [See the StorageGRID docs for details](#).

When you set up backup to StorageGRID, the backup wizard prompts you for an S3 access key and secret key for a tenant account. The tenant account enables BlueXP backup and recovery 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 manually on the object store bucket.

Prepare to archive older backup files to public cloud storage

Tiering older backup files to archival storage saves money by using a less expensive storage class for backups that you may not need. StorageGRID is an on-premises (private cloud) solution that doesn't provide archival storage, but you can move older backup files to public cloud archival storage. When used in this fashion, data that is tiered to cloud storage, or restored from cloud storage, goes between StorageGRID and the cloud storage - BlueXP is not involved in this data transfer.

Current support enables you to archive backups to *AWS S3 Glacier/S3 Glacier Deep Archive* or *Azure Archive* storage.

ONTAP Requirements

- Your cluster must be using ONTAP 9.12.1 or greater.

StorageGRID Requirements

- Your StorageGRID must be using 11.4 or greater.
- Your StorageGRID must be [discovered and available in the BlueXP Canvas](#).

Amazon S3 requirements

- You'll need to sign up for an Amazon S3 account for the storage space where your archived backups will be located.
- You can choose to tier backups to AWS S3 Glacier or S3 Glacier Deep Archive storage. [Learn more about AWS archival tiers](#).
- StorageGRID should have full-control access to the bucket (`s3:*`); however, if this is not possible, the bucket policy must grant the following S3 permissions to StorageGRID:
 - `s3:AbortMultipartUpload`
 - `s3:DeleteObject`
 - `s3:GetObject`
 - `s3:ListBucket`
 - `s3:ListBucketMultipartUploads`
 - `s3:ListMultipartUploadParts`
 - `s3:PutObject`

◦ `s3:RestoreObject`

Azure Blob requirements

- You'll need to sign up for an Azure Subscription for the storage space where your archived backups will be located.
- The activation wizard enables you to use an existing Resource Group to manage the Blob container that will store the backups, or you can create a new Resource Group.

When defining the Archival settings for the backup policy for your cluster, you'll enter your cloud provider credentials and select the storage class that you want to use. BlueXP backup and recovery creates the cloud bucket when you activate backup for the cluster. The information required for AWS and Azure archival storage is shown below.

AWS	Azure
<input checked="" type="checkbox"/> Tier Backups to Archive	<input checked="" type="checkbox"/> Tier Backups to Archive
Cloud Provider <div>AWS</div>	Cloud Provider <div>AZURE</div>
Account <div>Select Account</div>	Azure Subscription <div>Select Account</div>
Region <div>Select Region</div>	Region <div>Select Region</div>
AWS Access Key <div>Enter AWS Access Key</div>	Resource Group Type <div>Select an Existing Resource Group</div>
AWS Secret Key <div>Enter AWS Secret Key</div>	Resource Group <div>Select Resource Group</div>
Archive After (Days) <div>(1-999)</div>	Archive After (Days) <div>(1-999)</div>
Storage Class <div>S3 Glacier</div>	Storage Class <div>Azure Archive</div>

The archival policy settings you select will generate an information lifecycle management (ILM) policy in StorageGRID, and add the settings as "rules".

- If there is an existing active ILM policy, new rules will be added to the ILM policy to move the data to the archive tier.
- If there is an existing ILM policy in the "proposed" state, the creation and activation of a new ILM policy will not be possible. [Learn more about StorageGRID ILM policies and rules.](#)

Activate backups on your ONTAP volumes

Activate backups at any time directly from your on-premises working environment.

A wizard takes you through the following major steps:

- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future working environments.

Start the wizard

Steps

1. Access the Activate backup and recovery wizard using one of the following ways:

- From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.

If the destination for your backups exists as a working environment on the Canvas, you can drag the ONTAP cluster onto the object storage.

- Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions (...)** option and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

The Introduction page of the wizard shows the protection options including local Snapshots, replication, and backups. If you did the second option in this step, the Define Backup Strategy page appears with one volume selected.

2. Continue with the following options:

- If you already have a BlueXP Connector, you're all set. Just select **Next**.
- If you don't already have a BlueXP Connector, the **Add a Connector** option appears. Refer to [Prepare your BlueXP Connector](#).

Select the volumes that you want to back up

Choose the volumes you want to protect. A protected volume is one that has one or more of the following: Snapshot policy, replication policy, backup to object policy.

You can choose to protect FlexVol or FlexGroup volumes; however, you cannot select a mix of these volumes when activating backup for a working environment. See how to [activate backup for additional volumes in the working environment](#) (FlexVol or FlexGroup) after you have configured backup for the initial volumes.



- You can activate a backup only on a single FlexGroup volume at a time.
- The volumes you select must have the same SnapLock setting. All volumes must have SnapLock Enterprise enabled or have SnapLock disabled.

Steps

Note that if the volumes you choose already have Snapshot or replication policies applied, then the policies you select later will overwrite these existing policies.

1. In the Select Volumes page, select the volume or volumes you want to protect.
 - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
 - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row. (☒ Volume Name).
 - To back up individual volumes, check the box for each volume (☒ Volume_1).
2. Select **Next**.

Define the backup strategy

Defining the backup strategy involves setting the following options:

- Whether you want one or all of the backup options: local Snapshots, replication, and backup to object storage

- Architecture
- Local Snapshot policy
- Replication target and policy



If the volumes you choose have different Snapshot and replication policies than the policies you select in this step, the existing policies will be overwritten.

- Backup to object storage information (provider, encryption, networking, backup policy, and export options).

Steps

1. In the Define backup strategy page, choose one or all of the following. All three are selected by default:
 - **Local Snapshots:** If you are performing replication or back up to object storage, local Snapshots must be created.
 - **Replication:** Creates replicated volumes on another ONTAP storage system.
 - **Backup:** Backs up volumes to object storage.
2. **Architecture:** If you chose both replication and backup, choose one of the following flows of information:
 - **Cascading:** Information flows from the primary to the secondary, and then from the secondary to object storage.
 - **Fan out:** Information flows from the primary to the secondary *and* from the primary to object storage.

For details about these architectures, refer to [Plan your protection journey](#).

3. **Local Snapshot:** Choose an existing Snapshot policy or create a new one.



To create a custom policy before activating the Snapshot, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

4. **Replication:** Set the following options:
 - **Replication target:** Select the destination working environment and SVM. Optionally, select the destination aggregate or aggregates and prefix or suffix that will be added to the replicated volume name.
 - **Replication policy:** Choose an existing replication policy or create one.



To create a custom policy before activating the replication, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- Select **Create**.

5. **Back up to Object:** If you selected **Backup**, set the following options:

- **Provider:** Select **StorageGRID**.
- **Provider settings:** Enter the provider gateway node FQDN details, port, access key and secret key.

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

- **Networking:** Choose 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 (not required when the Connector is installed in a "dark" site).



Selecting the correct IPspace ensures that BlueXP backup and recovery can set up a connection from ONTAP to your StorageGRID object storage.

- **Backup policy:** Select an existing Backup to object storage policy or create one.



To create a custom policy before activating the backup, refer to [Create a policy](#).

To create a policy, select **Create new policy** and do the following:

- Enter the name of the policy.
- Select up to 5 schedules, typically of different frequencies.
- For backup-to-object policies, set the DataLock and Ransomware Protection settings. For details on DataLock and Ransomware Protection, refer to [Backup-to-object policy settings](#).

If your cluster is using ONTAP 9.11.1 or greater, you can choose to protect your backups from deletion and ransomware attacks by configuring *DataLock and Ransomware Protection*. *DataLock* protects your backup files from being modified or deleted, and *Ransomware Protection* scans your backup files to look for evidence of a ransomware attack in your backup files.

- Select **Create**.

If your cluster is using ONTAP 9.12.1 or greater and your StorageGRID system is using version 11.4 or greater, you can choose to tier older backups to public cloud archive tiers after a certain number of days. Current support is for AWS S3 Glacier/S3 Glacier Deep Archive or Azure Archive storage tiers. [See how to configure your systems for this functionality](#).

- **Tier backup to public cloud:** Select the cloud provider that you want to tier backups to and enter the provider details.

Select or create a new StorageGRID cluster. For details about creating a StorageGRID cluster so BlueXP can discover it, refer to [StorageGRID documentation](#).

- **Export existing Snapshot copies to object storage as backup copies:** If there are any local Snapshot copies for volumes in this working environment that match the backup schedule label you just selected for this working environment (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

Review your selections

This is the chance to review your selections and make adjustments, if necessary.

Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies.
3. Select **Activate Backup**.

Result

BlueXP backup and recovery starts taking the initial backups of your volumes. The baseline transfer of the replicated volume and the backup file includes a full copy of the source data. Subsequent transfers contain differential copies of the primary storage data contained in Snapshot copies.

A replicated volume is created in the destination cluster that will be synchronized with the primary storage volume.

An S3 bucket is created 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.

You can also monitor the status of backup and restore jobs using the [Job Monitoring panel](#).

Show the API commands

You might want to display and optionally copy the API commands used in the Activate backup and recovery wizard. You might want to do this to automate backup activation in future working environments.

Steps

1. From the Activate backup and recovery wizard, select **View API request**.
2. To copy the commands to the clipboard, select the **Copy** icon.

What's next?

- You can [manage your backup files and backup policies](#). This includes starting and stopping backups, deleting backups, adding and changing the backup schedule, and more.
- You can [manage cluster-level backup settings](#). This includes changing the network bandwidth available to upload backups to object storage, changing the automatic backup setting for future volumes, and more.
- You can also [restore volumes, folders, or individual files from a backup file](#) to an on-premises ONTAP system.

Manage backups for your ONTAP systems

You can manage backups for your Cloud Volumes ONTAP and on-premises ONTAP systems by changing the backup schedule, enabling/disabling volume backups, pausing backups, deleting backups, and more. This includes all types of backups, including Snapshot copies, replicated volumes, and backup files in object storage.



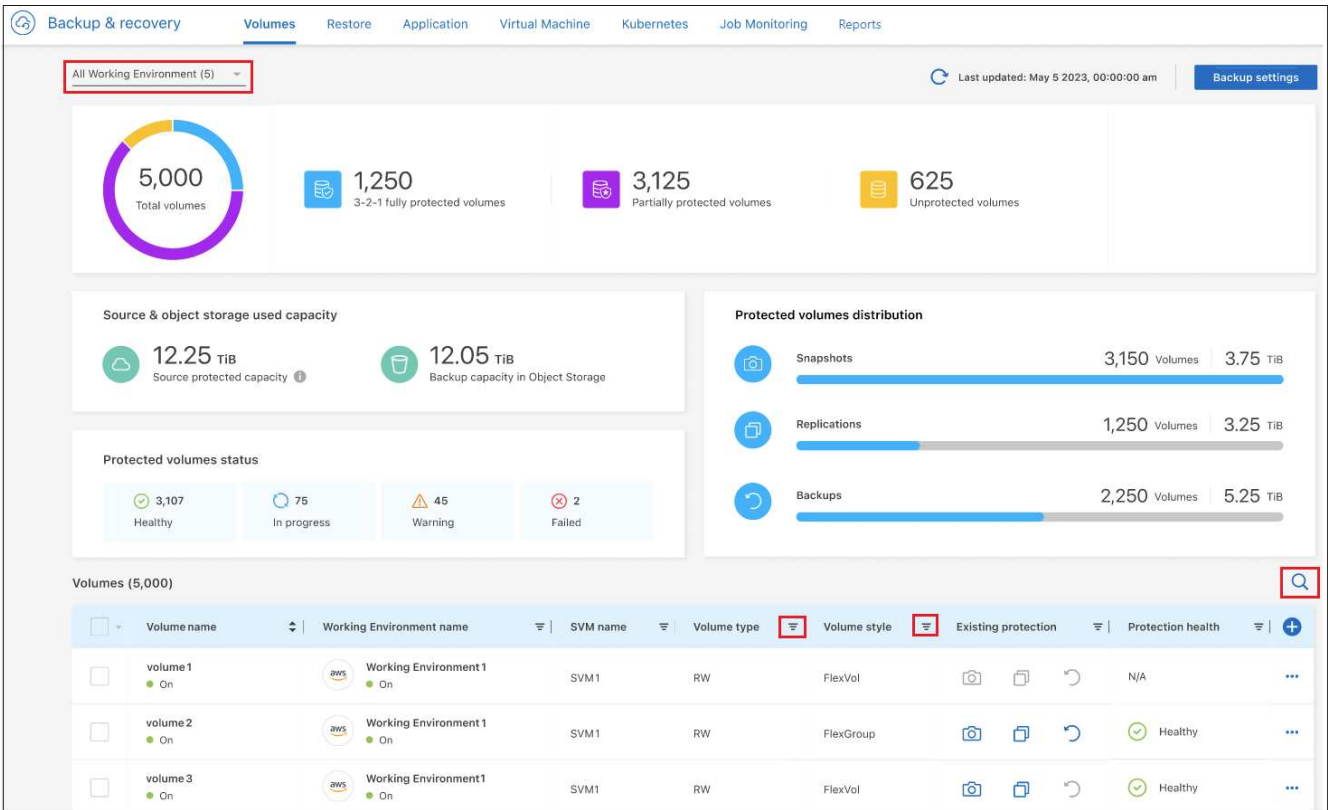
Do not manage or change backup files directly on your storage systems or from your cloud provider environment. This may corrupt the files and will result in an unsupported configuration.

View the backup status of volumes in your working environments

You can view a list of all the volumes that are currently being backed up in the Volumes Backup Dashboard. This includes all types of backups, including Snapshot copies, replicated volumes, and backup files in object storage. You can also view the volumes in those working environments that are not currently being backed up.

Steps

- 1. From the BlueXP menu, select **Protection > Backup and recovery**.
- 2. Click the **Volumes** tab to view the list of backed up volumes for your Cloud Volumes ONTAP and on-premises ONTAP systems.



- 3. If you are looking for specific volumes in certain working environments, you can refine the list by working environment and volume. You can also use the search filter, or you can sort the columns based on volume style (FlexVol or FlexGroup), volume type, and more.

To show additional columns (aggregates, security style (Windows or UNIX), snapshot policy, replication policy, and backup policy), select **+**.

- 4. Review the status of the protection options in the "Existing protection" column. The 3 icons stand for "Local Snapshot copies", "Replicated volumes", and "Backups in object storage".



Each icon is blue when that backup type is activated, and it's grey when the backup type is inactive. You can hover your cursor over each icon to see the backup policy that is being used, and other pertinent

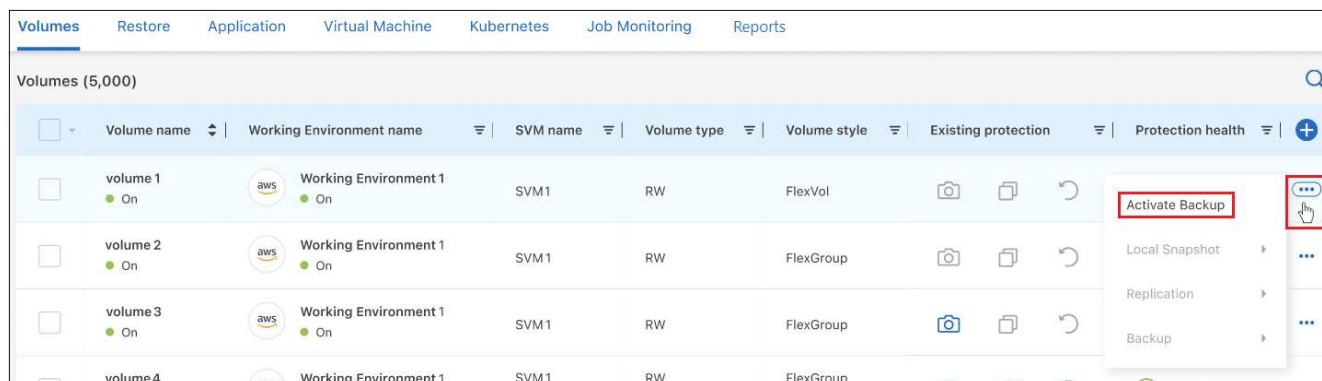
information for each type of backup.

Activate backup on additional volumes in a working environment

If you activated backup only on some of the volumes in a working environment when you first enabled BlueXP backup and recovery, you can activate backups on additional volumes later.

Steps

1. From the **Volumes** tab, identify the volume on which you want to activate backups, select the Actions menu **...** at the end of the row, and select **Activate backup**.



2. In the *Define backup strategy* page, select the backup architecture, and then define the policies and other details for Local Snapshot copies, Replicated volumes, and Backup files. See the details for backup options from the initial volumes you activated in this working environment. Then click **Next**.
3. Review the backup settings for this volume, and then click **Activate Backup**.

If you want to activate backup on multiple volumes at the same time with identical backup settings, see [Edit backup settings on multiple volumes](#) for details.


Change the backup settings assigned to existing volumes

You can change the backup policies assigned to your existing volumes that have assigned policies. You can change the policies for your Local Snapshot copies, Replicated volumes, and Backup files. Any new Snapshot, replication, or backup policy that you want to apply to the volumes must already exist.

Edit backup settings on a single volume

Steps

1. From the **Volumes** tab, identify the volume that you want to make policy changes, select the Actions menu **...** at the end of the row, and select **Edit backup strategy**.



2. In the *Edit backup strategy* page, make changes to the existing backup policies for Local Snapshot copies, Replicated volumes, and Backup files and click **Next**.

If you enabled *DataLock and Ransomware Protection* for cloud backups in the initial backup policy when activating BlueXP backup and recovery for this cluster, you'll only see other policies that have been configured with DataLock. And if you did not enable *DataLock and Ransomware Protection* when activating BlueXP backup and recovery, you'll only see other cloud backup policies that don't have DataLock configured.

3. Review the backup settings for this volume, and then click **Activate Backup**.

Edit backup settings on multiple volumes


If you want to use the same backup settings on multiple volumes, you can activate or edit backup settings on multiple volumes at the same time. You can select volumes that have no backup settings, only Snapshot settings, only backup to cloud settings, and so on, and make bulk changes across all these volumes with diverse backup settings.

When working with multiple volumes, all volumes must have these common characteristics:

- same working environment
- same style (FlexVol or FlexGroup volume)
- same type (Read-write or Data Protection volume)

Steps

1. From the **Volumes** tab, filter by the working environment on which the volumes reside.
2. Select all the volumes on which you want to manage backup settings.
3. Depending on the type of backup action you want to configure, click the button in the Bulk actions menu:



Backup action...	Click this button...
Manage Snapshot backup settings	Manage Local Snapshots
Manage Replication backup settings	Manage Replication
Manage Backup to cloud backup settings	Manage Backup
Manage multiple types of backup settings. This option enables you to change the backup architecture as well.	Manage Backup and Recovery

- In the backup page that appears, make changes to the existing backup policies for Local Snapshot copies, Replicated volumes, or Backup files and click **Save**.

If you enabled *DataLock and Ransomware Protection* for cloud backups in the initial backup policy when activating BlueXP backup and recovery for this cluster, you'll only see other policies that have been configured with DataLock. And if you did not enable *DataLock and Ransomware Protection* when activating BlueXP backup and recovery, you'll only see other cloud backup policies that don't have DataLock configured.

Create a manual volume backup at any time

You can create an on-demand backup at any time to capture the current state of the volume. This can be useful if very important changes have been made to a volume and you don't want to wait for the next scheduled backup to protect that data. You can also use this functionality to create a backup for a volume that is not currently being backed up and you want to capture its current state.

You can create an ad-hoc Snapshot copy or backup to object of a volume. You can't create an ad-hoc replicated volume.

The backup name includes the timestamp so you can identify your on-demand backup from other scheduled backups.

If you enabled *DataLock and Ransomware Protection* when activating BlueXP backup and recovery for this cluster, the on-demand backup also will be configured with DataLock, and the retention period will be 30 days. Ransomware scans are not supported for ad-hoc backups. [Learn more about DataLock and Ransomware protection.](#)

Note that when creating an ad-hoc backup, a Snapshot is created on the source volume. Since this Snapshot is not part of a normal Snapshot schedule, it will not rotate off. You may want to manually delete this Snapshot from the source volume once the backup is complete. This will allow blocks related to this Snapshot to be freed


up. The name of the Snapshot will begin with `cbs-snapshot-adhoc-`. [See how to delete a Snapshot using the ONTAP CLI.](#)



On-demand volume backup isn't supported on data protection volumes.

Steps

- 1. From the **Volumes** tab, click **...** for the volume and select **Backup > Create Ad-hoc Backup**.




The Backup Status column for that volume displays "In Progress" until the backup is created.

View the list of backups for each volume


You can view the list of all backup files that exist for each volume. This page displays details about the source volume, destination location, and backup details such as last backup taken, the current backup policy, backup file size, and more.

Steps

- 1. From the **Volumes** tab, click **...** for the source volume and select **View volume details**.



The details for the volume and the list of Snapshot copies are displayed by default.



2. Select **Snapshot**, **Replication**, or **Backup** to see the list of all backup files for each type of backup.



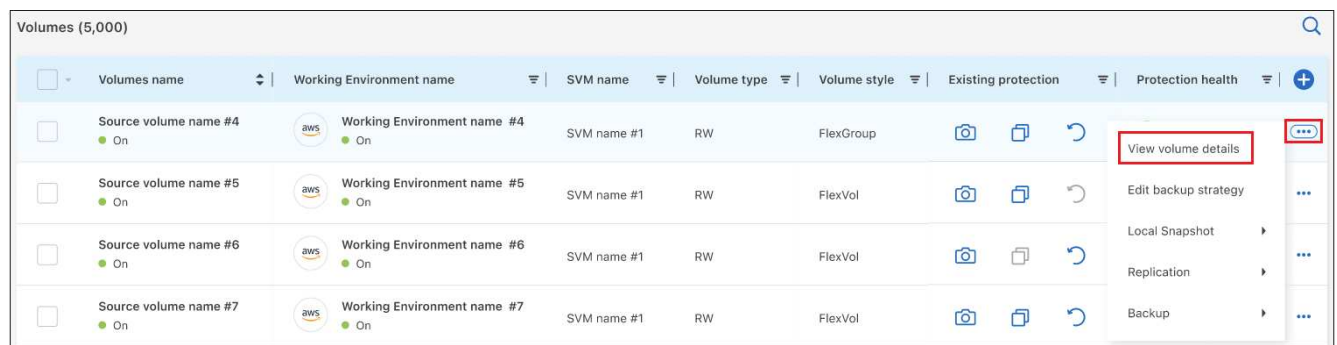
Run a ransomware scan on a volume backup in object storage

NetApp ransomware protection software scans your backup files to look for evidence of a ransomware attack when a backup to object file is created, and when data from a backup file is being restored. You can also run an on-demand ransomware protection scan at any time to verify the usability of a specific backup file in object storage. This can be useful if you have had a ransomware issue on a particular volume and you want to verify that the backups for that volume are not affected.

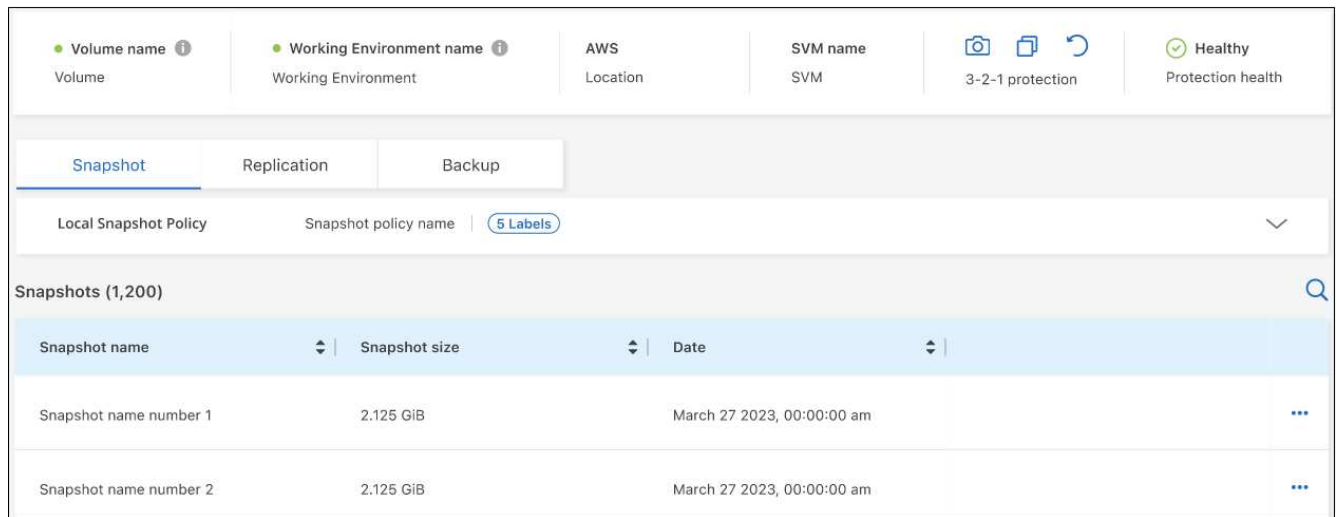
This feature is available only if the volume backup was created from a system with ONTAP 9.11.1 or greater, and if you enabled *DataLock and Ransomware Protection* in the backup to object policy.

Steps

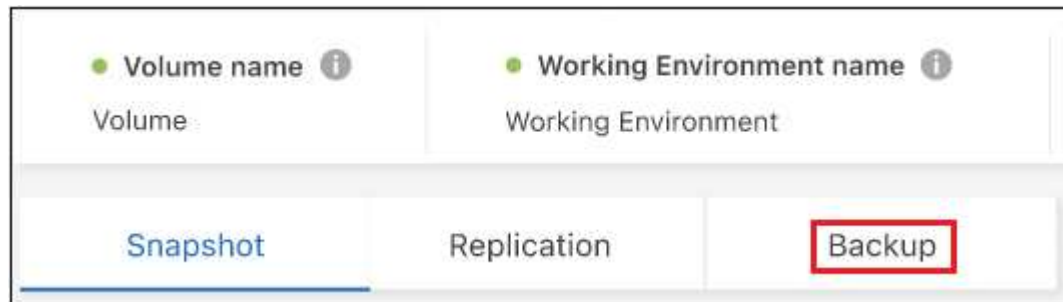
1. From the **Volumes** tab, click **...** for the source volume and select **View volume details**.




The details for the volume are displayed.



2. Select **Backup** to see the list of backup files in object storage.



3. Click ... for the volume backup file you want to scan for ransomware and click **Scan for Ransomware**.




The Ransomware Protection column will show that the scan is In Progress.

Manage the replication relationship with the source volume

After you set up data replication between two systems, you can manage the data replication relationship.

Steps

1. From the **Volumes** tab, click ... for the source volume and select the **Replication** option. You can see all of the available options.
2. Select the replication action that you want to perform.



The following table describes the available actions:

Action	Description
View Replication	Shows you details about the volume relationship: transfer information, last transfer information, details about the volume, and information about the protection policy assigned to the relationship.
Update Replication	Starts an incremental transfer to update the destination volume to be synchronized with the source volume.
Pause Replication	Pause the incremental transfer of Snapshot copies to update the destination volume. You can Resume later if you want to restart the incremental updates.
Break Replication	<p>Breaks the relationship between the source and destination volumes, and activates the destination volume for data access - makes it read-write.</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>Learn how to configure a destination volume for data access and reactivate a source volume in the ONTAP documentation</p>
Abort Replication	Disables backups of this volume to the destination system, and it also disables the ability to restore a volume. Any existing backups will not be deleted. This does not delete the data protection relationship between the source and destination volumes.
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>
Delete Relationship	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 - meaning it does not make it read-write. This action also deletes the cluster peer relationship and the storage VM (SVM) peer relationship, if there are no other data protection relationships between the systems.

Result

After you select an action, BlueXP updates the relationship.

Edit an existing backup-to-cloud policy


You can change the attributes for a backup policy that is currently applied to volumes in a working environment. Changing the backup policy affects all existing volumes that are using the policy.



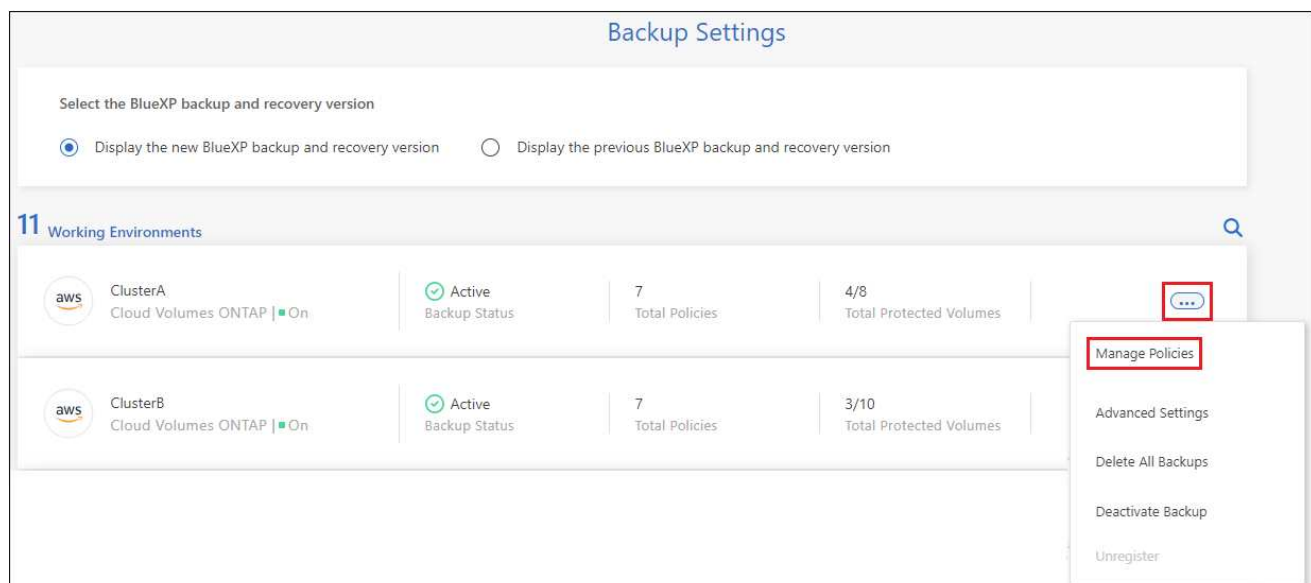
- If you enabled *DataLock and Ransomware Protection* in the initial policy when activating BlueXP backup and recovery for this cluster, any policies that you edit must be configured with the same DataLock setting (Governance or Compliance). And if you did not enable *DataLock and Ransomware Protection* when activating BlueXP backup and recovery, you can't enable DataLock now.
- When creating backups on AWS, if you chose *S3 Glacier* or *S3 Glacier Deep Archive* in your first backup policy when activating BlueXP backup and recovery, then that tier will be the only archive tier available when editing backup policies. And if you selected no archive tier in your first backup policy, then *S3 Glacier* will be your only archive option when editing a policy.

Steps

1. From the **Volumes** tab, select **Backup Settings**.



2. From the *Backup Settings* page, click **...** for the working environment where you want to change the policy settings, and select **Manage Policies**.



- From the *Manage Policies* page, click **Edit** for the backup policy you want to change in that working environment.

Manage Policies

Add New Policy

Working Environment: ClusterB

Only Custom policies are editable

7 Policies

hourly_bp
Custom Policy

Edit

2 Labels: Hourly (10), Daily (10)
Labels & Retention

None
DataLock & Ransomware Protection

Not Active
Archival Policy

3 out of 10
Associated Volumes

- From the *Edit Policy* page, click  to expand the *Labels & Retention* section to change the schedule and/or backup retention, and click **Save**.

Edit Policy

Working Environment: ClusterB

Name

hourly_bp

Labels & Retention

10 Hourly | 10 Daily

DataLock & Ransomware Protection

None

Archival Policy

Disabled

If your cluster is running ONTAP 9.10.1 or greater, you also have the option to enable or disable tiering of backups to archival storage after a certain number of days.

[Learn more about using AWS archival storage.](#)

[Learn more about using Azure archival storage.](#)

[Learn more about using Google archival storage.](#) (Requires ONTAP 9.12.1.)

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.

Azure

☒ Tier Backups to Archival

Archive after (Days)

30

Access Tier

Azure Archive

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.

AWS

☒ Tier Backups to Archival

Archive after (Days)

30

Storage Class

S3 Glacier
S3 Glacier
S3 Glacier Deep Archive

Archival Policy

Backups reside in Google Cloud Standard storage for frequently accessed data. Optionally, you can tier backups to Google Cloud Archive storage for further cost optimization.

Google

☒ Tier Backups to Archival

Archive after (Days)

30

Storage Class

Google Cloud Archive

Note that any backup files that have been tiered to archival storage are left in that tier if you stop tiering backups to archive - they are not automatically moved back to the standard tier. Only new volume backups will reside in the standard tier.

Add a new backup-to-cloud policy

When you enable BlueXP backup and recovery for a working environment, all the volumes you initially select are backed up using the default backup policy that you defined. If you want to assign different backup policies to certain volumes that have different recovery point objectives (RPO), you can create additional policies for that cluster and assign those policies to other volumes.


If you want to apply a new backup policy to certain volumes in a working environment, you first need to add the backup policy to the working environment. Then you can [apply the policy to volumes in that working environment](#).



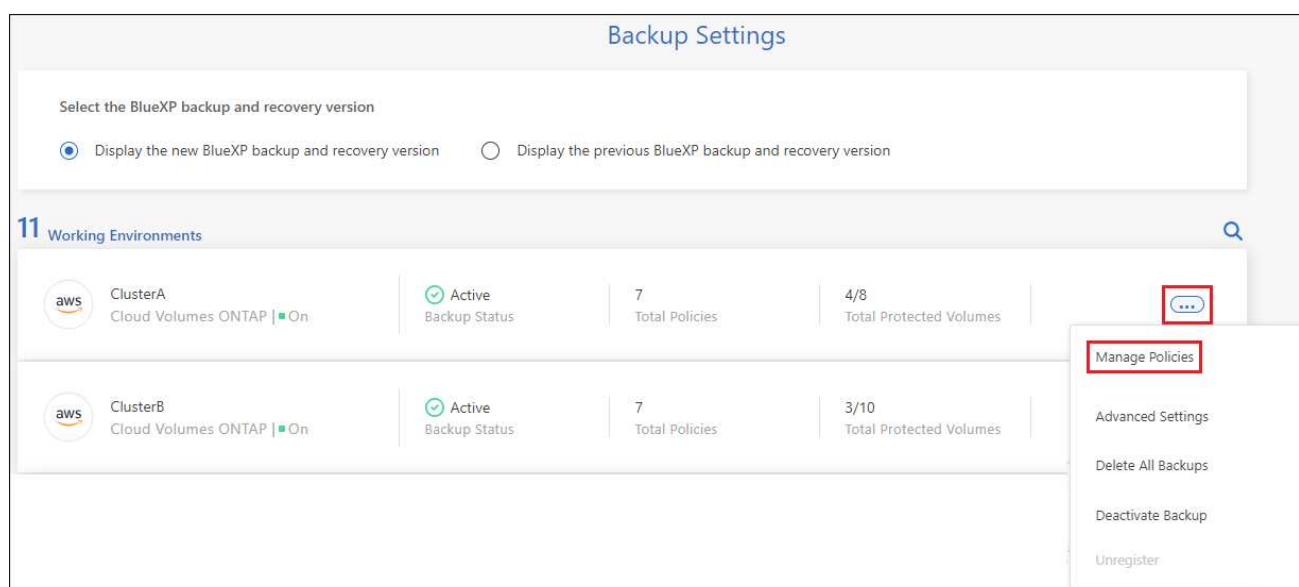
- If you enabled *DataLock and Ransomware Protection* in the initial policy when activating BlueXP backup and recovery for this cluster, any additional policies you create must be configured with the same DataLock setting (Governance or Compliance). And if you did not enable *DataLock and Ransomware Protection* when activating BlueXP backup and recovery, you can't create new policies that use DataLock.
- When creating backups on AWS, if you chose *S3 Glacier* or *S3 Glacier Deep Archive* in your first backup policy when activating BlueXP backup and recovery, then that tier will be the only archive tier available for future backup policies for that cluster. And if you selected no archive tier in your first backup policy, then *S3 Glacier* will be your only archive option for future policies.

Steps

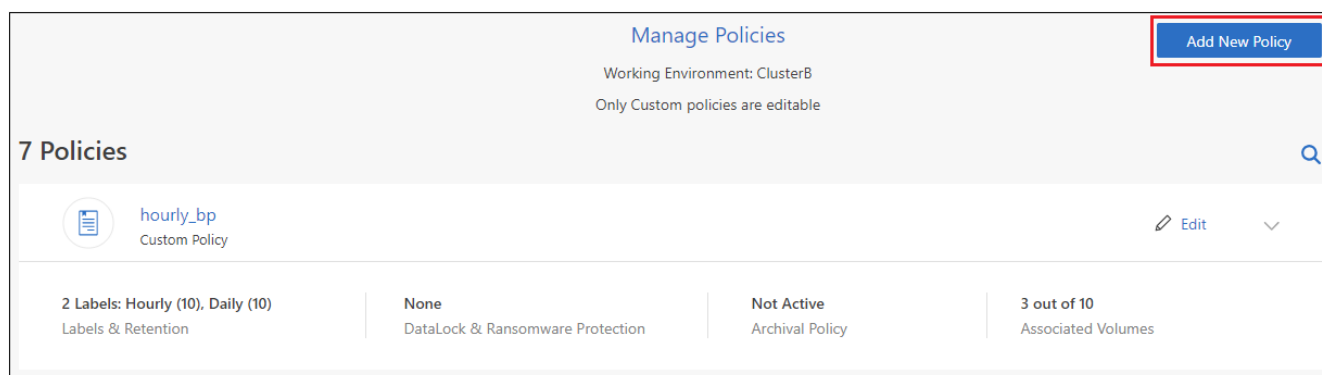
1. From the **Volumes** tab, select **Backup Settings**.



2. From the *Backup Settings* page, click ... for the working environment where you want to add the new policy, and select **Manage Policies**.



3. From the *Manage Policies* page, click **Add New Policy**.



4. From the *Add New Policy* page, click ▼ to expand the *Labels & Retention* section to define the schedule and backup retention, and click **Save**.

Add New Policy	
Working Environment: Working Environment 1	
Name	Default_Policy_Name
Labels & Retention	30 Daily
DataLock & Ransomware Protection	None
Archival Policy	Disabled

If your cluster is running ONTAP 9.10.1 or greater, you also have the option to enable or disable tiering of backups to archival storage after a certain number of days.

[Learn more about using AWS archival storage.](#)

[Learn more about using Azure archival storage.](#)

[Learn more about using Google archival storage.](#) (Requires ONTAP 9.12.1.)

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:

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

Archival Policy

Backups reside in Google Cloud Standard storage for frequently accessed data. Optionally, you can tier backups to Google Cloud Archive storage for further cost optimization.

☒ Tier Backups to Archival

Archive after (Days): Storage Class:

Delete backups

BlueXP backup and recovery enables you to delete a single backup file, delete all backups for a volume, or delete all backups of all volumes in a working environment. You might want to delete all backups if you no longer need the backups, or if you deleted the source volume and want to remove all backups.

Note that you can't delete backup files that you have locked using DataLock and Ransomware protection. The "Delete" option will be unavailable from the UI if you have selected one or more locked backup files.



If you plan to delete a working environment or cluster that has backups, you must delete the backups **before** deleting the system. BlueXP backup and recovery 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. You'll continue to be charged for object storage costs for any remaining backups.

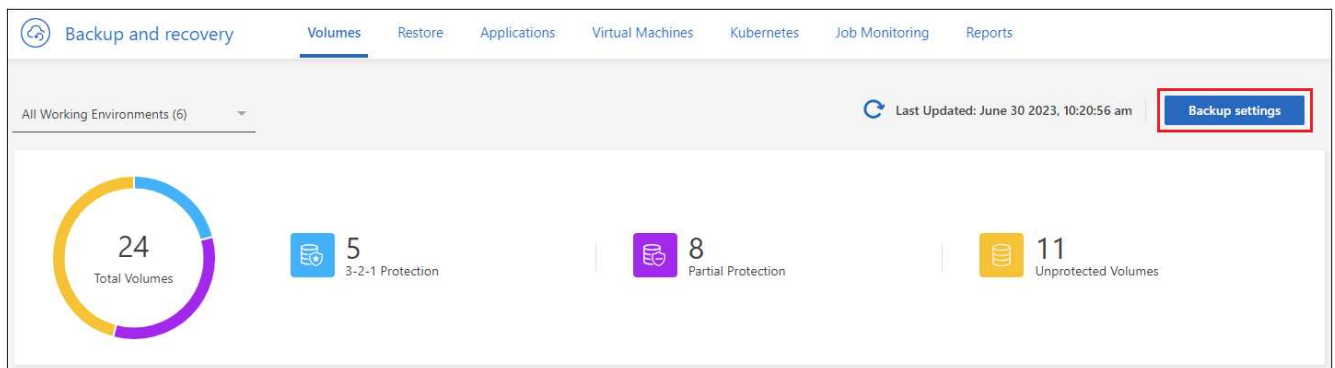
Delete all backup files for a working environment

Deleting all backups on object storage for a working environment does not disable future backups of volumes in this working environment. If you want to stop creating backups of all volumes in a working environment, you can deactivate backups [as described here](#).

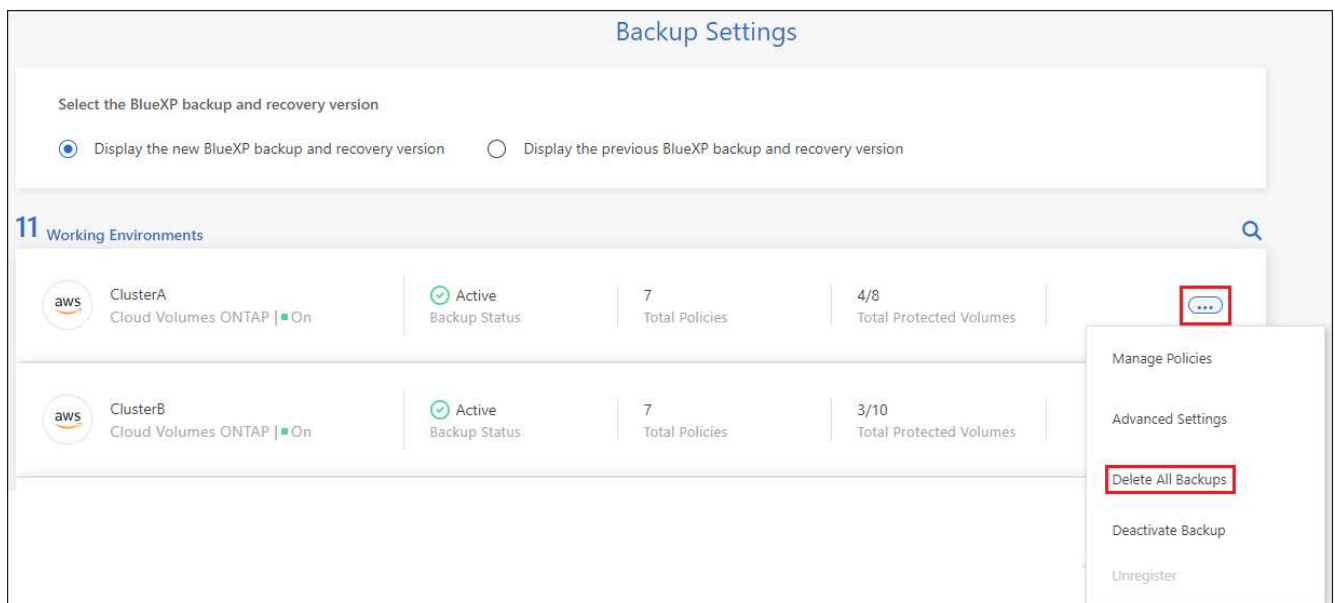
Note that this action does not affect Snapshot copies or replicated volumes - these types of backup files are not deleted.

Steps

1. From the **Volumes** tab, select **Backup Settings**.



2. Click **...** for the working environment where you want to delete all backups and select **Delete All Backups**.



3. In the confirmation dialog box, enter the name of the working environment and click **Delete**.

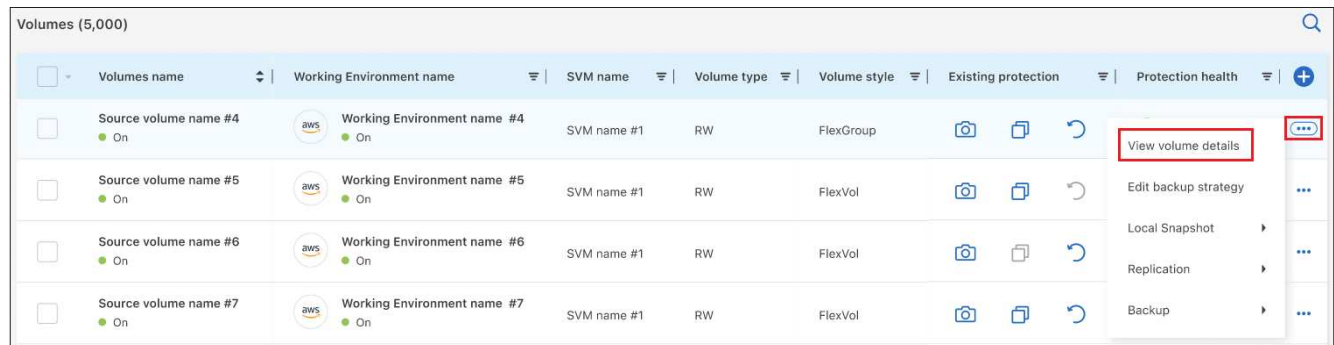
Delete a single backup file for a volume

You can delete a single backup file if you no longer need it. This includes deleting a single backup of a volume Snapshot copy or of a backup in object storage.

You can't delete replicated volumes (data protection volumes).

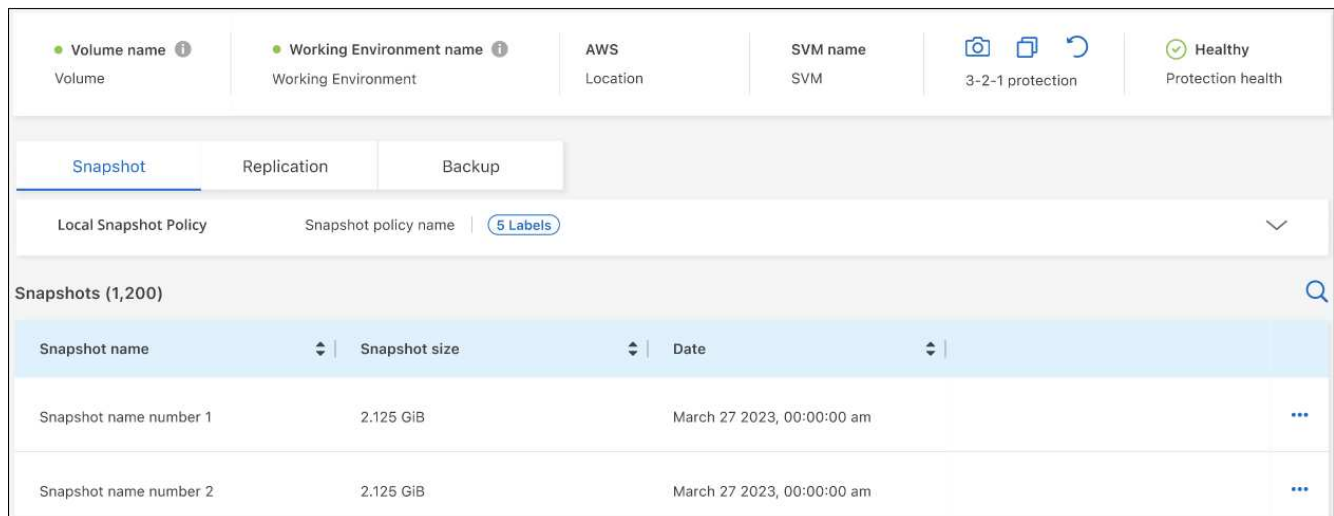
Steps

1. From the **Volumes** tab, click **...** for the source volume and select **View volume details**.



Volumes name	Working Environment name	SVM name	Volume type	Volume style	Existing protection	Protection health
Source volume name #4 On	Working Environment name #4 On	SVM name #1	RW	FlexGroup		
Source volume name #5 On	Working Environment name #5 On	SVM name #1	RW	FlexVol		
Source volume name #6 On	Working Environment name #6 On	SVM name #1	RW	FlexVol		
Source volume name #7 On	Working Environment name #7 On	SVM name #1	RW	FlexVol		

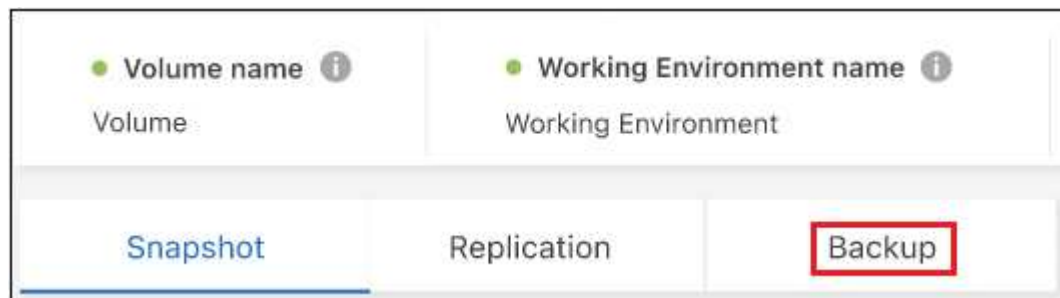
The details for the volume are displayed, and you can select **Snapshot**, **Replication**, or **Backup** to see the list of all backup files for the volume. By default, the available Snapshot copies are displayed.



Volume name	Working Environment name	AWS Location	SVM name	Protection health
Volume	Working Environment		SVM	Healthy

Snapshot name	Snapshot size	Date
Snapshot name number 1	2.125 GiB	March 27, 2023, 00:00:00 am
Snapshot name number 2	2.125 GiB	March 27, 2023, 00:00:00 am

2. Select **Snapshot** or **Backup** to see the type of backup files that you want to delete.



Volume name	Working Environment name
Volume	Working Environment

Snapshot	Replication	Backup

- Click **...** for the volume backup file you want to delete and click **Delete**. The screenshot below is from a backup file in object storage.

Backup name	Backup size	Date	Ransomware protection	Storage class	Snapmirror label	
Backup Name Number 1	2.125 GiB	March 27 2023, 00:00:00 am	Compliance	None		...
Backup name number 2	2.125 GiB	March 27 2023, 00:00:00 am	None	None		...
Backup name number 3	2.125 GiB	March 27 2023, 00:00:00 am	Compliance	None		...

Scan for Ransomware
 Restore
Delete

- In the confirmation dialog box, click **Delete**.

Delete volume backup relationships

Deleting the backup relationship for a volume provides you with an archiving mechanism if you want to stop the creation of new backup files and delete the source volume, but retain all the existing backup files. This gives you the ability to restore the volume from the backup file in the future, if needed, while clearing space from your source storage system.

You don't necessarily need to delete the source volume. You can delete the backup relationship for a volume and retain the source volume. In this case you can "Activate" backup on the volume at a later time. The original baseline backup copy continues to be used in this case - a new baseline backup copy is not created and exported to the cloud. Note that if you do reactivate a backup relationship, the volume is assigned the default backup policy.

This feature is available only if your system is running ONTAP 9.12.1 or greater.

You can't delete the source volume from the BlueXP backup and recovery user interface. However, you can open the Volume Details page on the Canvas, and [delete the volume from there](#).



You can't delete individual volume backup files once the relationship has been deleted. You can, however, you can delete all backups for the volume.

Steps

- From the **Volumes** tab, click **...** for the source volume and select **Backup > Delete relationship**.

Volume name	Working Environment name	SVM name	Volume type	Volume style	Existing protection	Protection health	
volume 4	Working Environment 4	SVM 1	RW	FlexGroup			...
volume 5	Working Environment 5	SVM 1	RW	FlexVol			...
volume 6	Working Environment 5	SVM 1	RW	FlexVol			...
volume 7	Working Environment 5	SVM 1	RW	FlexVol			...

View volume details
 Edit backup strategy
 Local Snapshot
 Replication
Delete relationship
Backup


Deactivate BlueXP backup and recovery for a working environment

Deactivating BlueXP backup and recovery 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. This does not unregister the backup service from this working environment - it basically allows you to pause all backup and restore activity for a period of time.


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. From the **Volumes** tab, select **Backup Settings**.



2. From the *Backup Settings* page, click ... for the working environment where you want to disable backups and select **Deactivate Backup**.



3. In the confirmation dialog box, click **Deactivate**.



An **Activate Backup** button appears for that working environment while backup is disabled. You can click this button when you want to re-enable backup functionality for that working environment.

Unregister BlueXP backup and recovery for a working environment

You can unregister BlueXP backup and recovery for a working environment if you no longer want to use backup functionality and you want to stop being charged for backups in that working environment. Typically this feature is used when you're planning to delete a working environment, and you want to cancel the backup service.

You can also use this feature if you want to change the destination object store where your cluster backups are being stored. After you unregister BlueXP backup and recovery for the working environment, then you can enable BlueXP backup and recovery for that cluster using the new cloud provider information.

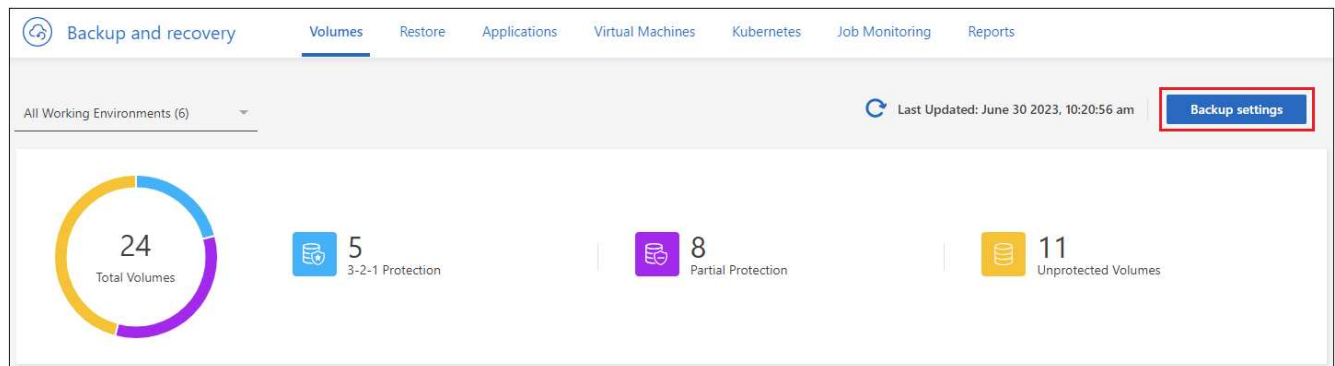
Before you can unregister BlueXP backup and recovery, you must perform the following steps, in this order:

- Deactivate BlueXP backup and recovery for the working environment
- Delete all backups for that working environment

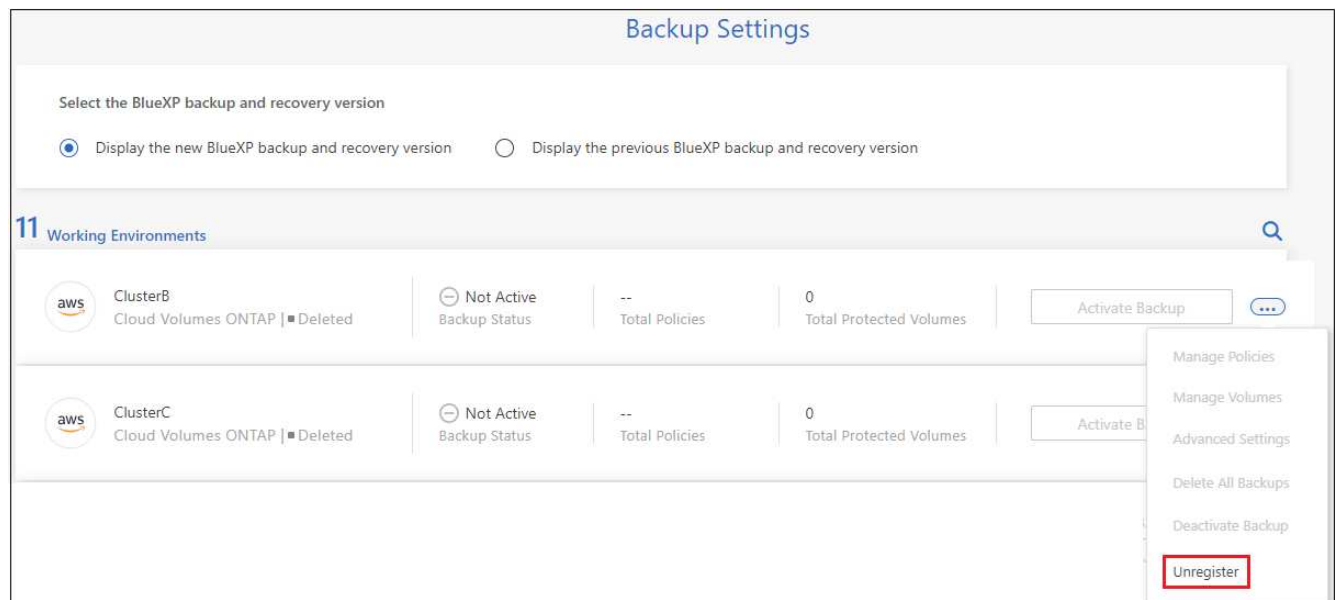
The unregister option is not available until these two actions are complete.

Steps

1. From the **Volumes** tab, select **Backup Settings**.



2. From the *Backup Settings* page, click ... for the working environment where you want to unregister the backup service and select **Unregister**.



3. In the confirmation dialog box, click **Unregister**.

Restore ONTAP data from backup files

Backups of your ONTAP volume data are available from the locations where you created backups: Snapshot copies, replicated volumes, and backups stored in object storage. You can restore data from a specific point in time from any of these backup locations. You can restore an entire ONTAP volume from a backup file, or if you only need to restore a few files, you can restore a folder or individual files.

- You can restore a **volume** (as a new volume) to the original working environment, to a different working environment that's using the same cloud account, or to an on-premises ONTAP system.
- You can restore a **folder** to a volume in the original working environment, to a volume in a different working environment that's using the same cloud account, or to a volume on an on-premises ONTAP system.
- You can restore **files** to a volume in the original working environment, to a volume in a different working environment that's using the same cloud account, or to a volume on an on-premises ONTAP system.

A valid BlueXP backup and recovery license is required to restore data from backup files to a production system.


To summarize, these are the valid flows you can use to restore volume data to an ONTAP working environment:

- Backup file → restored volume
- Replicated volume → restored volume
- Snapshot copy → restored volume




For limitations related to restoring ONTAP data, see [Backup and restore limitations for ONTAP volumes](#).

The Restore Dashboard

You use the Restore Dashboard to perform volume, folder, and file restore operations. You access the Restore Dashboard by clicking **Backup and recovery** from the BlueXP menu, and then clicking the **Restore** tab. You can also click  > **View Restore Dashboard** from the Backup and recovery service from the Services panel.



BlueXP backup and recovery must already be activated for at least one working environment and initial backup files must exist.



As you can see, the Restore Dashboard provides 2 different ways to restore data from backup files: **Browse & Restore** and **Search & Restore**.

Comparing Browse & Restore and Search & Restore

In broad terms, *Browse & Restore* is typically better when you need to restore a specific volume, folder, or file from the last week or month — and you know the name and location of the file, and the date when it was last in good shape. *Search & Restore* is typically better when you need to restore a volume, folder, or file, but you don't remember the exact name, or the volume in which it resides, or the date when it was last in good shape.

This table provides a feature comparison of the 2 methods.

Browse & Restore	Search & Restore
Browse through a folder-style structure to find the volume, folder, or file within a single backup file.	Search for a volume, folder, or file across all backup files by partial or full volume name, partial or full folder/file name, size range, and additional search filters.
Does not handle file recovery if the file has been deleted or renamed, and the user doesn't know the original file name	Handles newly created/deleted/renamed directories and newly created/deleted/renamed files
No additional cloud provider resources required	When you restore from the cloud, additional bucket and public cloud provider resources required per account.
No additional cloud provider costs required	When you restore from the cloud, additional costs are required when scanning your backups and volumes for search results.

Browse & Restore	Search & Restore
Quick restore is supported.	Quick restore is not supported.

This table provides a list of valid restore operations based on the location where your backup files reside.

Backup Type	Browse & Restore			Search & Restore		
	Restore volume	Restore files	Restore folder	Restore volume	Restore files	Restore folder
Snapshot copy	Yes	No	No	Yes	Yes	Yes
Replicated volume	Yes	No	No	Yes	Yes	Yes
Backup file	Yes	Yes	Yes	Yes	Yes	Yes

Before you can use either restore method, make sure you have configured your environment for the unique resource requirements. Those requirements are described in the sections below.

See the requirements and restore steps for the type of restore operation you want to use:

- [Restore volumes using Browse & Restore](#)
- [Restore folders and files using Browse & Restore](#)
- [Restore volumes, folders, and files using Search & Restore](#)

Restore ONTAP data using Browse & Restore

Before you start restoring a volume, folder, or file, you should know the name of the volume from which you want to restore, the name of the working environment and SVM where the volume resides, and the approximate date of the backup file that you want to restore from. You can restore ONTAP data from a Snapshot copy, a replicated volume, or from backups stored in object storage.

Note: If the backup file containing the data that you want to restore resides in archival cloud storage (starting with ONTAP 9.10.1), the restore operation will take a longer amount of time and will incur a cost. Additionally, the destination cluster must also be running ONTAP 9.10.1 or greater for volume restore, 9.11.1 for file restore, 9.12.1 for Google Archive and StorageGRID, and 9.13.1 for folder restore.

[Learn more about restoring from AWS archival storage.](#)
[Learn more about restoring from Azure archival storage.](#)
[Learn more about restoring from Google archival storage.](#)



The High priority isn't supported when restoring data from Azure archival storage to StorageGRID systems.

Browse & Restore supported working environments and object storage providers

You can restore ONTAP data from a backup file that resides in a secondary working environment (a replicated volume) or in object storage (a backup file) to the following working environments. Snapshot copies reside on the source working environment and can be restored only to that same system.

Note: You can restore a volume from any type of backup file, but you can restore a folder or individual files

only from a backup file in object storage at this time.

From Object Store (Backup)	From Primary (Snapshot)	From Secondary System (Replication)	To Destination Working Environment
Amazon S3	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Azure Blob
Cloud Volumes ONTAP in Azure On-premises ONTAP system	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Google Cloud Storage	Cloud Volumes ONTAP in Google On-premises ONTAP system
Cloud Volumes ONTAP in Google On-premises ONTAP system	NetApp StorageGRID	On-premises ONTAP system	On-premises ONTAP system Cloud Volumes ONTAP
To on-premises ONTAP system	ONTAP S3	On-premises ONTAP system	On-premises ONTAP system Cloud Volumes ONTAP

For Browse & Restore, the Connector can be installed in the following locations:

- For Amazon S3, the Connector can be deployed in AWS or in your premises
- For Azure Blob, the Connector can be deployed in Azure or in your premises
- For Google Cloud Storage, the Connector must be deployed in your Google Cloud Platform VPC
- For StorageGRID, the Connector must be deployed in your premises; with or without internet access
- For ONTAP S3, the Connector can be deployed in your premises (with or without internet access) or in a cloud provider environment

Note that references to "on-premises ONTAP systems" includes FAS, AFF, and ONTAP Select systems.



If the ONTAP version on your system is less than 9.13.1, then you can't restore folders or files if the backup file has been configured with DataLock & Ransomware. In this case, you can restore the entire volume from the backup file and then access the files you need.

Restore volumes using Browse & Restore


When you restore a volume from a backup file, BlueXP backup and recovery creates a *new* volume using the data from the backup. When using a backup from object storage, you can restore the data to a volume in the original working environment, to a different working environment that's located in the same cloud account as the source working environment, or to an on-premises ONTAP system.

When restoring a cloud backup to a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater or to an on-premises ONTAP system running ONTAP 9.14.1, you'll have the option to perform a *quick restore* operation. The quick restore is ideal for disaster recovery situations where you need to provide access to a volume as soon as possible. A quick restore restores the metadata from the backup file to a volume instead of restoring the entire backup file. Quick restore is not recommended for performance or latency-sensitive applications, and it is not supported with backups in archived storage.




Quick restore is supported for FlexGroup volumes only if the source system from which the cloud backup was created was running ONTAP 9.12.1 or greater. And it is supported for SnapLock volumes only if the source system was running ONTAP 9.11.0 or greater.

When restoring from a replicated volume, you can restore the volume to the original working environment or to a Cloud Volumes ONTAP or on-premises ONTAP system.



As you can see, you'll need to know the source working environment name, storage VM, volume name, and backup file date to perform a volume restore.

The following video shows a quick walkthrough of restoring a volume:



Steps

1. From the BlueXP menu, select **Protection > Backup and recovery**.
2. Click the **Restore** tab and the Restore Dashboard is displayed.
3. From the *Browse & Restore* section, click **Restore Volume**.



4. In the *Select Source* page, navigate to the backup file for the volume you want to restore. Select the **Working Environment**, the **Volume**, and the **Backup** file that has the date/time stamp from which you want to restore.

The **Location** column shows whether the backup file (Snapshot) is **Local** (a Snapshot copy on the source system), **Secondary** (a replicated volume on a secondary ONTAP system), or **Object Storage** (a backup file in object storage). Choose the file that you want to restore.


	Snapshot Name	Location	Date	Size	Ransomware Scan	Storage Class
<input type="radio"/>	Backup 1	Local	June 12 2022, 00:00:00	12.125 TiB	N/A	N/A
<input checked="" type="radio"/>	Backup 2	Local	June 12 2022, 00:00:00	12.125 TiB	N/A	N/A
<input type="radio"/>	Backup 3	Local	June 12 2022, 00:00:00	12.125 TiB	N/A	N/A
<input type="radio"/>	Backup 4	Object Storage	June 12 2022, 00:00:00	12.125 TiB	Protected	Standard

5. Click **Next**.

Note that if you select a backup file in object storage, and ransomware protection is active for that backup (if you enabled DataLock and Ransomware Protection in the backup policy), then you are prompted to run an additional ransomware scan on the backup file before restoring the data. We recommend that you scan the backup file for ransomware. (You'll incur extra egress costs from your cloud provider to access the

contents of the backup file.)

6. In the *Select Destination* page, select the **Working Environment** where you want to restore the volume.



7. When restoring a backup file from object storage, if you select an on-premises ONTAP system and you haven't already configured the cluster connection to the object storage, you are prompted for additional information:
- When restoring from Amazon S3, select the IPspace in the ONTAP cluster where the destination volume will reside, enter the access key and secret key for the user you created to give the ONTAP cluster access to the S3 bucket, and optionally choose a private VPC endpoint for secure data transfer.
 - When restoring from Azure Blob, select the IPspace in the ONTAP cluster where the destination volume will reside, select the Azure Subscription to access the object storage, and optionally choose a private endpoint for secure data transfer by selecting the VNet and Subnet.
 - When restoring from Google Cloud Storage, select the Google Cloud Project and the Access Key and Secret Key to access the object storage, the region where the backups are stored, and the IPspace in the ONTAP cluster where the destination volume will reside.
 - When restoring from StorageGRID, enter the FQDN of the StorageGRID server and the port that ONTAP should use for HTTPS communication with StorageGRID, select the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume will reside.
 - When restoring from ONTAP S3, enter the FQDN of the ONTAP S3 server and the port that ONTAP should use for HTTPS communication with ONTAP S3, select the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume will reside.
8. Enter the name you want to use for the restored volume, and select the Storage VM and Aggregate where the volume will reside. When restoring a FlexGroup volume you'll need to select multiple aggregates. By default, **<source_volume_name>_restore** is used as the volume name.

Select Destination

✓ Selected Working Environment
 Working Environment Name 2

📄 Destination Volume >
 General_restore

i A new volume will be created in the working environment based on the backup you selected

Volume Name

Storage VM

Aggregate

Restore Priority

Volume Information
Volume Size: 50.00 GB
Backup Policy: CloudBackupService
Protocol: NFS
Disk Type: RW

When restoring a backup from object storage to a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater or to an on-premises ONTAP system running ONTAP 9.14.1, you'll have the option to perform a *quick restore* operation.

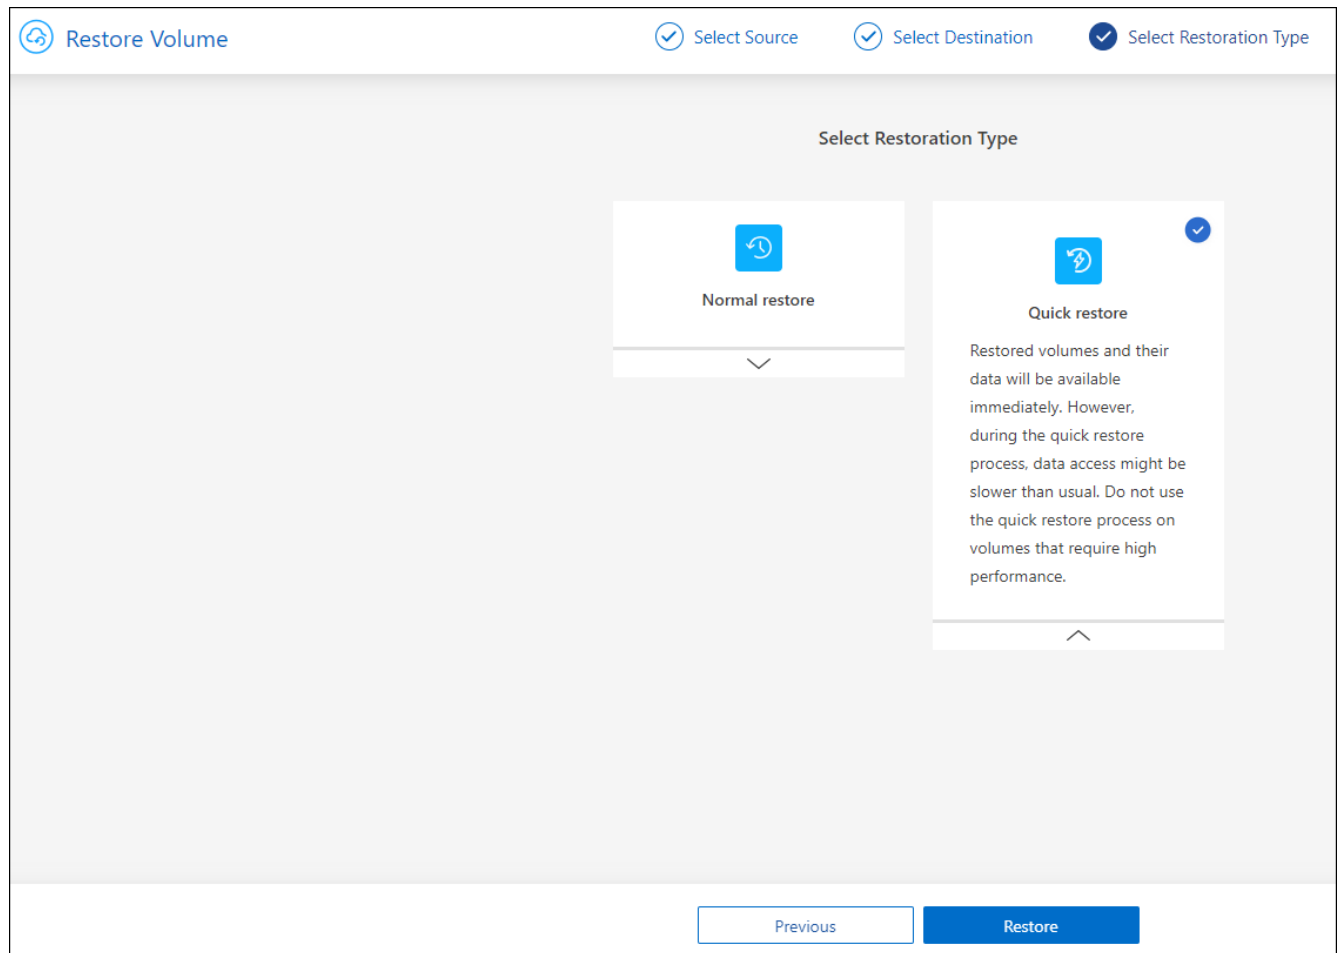
And if you are restoring the volume from a backup file that resides in an archival storage tier (available starting with ONTAP 9.10.1), then you can select the Restore Priority.

[Learn more about restoring from AWS archival storage.](#)

[Learn more about restoring from Azure archival storage.](#)

[Learn more about restoring from Google archival storage.](#) Backup files in the Google Archive storage tier are restored almost immediately, and require no Restore Priority.

9. Click **Next** to choose whether you want to do a Normal restore or a Quick Restore process:



- **Normal restore:** Use normal restore on volumes that require high performance. Volumes will not be available until the restore process is complete.
- **Quick restore:** Restored volumes and data will be available immediately. Do not use this on volumes that require high performance because during the quick restore process, access to the data might be slower than usual.

10. Click **Restore** and you are returned to the Restore Dashboard so you can review the progress of the restore operation.

Result

BlueXP backup and recovery creates a new volume based on the backup you selected.

Note that restoring a volume from a backup file that resides in archival storage can take many minutes or hours depending on the archive tier and the restore priority. You can click the **Job Monitoring** tab to see the restore progress.

Restore folders and files using Browse & Restore

If you need to restore only a few files from an ONTAP volume backup, you can choose to restore a folder or individual files instead of restoring the entire volume. You can restore folders and files to an existing volume in the original working environment, or to a different working environment that's using the same cloud account. You can also restore folders and files to a volume on an on-premises ONTAP system.



You can restore a folder or individual files only from a backup file in object storage at this time. Restoring files and folders is not currently supported from a local Snapshot copy or from a backup file that resides in a secondary working environment (a replicated volume).

If you select multiple files, all the files are restored to the same destination volume that you choose. So if you want to restore files to different volumes, you'll need to run the restore process multiple times.

When using ONTAP 9.13.0 or greater, you can restore a folder along with all files and sub-folders within it. When using a version of ONTAP before 9.13.0, only files from that folder are restored - no sub-folders, or files in sub-folders, are restored.



- If the backup file has been configured with DataLock & Ransomware protection, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the entire volume from the backup file and then access the folder and files you need.
- If the backup file resides in archival storage, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the folder from a newer backup file that has not been archived, or you can restore the entire volume from the archived backup and then access the folder and files you need.
- With ONTAP 9.15.1, you can restore FlexGroup folders using the "Browse and restore" option. This feature is in a Technology Preview mode.

You can test it using a special flag described in the [BlueXP backup and recovery July 2024 Release blog](#).


Prerequisites

- The ONTAP version must be 9.6 or greater to perform *file* restore operations.
- The ONTAP version must be 9.11.1 or greater to perform *folder* restore operations. ONTAP version 9.13.1 is required if the data is in archival storage, or if the backup file is using DataLock and Ransomware protection.
- The ONTAP version must be 9.15.1 p2 or greater to restore FlexGroup directories using the Browse and restore option.

Folder and file restore process

The process goes like this:

1. When you want to restore a folder, or one or more files, from a volume backup, click the **Restore** tab, and click **Restore Files or Folder** under *Browse & Restore*.
2. Select the source working environment, volume, and backup file in which the folder or file(s) reside.
3. BlueXP backup and recovery displays the folders and files that exist within the selected backup file.
4. Select the folder or file(s) that you want to restore from that backup.
5. Select the destination location where you want the folder or file(s) to be restored (the working environment, volume, and folder), and click **Restore**.
6. The file(s) are restored.




As you can see, you need to know the working environment name, volume name, backup file date, and folder/file name to perform a folder or file restore.

Restore folders and files

Follow these steps to restore folders or files to a volume from an ONTAP volume backup. You should know the name of the volume and the date of the backup file that you want to use to restore the folder or file(s). This functionality uses Live Browsing so that you can view the list of directories and files within each backup file.


The following video shows a quick walkthrough of restoring a single file:



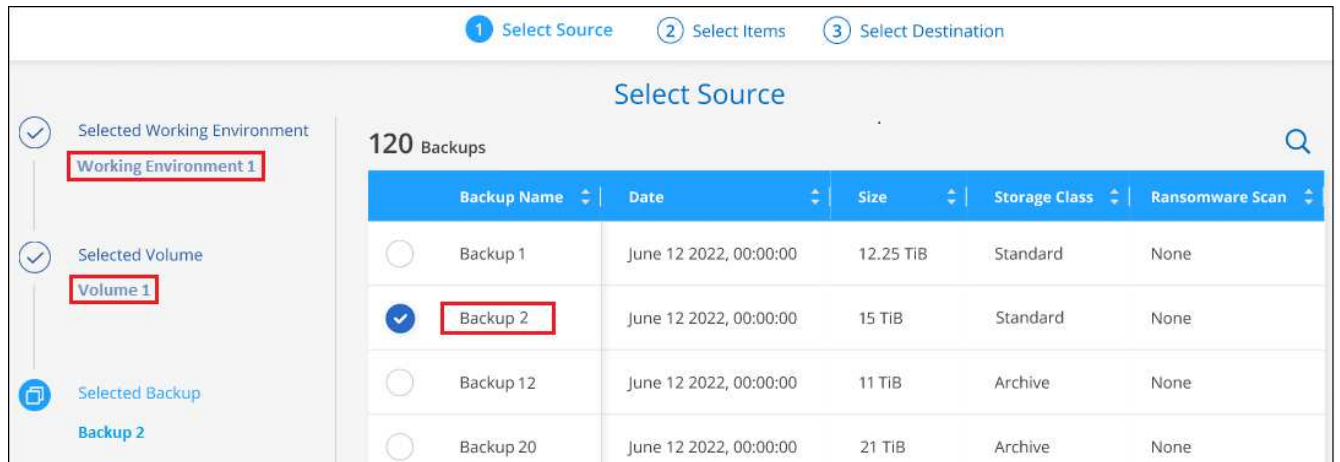
Steps

1. From the BlueXP menu, select **Protection > Backup and recovery**.
2. Click the **Restore** tab and the Restore Dashboard is displayed.

3. From the *Browse & Restore* section, click **Restore Files or Folder**.



4. In the *Select Source* page, navigate to the backup file for the volume that contains the folder or files you want to restore. Select the **Working Environment**, the **Volume**, and the **Backup** that has the date/time stamp from which you want to restore files.



5. Click **Next** and the list of folders and files from the volume backup are displayed.

If you are restoring folders or files from a backup file that resides in an archival storage tier, then you can select the Restore Priority.

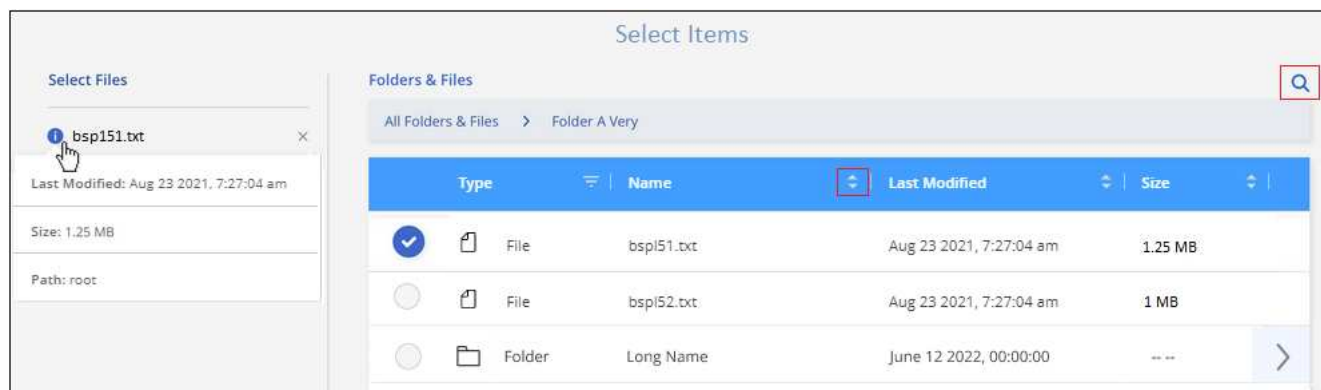
[Learn more about restoring from AWS archival storage.](#)

[Learn more about restoring from Azure archival storage.](#)


[Learn more about restoring from Google archival storage.](#) Backup files in the Google Archive storage tier are restored almost immediately, and require no Restore Priority.

And if ransomware protection is active for the backup file (if you enabled DataLock and Ransomware Protection in the backup policy), then you are prompted to run an additional ransomware scan on the backup file before restoring the data. We recommend that you scan the backup file for ransomware. (You'll

incur extra egress costs from your cloud provider to access the contents of the backup file.)




6. In the *Select Items* page, select the folder or file(s) that you want to restore and click **Continue**. To assist you in finding the item:

- You can click the folder or file name if you see it.
- You can click the search icon and enter the name of the folder or file to navigate directly to the item.
- You can navigate down levels in folders using the  button at the end of the row to find specific files.

As you select files they are added to the left side of the page so you can see the files that you have already chosen. You can remove a file from this list if needed by clicking the **x** next to the file name.

7. In the *Select Destination* page, select the **Working Environment** where you want to restore the items.

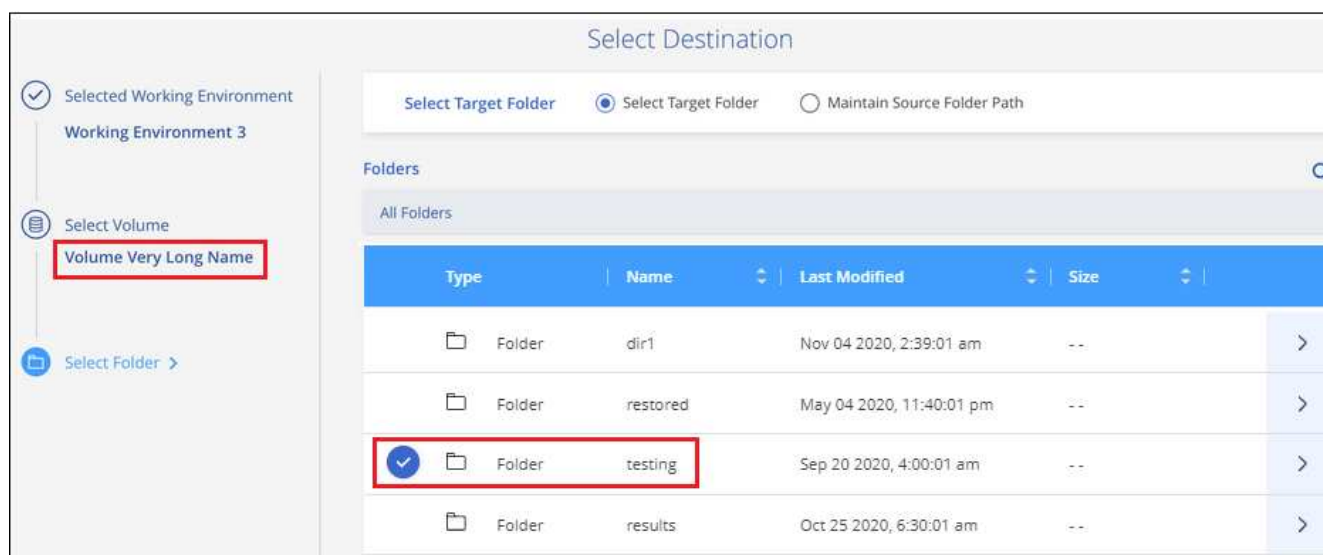


If you select an on-premises cluster and you haven't already configured the cluster connection to the object storage, you are prompted for additional information:


- When restoring from Amazon S3, enter the IPspace in the ONTAP cluster where the destination volume resides, and the AWS Access Key and Secret Key needed to access the object storage. You can also select a Private Link Configuration for the connection to the cluster.
- When restoring from Azure Blob, enter the IPspace in the ONTAP cluster where the destination volume resides. You can also select a Private Endpoint Configuration for the connection to the cluster.
- When restoring from Google Cloud Storage, enter the IPspace in the ONTAP cluster where the destination volumes reside, and the Access Key and Secret Key needed to access the object storage.

- When restoring from StorageGRID, enter the FQDN of the StorageGRID server and the port that ONTAP should use for HTTPS communication with StorageGRID, enter the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume resides.

8. Then select the **Volume** and the **Folder** where you want to restore the folder or file(s).



You have a few options for the location when restoring folders and file(s).

- When you have chosen **Select Target Folder**, as shown above:
 - You can select any folder.
 - You can hover over a folder and click  at the end of the row to drill down into subfolders, and then select a folder.
- If you have selected the same destination Working Environment and Volume as where the source folder/file was located, you can select **Maintain Source Folder Path** to restore the folder, or file(s), to the same folder where they existed in the source structure. All the same folders and sub-folders must already exist; folders are not created. When restoring files to their original location, you can choose to overwrite the source file(s) or to create new file(s).

9. Click **Restore** and you are returned to the Restore Dashboard so you can review the progress of the restore operation. You can also click the **Job Monitoring** tab to see the restore progress.

Restore ONTAP data using Search & Restore

You can restore a volume, folder, or files from an ONTAP backup file using Search & Restore. Search & Restore enables you to search for a specific volume, folder, or file from all backups, and then perform a restore. You don't need to know the exact working environment name, volume name, or file name - the search looks through all volume backup files.

The search operation looks across all local Snapshot copies that exist for your ONTAP volumes, all replicated volumes on secondary storage systems, and all backup files that exist in object storage. Since restoring data from a local Snapshot copy or replicated volume can be faster and less costly than restoring from a backup file in object storage, you may want to restore data from these other locations.

When you restore a *full volume* from a backup file, BlueXP backup and recovery creates a *new* volume using the data from the backup. You can restore the data as a volume in the original working environment, to a different working environment that's located in the same cloud account as the source working environment, or

to an on-premises ONTAP system.

You can restore *folders or files* to the original volume location, to a different volume in the same working environment, to a different working environment that's using the same cloud account, or to a volume on an on-premises ONTAP system.

When using ONTAP 9.13.0 or greater, you can restore a folder along with all files and sub-folders within it. When using a version of ONTAP before 9.13.0, only files from that folder are restored - no sub-folders, or files in sub-folders, are restored.

If the backup file for the volume that you want to restore resides in archival storage (available starting with ONTAP 9.10.1), the restore operation will take a longer amount of time and will incur additional cost. Note that the destination cluster must also be running ONTAP 9.10.1 or greater for volume restore, 9.11.1 for file restore, 9.12.1 for Google Archive and StorageGRID, and 9.13.1 for folder restore.

[Learn more about restoring from AWS archival storage.](#)

[Learn more about restoring from Azure archival storage.](#)

[Learn more about restoring from Google archival storage.](#)



- If the backup file in object storage has been configured with DataLock & Ransomware protection, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the entire volume from the backup file and then access the folder and files you need.
- If the backup file in object storage resides in archival storage, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the folder from a newer backup file that has not been archived, or you can restore the entire volume from the archived backup and then access the folder and files you need.
- The "High" restore priority is not supported when restoring data from Azure archival storage to StorageGRID systems.
- Restoring folders is not currently supported from volumes in ONTAP S3 object storage.

Before you start, you should have some idea of the name or location of the volume or file you want to restore.

The following video shows a quick walkthrough of restoring a single file:

Cloud Backup : Search and Restore

Indexed Catalog Preview Feature

February 2022

© 2022 NetApp, Inc. All rights reserved.



Search & Restore supported working environments and object storage providers

You can restore ONTAP data from a backup file that resides in a secondary working environment (a replicated volume) or in object storage (a backup file) to the following working environments. Snapshot copies reside on the source working environment and can be restored only to that same system.

Note: You can restore volumes and files from any type of backup file, but you can restore a folder only from backup files in object storage at this time.

Backup File Location		Destination Working Environment
Object Store (Backup)	Secondary System (Replication)	
Amazon S3	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Cloud Volumes ONTAP in AWS On-premises ONTAP system
Azure Blob	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Cloud Volumes ONTAP in Azure On-premises ONTAP system
Google Cloud Storage	Cloud Volumes ONTAP in Google On-premises ONTAP system	Cloud Volumes ONTAP in Google On-premises ONTAP system
NetApp StorageGRID	On-premises ONTAP system Cloud Volumes ONTAP	On-premises ONTAP system
ONTAP S3	On-premises ONTAP system Cloud Volumes ONTAP	On-premises ONTAP system

For Search & Restore, the Connector can be installed in the following locations:

- For Amazon S3, the Connector can be deployed in AWS or in your premises
- For Azure Blob, the Connector can be deployed in Azure or in your premises
- For Google Cloud Storage, the Connector must be deployed in your Google Cloud Platform VPC
- For StorageGRID, the Connector must be deployed in your premises; with or without internet access
- For ONTAP S3, the Connector can be deployed in your premises (with or without internet access) or in a cloud provider environment

Note that references to "on-premises ONTAP systems" includes FAS, AFF, and ONTAP Select systems.

Prerequisites

- Cluster requirements:
 - The ONTAP version must be 9.8 or greater.
 - The storage VM (SVM) on which the volume resides must have a configured data LIF.
 - NFS must be enabled on the volume (both NFS and SMB/CIFS volumes are supported).
 - The SnapDiff RPC Server must be activated on the SVM. BlueXP does this automatically when you enable Indexing on the working environment. (SnapDiff is the technology that quickly identifies the file and directory differences between Snapshot copies.)

- AWS requirements:
 - Specific Amazon Athena, AWS Glue, and AWS S3 permissions must be added to the user role that provides BlueXP with permissions. [Make sure all the permissions are configured correctly.](#)

Note that if you were already using BlueXP backup and recovery with a Connector you configured in the past, you'll need to add the Athena and Glue permissions to the BlueXP user role now. They are required for Search & Restore.

- Azure requirements:
 - You must register the Azure Synapse Analytics Resource Provider (called "Microsoft.Synapse") with your Subscription. [See how to register this resource provider for your subscription.](#) You must be the Subscription **Owner** or **Contributor** to register the resource provider.
 - Specific Azure Synapse Workspace and Data Lake Storage Account permissions must be added to the user role that provides BlueXP with permissions. [Make sure all the permissions are configured correctly.](#)

Note that if you were already using BlueXP backup and recovery with a Connector you configured in the past, you'll need to add the Azure Synapse Workspace and Data Lake Storage Account permissions to the BlueXP user role now. They are required for Search & Restore.

- The Connector must be configured **without** a proxy server for HTTP communication to the internet. If you have configured an HTTP proxy server for your Connector, you can't use Search & Replace functionality.
- Google Cloud requirements:
 - Specific Google BigQuery permissions must be added to the user role that provides BlueXP with permissions. [Make sure all the permissions are configured correctly.](#)

Note that if you were already using BlueXP backup and recovery with a Connector you configured in the past, you'll need to add the BigQuery permissions to the BlueXP user role now. They are required for Search & Restore.

- StorageGRID and ONTAP S3 requirements:

Depending on your configuration, there are 2 ways that Search & Restore is implemented:

- If there are no cloud provider credentials in your account, then the Indexed Catalog information is stored on the Connector.
- If you are using a Connector in a private (dark) site, then the Indexed Catalog information is stored on the Connector (requires Connector version 3.9.25 or greater).
- If you have [AWS credentials](#) or [Azure credentials](#) in the account, then the Indexed Catalog is stored at the cloud provider, just like with a Connector deployed in the cloud. (If you have both credentials, AWS is selected by default.)

Even though you are using an on-premises Connector, the cloud provider requirements must be met for both Connector permissions and cloud provider resources. See the AWS and Azure requirements above when using this implementation.

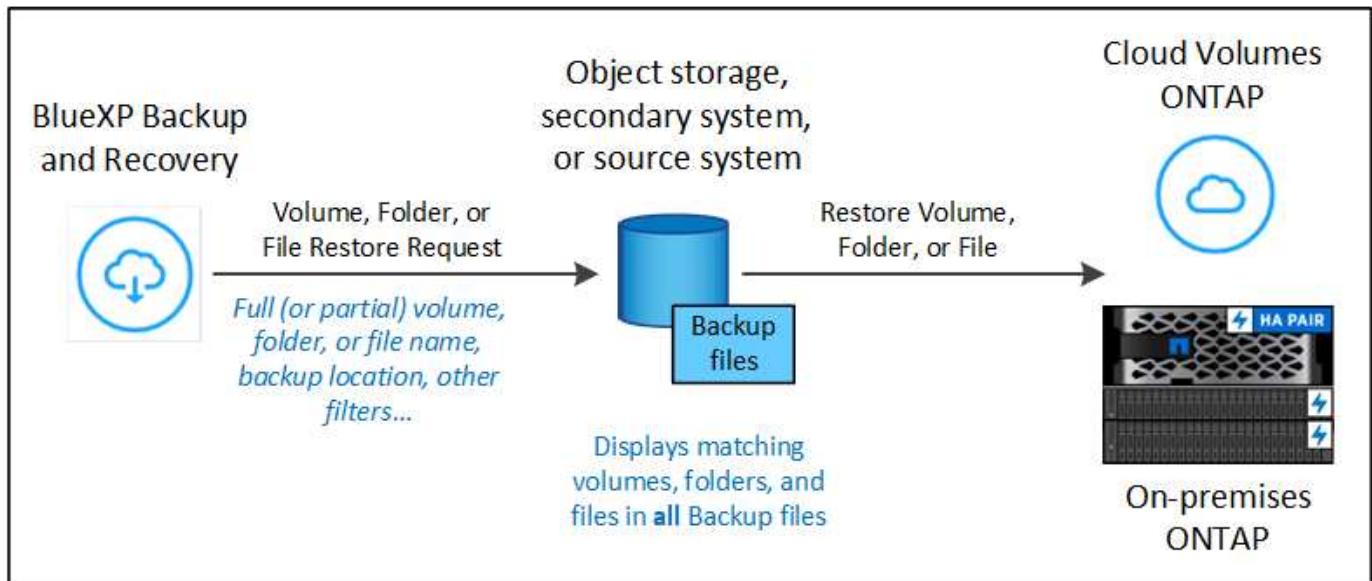
Search & Restore process

The process goes like this:

1. Before you can use Search & Restore, you need to enable "Indexing" on each source working environment from which you'll want to restore volume data. This allows the Indexed Catalog to track the backup files for every volume.
2. When you want to restore a volume or files from a volume backup, under *Search & Restore*, click **Search & Restore**.
3. Enter the search criteria for a volume, folder, or file by partial or full volume name, partial or full file name, backup location, size range, creation date range, other search filters, and click **Search**.

The Search Results page displays all the locations that have a file or volume that matches your search criteria.

4. Click **View All Backups** for the location you want to use to restore the volume or file, and then click **Restore** on the actual backup file you want to use.
5. Select the location where you want the volume, folder, or file(s) to be restored and click **Restore**.
6. The volume, folder, or file(s) are restored.



As you can see, you really only need to know a partial name and BlueXP backup and recovery searches through all backup files that match your search.

Enable the Indexed Catalog for each working environment

Before you can use Search & Restore, you need to enable "Indexing" on each source working environment from which you're planning to restore volumes or files. This allows the Indexed Catalog to track every volume and every backup file - making your searches very quick and efficient.

When you enable this functionality, BlueXP backup and recovery enables SnapDiff v3 on the SVM for your volumes, and it performs the following actions:

- For backups stored in AWS, it provisions a new S3 bucket and the [Amazon Athena interactive query service](#) and [AWS Glue serverless data integration service](#).
- For backups stored in Azure, it provisions an Azure Synapse workspace and a Data Lake file system as the container that will store the workspace data.
- For backups stored in Google Cloud, it provisions a new bucket, and the [Google Cloud BigQuery services](#) are provisioned on an account/project level.
- For backups stored in StorageGRID or ONTAP S3, it provisions space on the Connector, or on the cloud provider environment.

If Indexing has already been enabled for your working environment, go to the next section to restore your data.

To enable Indexing for a working environment:

- If no working environments have been indexed, on the Restore Dashboard under *Search & Restore*, click **Enable Indexing for Working Environments**, and click **Enable Indexing** for the working environment.
- If at least one working environment has already been indexed, on the Restore Dashboard under *Search & Restore*, click **Indexing Settings**, and click **Enable Indexing** for the working environment.

After all the services are provisioned and the Indexed Catalog has been activated, the working environment is shown as "Active".



Depending on the size of the volumes in the working environment, and the number of backup files in all 3 backup locations, the initial indexing process could take up to an hour. After that it is transparently updated hourly with incremental changes to stay current.

Restore volumes, folders, and files using Search & Restore

After you have [enabled Indexing for your working environment](#), you can restore volumes, folders, and files using Search & Restore. This allows you to use a broad range of filters to find the exact file or volume that you want to restore from all backup files.

Steps

1. From the BlueXP menu, select **Protection > Backup and recovery**.
2. Click the **Restore** tab and the Restore Dashboard is displayed.
3. From the *Search & Restore* section, click **Search & Restore**.



4. From the Search to Restore page:
 - a. In the *Search bar*, enter a full or partial volume name, folder name, or file name.
 - b. Select the type of resource: **Volumes**, **Files**, **Folders**, or **All**.
 - c. In the *Filter by* area, select the filter criteria. For example, you can select the working environment where the data resides and the file type, for example a .JPEG file. Or you can select the type of Backup Location if you want to search for results only within available Snapshot copies or backup files in object storage.
5. Click **Search** and the Search Results area displays all the resources that have a file, folder, or volume that matches your search.

Search to Restore

Browse to Restore

Search bar: mysearch 1a

Resource Type: All Resources 1b

Filter by: 1c

- Working Environment: Working Environment 2, Working Environment 4
- File Type: JPEG
- Size: August 10, 2021 - September 10, 2021
- Creation: Cloud, Primary (Local)


Backup Location: Clear All Filters

Search 2

100 Resources

Resource Name	Source Path	Size	Last Backup	Backups	
Volume 1	WorkingEnvironment\SVMName\...	2.25 GiB	June 12 2022, 00:00:00	10	3 View All Backups
Volume 2	WorkingEnvironment\SVMName\...	25.125 GiB	June 12 2022, 00:00:00	100	View All Backups

6. Locate the resource that has the data you want to restore and click **View All Backups** to display all the backup files that contain the matching volume, folder, or file.



7. Locate the backup file that you want to use to restore the data and click **Restore**.

Note that the results identify local volume Snapshot copies and remote Replicated volumes that contain the file in your search. You can choose to restore from the cloud backup file, from the Snapshot copy, or from the Replicated volume.

8. Select the destination location where you want the volume, folder, or file(s) to be restored and click **Restore**.

- For volumes, you can select the original destination working environment or you can select an alternate working environment. When restoring a FlexGroup volume you'll need to choose multiple aggregates.
- For folders, you can restore to the original location or you can select an alternate location; including the working environment, volume, and folder.
- For files, you can restore to the original location or you can select an alternate location; including the working environment, volume, and folder. When selecting the original location, you can choose to overwrite the source file(s) or to create new file(s).

If you select an on-premises ONTAP system and you haven't already configured the cluster connection to the object storage, you are prompted for additional information:

- When restoring from Amazon S3, select the IPspace in the ONTAP cluster where the destination volume will reside, enter the access key and secret key for the user you created to give the ONTAP cluster access to the S3 bucket, and optionally choose a private VPC endpoint for secure data transfer. [See details about these requirements.](#)
- When restoring from Azure Blob, select the IPspace in the ONTAP cluster where the destination volume will reside, and optionally choose a private endpoint for secure data transfer by selecting the VNet and Subnet. [See details about these requirements.](#)
- When restoring from Google Cloud Storage, select the IPspace in the ONTAP cluster where the destination volume will reside, and the Access Key and Secret Key to access the object storage. [See details about these requirements.](#)
- When restoring from StorageGRID, enter the FQDN of the StorageGRID server and the port that ONTAP should use for HTTPS communication with StorageGRID, enter the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume resides. [See details about these requirements.](#)
- When restoring from ONTAP S3, enter the FQDN of the ONTAP S3 server and the port that ONTAP should use for HTTPS communication with ONTAP S3, select the Access Key and Secret

Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume will reside. [See details about these requirements.](#)

Results

The volume, folder, or file(s) are restored and you are returned to the Restore Dashboard so you can review the progress of the restore operation. You can also click the **Job Monitoring** tab to see the restore progress.

For restored volumes, you can [manage the backup settings for this new volume](#) as required.

Copyright information

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

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

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

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

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

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

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

Trademark information

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