



BlueXP backup and recovery documentation

BlueXP backup and recovery

NetApp
February 11, 2024

This PDF was generated from <https://docs.netapp.com/us-en/bluexp-backup-recovery/index.html> on February 11, 2024. Always check docs.netapp.com for the latest.

Table of Contents

BlueXP backup and recovery documentation	1
Release notes	2
What's new with BlueXP backup and recovery	2
Known limitations	15
Get started	18
Learn about BlueXP backup and recovery	18
Set up licensing for BlueXP backup and recovery	20
Monitor data protection	27
Report on data protection coverage	27
Monitor the status of backup and restore jobs	29
Back up and restore ONTAP data	35
Protect your ONTAP volume data using BlueXP backup and recovery	35
Plan your protection journey	44
Manage backup policies for ONTAP volumes	51
Back up Cloud Volumes ONTAP data to Amazon S3	54
Back up Cloud Volumes ONTAP data to Azure Blob storage	65
Back up Cloud Volumes ONTAP data to Google Cloud Storage	76
Back up on-premises ONTAP data to Amazon S3	86
Back up on-premises ONTAP data to Azure Blob storage	102
Back up on-premises ONTAP data to Google Cloud Storage	114
Back up on-premises ONTAP data to ONTAP S3	126
Back up on-premises ONTAP data to StorageGRID	136
Manage backups for your ONTAP systems	146
Restore ONTAP data from backup files	165
Manage cluster-level backup settings	186
Back up and restore on-premises applications data	190
Protect your on-premises applications data	190
Register SnapCenter Server	191
Create a policy to back up applications	192
Back up on-premises applications data to Amazon Web Services	193
Back up on-premises applications data to Microsoft Azure	194
Back up on-premises applications data to Google Cloud Platform	195
Back up on-premises applications data to StorageGRID	196
Manage protection of applications	197
Restore on-premises applications data	201
Back up and restore cloud-native applications data	211
Protect your cloud-native applications data	211
Back up cloud-native Oracle databases	214
Back up cloud-native SAP HANA databases	227
Back up cloud-native SQL Server databases using REST APIs	236
Restore cloud-native Oracle databases	248
Restore cloud-native SAP HANA databases	250
Restore Microsoft SQL Server database	252

Clone cloud-native Oracle databases	255
Refresh SAP HANA target system	263
Manage protection of cloud-native application data	265
Back up and restore virtual machines data	271
Protect your virtual machines data	271
Register SnapCenter Plug-in for VMware vSphere host	272
Create a policy to back up datastores	273
Back up datastores to Amazon Web Services	274
Back up datastores to Microsoft Azure	275
Back up datastores to Google Cloud Platform	275
Back up datastores to StorageGRID	276
Manage protection of datastores and virtual machines data	277
Restore virtual machines data from the cloud	279
Back up and restore Kubernetes data	282
Protect your Kubernetes cluster data using BlueXP backup and recovery	282
Backing up Kubernetes persistent volume data to Amazon S3	285
Backing up Kubernetes persistent volume data to Azure Blob storage	292
Backing up Kubernetes persistent volume data to Google Cloud storage	297
Managing backups for your Kubernetes systems	302
Restoring Kubernetes data from backup files	313
BlueXP backup and recovery APIs	316
Getting started	316
Example using the APIs	318
API reference	320
Reference	322
BlueXP backup and recovery policy configuration settings	322
AWS S3 archival storage classes and restore retrieval times	330
Azure archival tiers and restore retrieval times	332
Google archival storage classes and restore retrieval times	333
Configure backup for multi-account access in Azure	333
Restore BlueXP backup and recovery data in a dark site	340
Restart the BlueXP backup and recovery service	344
Knowledge and support	346
Register for support	346
Get help	350
Legal notices	356
Copyright	356
Trademarks	356
Patents	356
Privacy policy	356
Open source	356

BlueXP backup and recovery documentation

Release notes

What's new with BlueXP backup and recovery

Learn what's new in BlueXP backup and recovery.

01 February 2024

Enhancements to BlueXP backup and recovery for Virtual Machines

- Support restoring virtual machines to an alternate location
- Support for unprotecting datastores

15 December 2023

Reports available for local Snapshot and replication Snapshot copies

Previously, you could generate reports on backup copies only. Now, you can create reports on local Snapshot copies and replication Snapshot copies as well.

With these reports, you can do the following:

- Ensure that critical data is protected according to your organizational policy.
- Ensure that backups ran smoothly for a group of volumes.
- Provide proof of protection on your production data.

Refer to [Report on data protection coverage](#).

Custom tagging available on volumes for sorting and filtering

You can now add custom tags to volumes starting in ONTAP 9.13.1 so that you can group volumes together within and across working environments. Doing this enables you to sort volumes in the BlueXP backup and recovery UI pages and filter in reports.

Catalog backups retained for 30 days

Previously, Catalog.zip backups were retained for 7 days. Now, they are retained for 30 days.

Refer to [Restore BlueXP backup and recovery data in dark sites](#).

23 October 2023

3-2-1 backup policy creation during backup activation

Previously, custom policies had to be created before you initiated a Snapshot, replication, or backup. Now you can create a policy during the backup activation process using the BlueXP backup and recovery UI.

[Learn more about policies](#).

Support for on-demand quick restore of ONTAP volumes

BlueXP backup and recovery now provides the ability to perform a "quick restore" of a volume from cloud storage to a Cloud Volumes ONTAP system. 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.

The Cloud Volumes ONTAP destination system must be running ONTAP version 9.13.0 or greater. [Learn more about restoring data.](#)

The BlueXP backup and recovery Job Monitor also shows information about the progress of quick restore jobs.

Support for scheduled jobs in the Job Monitor

The BlueXP backup and recovery Job Monitor previously monitored scheduled volume-to-object-store backup and restore jobs but not local Snapshot, replication, backup, and restore jobs that were scheduled via the UI or API.

The BlueXP backup and recovery Job Monitor now includes scheduled jobs for local Snapshots, replications, and backups to object storage.

[Learn more about the updated Job Monitor.](#)

13 October 2023

Enhancements to BlueXP backup and recovery for applications (cloud-native)

- Microsoft SQL Server database
 - Supports backup, restore, and recovery of Microsoft SQL Server databases residing on Amazon FSx for NetApp ONTAP
 - All the operations are supported only through REST APIs.
- SAP HANA systems
 - During system refresh, the auto mount and unmount of the volumes are performed using workflows instead of scripts
 - Supports addition, removal, edit, delete, maintain, and upgrade of the plug-in host using UI

Enhancements to BlueXP backup and recovery for applications (hybrid)

- Supports data lock and ransomware protection
- Supports moving backups from StorageGrid to archival tier
- Supports backing up of MongoDB, MySQL, and PostgreSQL applications data from on-premises ONTAP systems to Amazon Web Services, Microsoft Azure, Google Cloud Platform, and StorageGRID. You can restore the data when required.

Enhancements to BlueXP backup and recovery for Virtual Machines

- Support for connector proxy deployment model

11 September 2023

New policies management for ONTAP data

This release includes the ability within the UI to create custom Snapshot policies, replication policies, and policies for backups to object storage for ONTAP data.

[Learn more about policies.](#)

Support for restoring files and folder from volumes in ONTAP S3 object storage

Previously, you couldn't restore files and folders using the "Browse & Restore" feature when volumes were backed up to ONTAP S3 object storage. This release removes that restriction.

[Learn more about restoring data.](#)

Ability to archive backup data immediately instead of first writing to standard storage

Now you can send your backup files immediately to archive storage instead of writing the data to standard cloud 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 environment.

Additional support for backing up and restoring SnapLock volumes

Backup and recovery now can back up both FlexVol and FlexGroup volumes that are configured using either SnapLock Compliance or SnapLock Enterprise protection modes. Your clusters must be running ONTAP 9.14 or greater for this support. Backing up FlexVol volumes using SnapLock Enterprise mode has been supported since ONTAP version 9.11.1. Earlier ONTAP releases provide no support for backing up SnapLock protection volumes.

[Learn more about protecting ONTAP data.](#)

1 August 2023



- Because of an important security enhancement, your Connector now requires outbound internet access to an additional endpoint in order to manage backup and recovery resources within your public cloud environment. If this endpoint has not been added to the "allowed" list in your firewall you'll see an error in the UI about "Service Unavailable" or "Failed to determine service status":

<https://netapp-cloud-account.auth0.com>

- A Backup and recovery PAYGO subscription is now required when you are using the "CVO Professional" package that enables you to bundle Cloud Volumes ONTAP and BlueXP backup and recovery. This was not required in the past. No charges will be incurred on the Backup and recovery subscription for eligible Cloud Volumes ONTAP systems, but it is required when configuring backup on any new volumes.

Support has been added to back up volumes to buckets on S3-configured ONTAP systems

Now you can use an ONTAP system which has been configured for the Simple Storage Service (S3) to back up volumes to object storage. This is supported for both on-premises ONTAP systems and Cloud Volumes ONTAP systems. This configuration is supported in cloud deployments and in on-premises locations without internet access (a "private" mode deployment).

[Learn more.](#)

Now you can include existing Snapshots from a protected volume in your backup files

In the past you've had the ability to include existing Snapshot copies from read-write volumes in your initial backup file to object storage (instead of starting with the most recent Snapshot copy). Existing Snapshot copies from read-only volumes (data protection volumes) were not included in the backup file. Now you can choose to include older Snapshot copies in the backup file for "DP" volumes.

The backup wizard displays a prompt at the end of the backup steps where you can select these "existing Snapshots".

BlueXP backup and recovery no longer supports auto backup of volumes added in the future

Previously you could check a box in the backup wizard to apply the selected backup policy to all future volumes added to the cluster. This feature has been removed based on user feedback and lack of usage of this feature. You'll need to manually enable backups for any new volumes added to the cluster.

The Job Monitoring page has been updated with new features

The Job Monitoring page now provides more information related to the 3-2-1 backup strategy. The service also provides additional alert notifications related to the backup strategy.

The "Backup lifecycle" Type filter has been renamed to "Retention". Use this filter to track the backup lifecycle and to identify the expiration of all backup copies. The "Retention" job type captures all Snapshot deletion jobs initiated on a volume that is protected by BlueXP backup and recovery.

[Learn more about the updated Job Monitor.](#)

6 July 2023

BlueXP backup and recovery now includes the ability to schedule and create Snapshot copies and replicated volumes

BlueXP backup and recovery now enables you to implement a 3-2-1 strategy where you can have 3 copies of your source data on 2 different storage systems along with 1 copy in the cloud. After activation, 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

[Learn more about the new full spectrum backup and restore capabilities.](#)

This new functionality also applies to recovery operations. You can perform restore operations from a Snapshot copy, from a replicated volume, or from a backup file in the cloud. This gives you the flexibility to choose the backup file that meets your recovery requirements, including cost and speed of recovery.

Note that this new functionality and user interface is supported only for clusters running ONTAP 9.8 or greater. If your cluster has an earlier version of software, you can continue using the previous version of BlueXP backup and recovery. However, we recommend that you upgrade to a supported version of ONTAP to get the newest features and functionality. To continue using the older version of the software, follow these steps:

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings* page, click the radio button for **Display the previous BlueXP backup and recovery version**.

Then you can manage your older clusters using the previous version of software.

Ability to create your storage container for backup to object storage

When you create backup files in object storage, by default, the backup and recovery service will create the buckets in object storage for you. You can create the buckets yourself if you want to use a certain name or assign special properties. If you want to create your own bucket, you must create it before starting the activation wizard. [Learn how to create your object storage buckets.](#)

This functionality is not currently supported when creating backup files to StorageGRID systems.

04 July 2023

Enhancements to BlueXP backup and recovery for applications (cloud-native)

- SAP HANA systems
 - Supports connect and copy restore of Non-Data Volumes and Global Non-Data volumes having Azure NetApp Files secondary protection
- Oracle databases
 - Supports restore of Oracle databases on Azure NetApp Files to alternate location
 - Supports Oracle Recovery Manager (RMAN) cataloging of backups of Oracle databases on Azure NetApp Files
 - Allows you to put the database host to maintenance mode to perform maintenance tasks

Enhancements to BlueXP backup and recovery for applications (hybrid)

- Supports restore to alternate location
- Allows you to mount Oracle database backups
- Supports moving backups from GCP to archival tier

Enhancements to BlueXP backup and recovery for virtual machines (hybrid)

- Supports protection of NFS and VMFS type of datastores
- Allows you to unregister the SnapCenter Plug-in for VMware vSphere host
- Supports refresh and discovery of latest datastores and backups

5 June 2023

FlexGroup volumes can be backed up and protected using DataLock and Ransomware protection

Backup policies for FlexGroup volumes now can use DataLock and Ransomware protection when the cluster is running ONTAP 9.13.1 or greater.

New reporting features

There is now a Reports tab where you can generate a Backup Inventory report, which includes all backups for a specific account, working environment, or SVM inventory. You can also create a Data Protection Job Activity report, which provides information about Snapshot, backup, clone, and restore operations that can help you with service level agreement monitoring. Refer to [Report on data protection coverage.](#)

Job Monitor enhancements

You can now review *backup lifecycle* as a Job Type on the Job Monitor page, helping you to track the entire backup lifecycle. You can also see details of all operations on the BlueXP Timeline. Refer to [Monitor the status of backup and restore jobs](#).

Additional notification alert for unmatched policy labels

A new backup alert has been added: "Backup files were not created because Snapshot policy labels do not match". If the *label* defined in a Backup policy does not have a matching *label* in the Snapshot policy, then no backup file will be created. You'll need to use System Manager or the ONTAP CLI to add the missing label to the volume Snapshot policy.

[Review all of the alerts that BlueXP backup and recovery can send.](#)

Automatic back up of critical BlueXP backup and recovery files in dark sites

When you're using BlueXP backup and recovery in a site with no internet access, known as a "private mode" deployment, the BlueXP backup and recovery information is stored only on the local Connector system. This new functionality automatically backs up critical BlueXP backup and recovery data to a bucket on your connected StorageGRID system so that you can restore this data onto a new Connector if necessary. [Learn more](#)

8 May 2023

Folder-level restore operations are now supported from archival storage and from locked backups

If a backup file has been configured with DataLock & Ransomware protection, or if the backup file resides in archival storage, now folder-level restore operations are supported if the cluster is running ONTAP 9.13.1 or greater.

Cross-region and cross-project customer-managed keys are supported when backing up volumes to Google Cloud

Now you can choose a bucket that's in a different project than the project of your customer-managed encryption keys (CMEK). [Learn more about setting up your own customer-managed encryption keys.](#)

AWS China regions are now supported for backup files

The AWS China Beijing (cn-north-1) and Ningxia (cn-northwest-1) regions are now supported as destinations for your backup files if the cluster is running ONTAP 9.12.1 or greater.

Note that the IAM policies assigned to the BlueXP Connector need to change the AWS Resource Name "arn" under all *Resource* sections from "aws" to "aws-cn"; for example "arn:aws-cn:s3:::netapp-backup-*". See [Backing up Cloud Volumes ONTAP data to Amazon S3](#) and [Backing up on-prem ONTAP data to Amazon S3](#) for details.

Enhancements to the Job Monitor

System-initiated jobs, such ongoing backup operations, are now available in the **Job Monitoring** tab for on-premises ONTAP systems running ONTAP 9.13.1 or greater. Earlier ONTAP versions will display only user-initiated jobs.

14 April 2023

Enhancements to BlueXP backup and recovery for applications (cloud-native)

- SAP HANA databases
 - Supports script based system refresh
 - Supports Single-File-Snapshot-Restore if Azure NetApp Files backup is configured
 - Supports plug-in upgrade
- Oracle databases
 - Enhancements to plug-in deployment by simplifying non-root sudo user configuration
 - Supports plug-in upgrade
 - Supports auto-discovery and policy driven protection of Oracle databases on Azure NetApp Files
 - Supports restore of Oracle database to original location with granular recovery

Enhancements to BlueXP backup and recovery for applications (hybrid)

- BlueXP backup and recovery for applications (hybrid) is driven from the SaaS control plane
- Modified the hybrid REST APIs to align with cloud-native APIs.
- Supports email notification

4 April 2023

Ability to back up data to the cloud from Cloud Volumes ONTAP systems in "Restricted" mode

Now you can back up data from Cloud Volumes ONTAP systems installed in AWS, Azure, and GCP commercial regions in "restricted mode". This requires that you first install the Connector in the "restricted" commercial region. [Learn more about BlueXP deployment modes](#). See [Backing up Cloud Volumes ONTAP data to Amazon S3](#) and [Backing up Cloud Volumes ONTAP data to Azure Blob](#).

Ability to back up your on-premises ONTAP volumes to ONTAP S3 using the API

New functionality in the APIs enable you to back up your volume snapshots to ONTAP S3 using BlueXP backup and recovery. This functionality is available only for On-Premises ONTAP systems at this time. For detailed instructions, see the Blog [Integration with ONTAP S3 as a destination](#).

Ability to change the zone-redundancy aspect of your Azure storage account from LRS to ZRS

When creating backups from Cloud Volumes ONTAP systems to Azure storage, 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 your data to be replicated between different zones. See the Microsoft instructions for [changing how your storage account is replicated](#).

Enhancements to the Job Monitor

- Both user-initiated backup and restore operations initiated from the BlueXP backup and recovery UI and API, and system-initiated jobs, such ongoing backup operations, are now available in the **Job Monitoring** tab for Cloud Volumes ONTAP systems running ONTAP 9.13.0 or greater. Earlier ONTAP versions will display only user-initiated jobs.
- In addition to being able to download a CSV file for reporting on all jobs, now you can download a JSON

file for a single job and see its details. [Learn more](#).

- Two new backup job alerts have been added: "Scheduled job failure" and "Restore job completes but with warnings". [Review all of the alerts that BlueXP backup and recovery can send](#).

9 March 2023

Folder-level restore operations now include all sub-folders and files

In the past when you restored a folder, only files from that folder were restored - no sub-folders, or files in sub-folders, were restored. Now, if you are using ONTAP 9.13.0 or greater, all the sub-folders and files in the selected folder are restored. This can save a great deal of time and money in cases where you have multiple nested folders in a top-level folder.

Ability to back up data from Cloud Volumes ONTAP systems in sites with limited outbound connectivity

Now you can back up data from Cloud Volumes ONTAP systems installed in AWS and Azure commercial regions to Amazon S3 or Azure Blob. This requires that you install the Connector in "restricted mode" on a Linux host in the commercial region, and that you deploy the Cloud Volumes ONTAP system there as well. See [Backing up Cloud Volumes ONTAP data to Amazon S3](#) and [Backing up Cloud Volumes ONTAP data to Azure Blob](#).

Multiple enhancements to the Job Monitor

- The Job Monitoring page has added advanced filtering so you can search for backup and restore jobs by time, workload (volumes, applications, virtual machines, or Kubernetes), job type, status, working environment, and storage VM. You can also enter free text to search for any resource, for example, "application_3". [See how to use the advanced filters](#).
- Both user-initiated backup and restore operations initiated from the BlueXP backup and recovery UI and API, and system-initiated jobs, such ongoing backup operations, are now available in the **Job Monitoring** tab for Cloud Volumes ONTAP systems running ONTAP 9.13.0 or greater. Earlier versions of Cloud Volumes ONTAP systems, and on-premises ONTAP systems, will display only user-initiated jobs at this time.

6 February 2023

Ability to move older backup files to Azure archival storage from StorageGRID systems

Now you can tier older backup files from StorageGRID systems to archival storage in Azure. This enables you to free up space on your StorageGRID systems, and save money by using an inexpensive storage class for old backup files.

This functionality is available if your on-prem cluster is using ONTAP 9.12.1 or greater and your StorageGRID system is using 11.4 or greater. [Learn more here](#).

DataLock and Ransomware protection can be configured for backup files in Azure Blob

DataLock and Ransomware Protection is now supported for backup files stored in Azure Blob. If your Cloud Volumes ONTAP or on-prem ONTAP system are running ONTAP 9.12.1 or greater, now you can lock your backup files and scan them to detect possible ransomware. [Learn more about how you can protect your backups by using DataLock and Ransomware protection](#).

Backup and restore FlexGroup volume enhancements

- Now you can choose multiple aggregates when restoring a FlexGroup volume. In the last release you could only select a single aggregate.
- FlexGroup volume restore is now supported on Cloud Volumes ONTAP systems. In the last release you could only restore to on-prem ONTAP systems.

Cloud Volumes ONTAP systems can move older backups to Google Archival storage

Backup files are initially created in the Google Standard storage class. Now you can use BlueXP backup and recovery to tier older backups to Google Archive storage for further cost optimization. The last release only supported this functionality with on-prem ONTAP clusters - now Cloud Volumes ONTAP systems deployed in Google Cloud are supported.

Volume Restore operations now enable you to select the SVM where you want to restore volume data

Now you restore volume data to different storage VMs in your ONTAP clusters. In the past there was no ability to choose the storage VM.

Enhanced support for volumes in MetroCluster configurations

When using ONTAP 9.12.1 GA or greater, backup is now supported when connected to the primary system in a MetroCluster configuration. The entire backup configuration is transferred to the secondary system so that backups to the cloud continue automatically after switchover.

[See Backup limitations for more information.](#)

9 January 2023

Ability to move older backup files to AWS S3 archival storage from StorageGRID systems

Now you can tier older backup files from StorageGRID systems to archival storage in AWS S3. This enables you to free up space on your StorageGRID systems, and save money by using an inexpensive storage class for old backup files. You can choose to tier backups to AWS S3 Glacier or S3 Glacier Deep Archive storage.

This functionality is available if your on-prem cluster is using ONTAP 9.12.1 or greater, and your StorageGRID system is using 11.3 or greater. [Learn more here.](#)

Ability to select your own customer-managed keys for data encryption on Google Cloud

When backing up data from your ONTAP systems to Google Cloud Storage, now you can select your own customer-managed keys for data encryption in the activation wizard instead of using the default Google-managed encryption keys. Just set up your customer-managed encryption keys in Google first, and then enter the details when activating BlueXP backup and recovery.

"Storage Admin" role no longer needed for the service account to create backups in Google Cloud Storage

In earlier releases, the "Storage Admin" role was required for the service account that enables BlueXP backup and recovery to access Google Cloud Storage buckets. Now you can create a custom role with a reduced set of permissions to be assigned to the service account. [See how to prepare your Google Cloud Storage for backups.](#)

Support has been added to restore data using Search & Restore in sites without internet access

If you are backing up data from an on-prem ONTAP cluster to StorageGRID in a site with no internet access, also known as a dark site or offline site, now you can use the Search & Restore option to restore data when necessary. This functionality requires that the BlueXP Connector (version 3.9.25 or greater) is deployed in the offline site.

[See how to restore ONTAP data using Search & Restore.](#)

[See how to install the Connector in your offline site.](#)

Ability to download the Job Monitoring results page as a .csv report

After you filter the Job Monitoring page to display the jobs and actions you are interested in, now you can generate and download a .csv file of that data. Then you can analyze the information, or send the report to other people in your organization. [See how to generate a Job Monitoring report.](#)

19 December 2022

Enhancements to Cloud Backup for Applications

- SAP HANA databases
 - Supports policy-based backup and restore of SAP HANA databases residing on Azure NetApp Files
 - Supports custom policies
- Oracle databases
 - Add hosts and deploy plug-in automatically
 - Supports custom policies
 - Supports policy-based backup, restore, and clone of Oracle databases residing on Cloud Volumes ONTAP
 - Supports policy-based backup and restore of Oracle databases residing on Amazon FSx for NetApp ONTAP
 - Supports restore of Oracle databases using connect-and-copy method
 - Supports Oracle 21c
 - Supports cloning of cloud-native Oracle database

Enhancements to Cloud Backup for Virtual Machines

- Virtual machines
 - Back up virtual machines from on-premises secondary storage
 - Supports custom policies
 - Supports Google Cloud Platform (GCP) to back up one or more datastores
 - Supports low cost cloud storage like Glacier, Deep Glacier, and Azure Archive

6 December 2022

Required Connector outbound internet access endpoint changes

Because of a change in Cloud Backup, you need to change the following Connector endpoints for successful

Cloud Backup operation:

Old endpoint	New endpoint
https://cloudmanager.cloud.netapp.com	https://api.bluelxp.netapp.com
https://*.cloudmanager.cloud.netapp.com	https://*.api.bluelxp.netapp.com

See the full list of endpoints for your [AWS](#), [Google Cloud](#), or [Azure](#) cloud environment.

Support for selecting the Google Archival storage class in the UI

Backup files are initially created in the Google Standard storage class. Now you can use the Cloud Backup UI to tier older backups to Google Archive storage after a certain number of days for further cost optimization.

This functionality is currently supported for on-prem ONTAP clusters using ONTAP 9.12.1 or greater. It is not currently available for Cloud Volumes ONTAP systems.

Support for FlexGroup volumes

Cloud Backup now supports backing up and restoring FlexGroup volumes. When using ONTAP 9.12.1 or greater, you can back up FlexGroup volumes to public and private cloud storage. If you have working environments that include FlexVol and FlexGroup volumes, once you update your ONTAP software, you can back up any of the FlexGroup volumes on those systems.

[See the full list of supported volume types.](#)

Ability to restore data from backups to a specific aggregate on Cloud Volumes ONTAP systems

In earlier releases you could select the aggregate only when restoring data to on-prem ONTAP systems. This functionality now works when restoring data to Cloud Volumes ONTAP systems.

2 November 2022

Ability to export older Snapshot copies into your baseline backup files

If there are any local Snapshot copies for volumes in your working environment that match your backup schedule labels (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.

This option is available when activating Cloud Backup for your working environments. You can also change this setting later in the [Advanced Settings page](#).

Cloud Backup can now be used for archiving volumes that you no longer need on the source system

Now you can delete the backup relationship for a volume. This 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. [Learn how.](#)

Support has been added to receive Cloud Backup alerts in email and in the Notification Center

Cloud Backup has been integrated into the BlueXP Notification service. You can display Cloud Backup notifications by clicking the notification bell in the BlueXP menu bar. You can also configure BlueXP to send

notifications by email as alerts so you can be informed of important system activity even when you're not logged into the system. The email can be sent to any recipients who need to be aware of backup and restore activity. [Learn how](#).

New Advanced Settings page enables you to change cluster-level backup settings

This new page enables you to change many cluster-level backup settings that you set when activating Cloud Backup for each ONTAP system. You can also modify some settings that are applied as "default" backup settings. The full set of backup settings that you can change includes:

- The storage keys that give your ONTAP system permission to access object storage
- The network bandwidth allocated to upload backups to object storage
- The automatic backup setting (and policy) for future volumes
- The archival storage class (AWS only)
- Whether historical Snapshot copies are included in your initial baseline backup files
- Whether "yearly" snapshots are removed from the source system
- The ONTAP IPspace that is connected to object storage (in case of incorrect selection during activation)

[Learn more about managing cluster-level backup settings](#).

Now you can restore backup files using Search & Restore when using an on-premises Connector

In the previous release, support was added for creating backup files to the public cloud when the Connector is deployed in your premises. In this release, support has been continued to allow using Search & Restore to restore backups from Amazon S3 or Azure Blob when the Connector is deployed in your premises. Search & Restore also supports restoring backups from StorageGRID systems to on-premises ONTAP systems now.

At this time, the Connector must be deployed in the Google Cloud Platform when using Search & Restore to restore backups from Google Cloud Storage.

Job Monitoring page has been updated

The following updates have been made to the [Job Monitoring page](#):

- A column for "Workload" is available so you can filter the page to view jobs for the following Backup services: Volumes, Applications, Virtual Machines, and Kubernetes.
- You can add new columns for "User Name" and "Job Type" if you want to view these details for a specific backup job.
- The Job Details page displays all the sub-jobs that are running to complete the main job.
- The page automatically refreshes every 15 minutes so that you'll always see the most recent job status results. And you can click the **Refresh** button to update the page immediately.

AWS cross-account backup enhancements

If you want to use a different AWS account for your Cloud Volumes ONTAP backups than you're using for the source volumes, you must add the destination AWS account credentials in BlueXP, and you must add the permissions "s3:PutBucketPolicy" and "s3:PutBucketOwnershipControls" to the IAM role that provides BlueXP with permissions. In the past you needed to configure many settings in the AWS Console - you don't need to do that anymore.

28 September 2022

Enhancements to Cloud Backup for Applications

- Supports Google Cloud Platform (GCP) and StorageGRID to back up application consistent snapshots
- Create custom policies
- Supports archival storage
- Back up SAP HANA applications
- Back up Oracle and SQL applications that are on VMware environment
- Back up applications from on-premises secondary storage
- Deactivate backups
- Unregister SnapCenter Server

Enhancements to Cloud Backup for Virtual Machines

- Supports StorageGRID to back up one or more datastores
- Create custom policies

19 September 2022

DataLock and Ransomware protection can be configured for backup files in StorageGRID systems

The last release introduced *DataLock and Ransomware Protection* for backups stored in Amazon S3 buckets. This release expands support to backup files stored in StorageGRID systems. If your cluster is using ONTAP 9.11.1 or greater, and your StorageGRID system is running version 11.6.0.3 or greater, this new backup policy option is available. [Learn more about how you can use DataLock and Ransomware protection to protect your backups.](#)

Note that you'll need to be running a Connector with version 3.9.22 or greater software. The Connector must be installed in your premises, and it can be installed in a site with or without internet access.

Folder-level restore is now available from your backup files

Now you can restore a folder from a backup file if you need access to all the files in that folder (directory or share). Restoring a folder is much more efficient than restoring an entire volume. This functionality is available for restore operations using both the Browse & Restore method and the Search & Restore method when using ONTAP 9.11.1 or greater. At this time you can select and restore only a single folder, and only files from that folder are restored - no sub-folders, or files in sub-folders, are restored.

File-level restore is now available from backups that have been moved to archival storage

In the past you could only restore volumes from backup files that had been moved to archival storage (AWS and Azure only). Now you can restore individual files from these archived backup files. This functionality is available for restore operations using both the Browse & Restore method and the Search & Restore method when using ONTAP 9.11.1 or greater.

File-level restore now provides the option to overwrite the original source file

In the past, a file restored to the original volume was always restored as a new file with the prefix "Restore_<file_name>". Now you can choose to overwrite the original source file when restoring the file to the

original location on the volume. This functionality is available for restore operations using both the Browse & Restore method and the Search & Restore method.

Drag and drop to enable Cloud Backup to StorageGRID systems

If the [StorageGRID](#) destination for your backups exists as a working environment on the Canvas, you can drag your on-prem ONTAP working environment onto the destination to initiate the Cloud Backup setup wizard.

Known limitations

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

Backup and restore limitations for ONTAP volumes

Replication limitations

- You can select only one FlexGroup volume at a time for replication. You'll need to activate backups separately for each FlexGroup volume.

There is no limitation for FlexVol volumes - you can select all FlexVol volumes in your working environment and assign the same backup policies.

- The following functionality is supported in the [BlueXP replication service](#), but not when using the replication feature of BlueXP backup and recovery:
 - There is no support for a cascade configuration where replication occurs from volume A to volume B and from volume B to volume C. Support includes replication from volume A to volume B.
 - There is no support for replicating data to and from FSx for ONTAP systems.
 - There is no support for creating a one-time replication of a volume.
- When creating replications from on-premises ONTAP systems, if the ONTAP version on the target Cloud Volumes ONTAP system is 9.8, 9.9, or 9.11, only mirror-vault policies are allowed.

Backup to object limitations

- When you create or edit a backup policy when no volumes are assigned to the policy, the number of retained backups can be a maximum of 1018. After you assign volumes to the policy, you can edit the policy to create up to 4000 backups.
- When backing up data protection (DP) volumes:
 - Relationships with the SnapMirror labels `app_consistent` and `all_source_snapshot` won't be backed up to cloud.
 - If you create local copies of Snapshots on the SnapMirror destination volume (irrespective of the SnapMirror labels used) these Snapshots will not be moved to the cloud as backups. At this time you'll need to create a Snapshot policy with the desired labels to the source DP volume in order for BlueXP backup and recovery to back them up.
- FlexGroup volume backups can't be moved to archival storage.
- FlexGroup volume backups can use DataLock and Ransomware protection if the cluster is running ONTAP 9.13.1 or greater.
- SVM-DR volume backup is supported with the following restrictions:

- Backups are supported from the ONTAP secondary only.
- The Snapshot policy applied to the volume must be one of the policies recognized by BlueXP backup and recovery, including daily, weekly, monthly, etc. The default "sm_created" policy (used for **Mirror All Snapshots**) is not recognized and the DP volume will not be shown in the list of volumes that can be backed up.
- MetroCluster support:
 - When you use ONTAP 9.12.1 GA or greater, backup is supported when connected to the primary system. The entire backup configuration is transferred to the secondary system so that backups to the cloud continue automatically after switchover. You don't need to set up backup on the secondary system (in fact, you are restricted from doing so).
 - When you use ONTAP 9.12.0 and earlier, backup is supported only from the ONTAP secondary system.
 - Backups of FlexGroup volumes are not supported at this time.
- Ad-hoc volume backup using the **Backup Now** button isn't supported on data protection volumes.
- SM-BC configurations are not supported.
- ONTAP doesn't support fan-out of SnapMirror relationships from a single volume to multiple object stores; therefore, this configuration is not supported by BlueXP backup and recovery.
- WORM/Compliance mode on an object store is supported on Amazon S3, Azure, and StorageGRID at this time. This is known as the DataLock feature, and it must be managed by using BlueXP backup and recovery settings, not by using the cloud provider interface.

Restore limitations

These limitations apply to both the Search & Restore and the Browse & Restore methods of restoring files and folders; unless called out specifically.

- Browse & Restore can restore up to 100 individual files at a time.
- Search & Restore can restore 1 file at a time.
- When using ONTAP 9.13.0 or greater, Browse & Restore and Search & Restore can restore a folder along with all files and sub-folders within it.

When using a version of ONTAP greater than 9.11.1 but before 9.13.0, the restore operation can restore only the selected folder and the files in that folder - no sub-folders, or files in sub-folders, are restored.

When using a version of ONTAP before 9.11.1, folder restore is not supported.

- Directory/folder restore is supported for data that resides in archival storage only when the cluster is running ONTAP 9.13.1 and greater.
- Directory/folder restore is supported for data that is protected using DataLock only when the cluster is running ONTAP 9.13.1 and greater.
- Directory/folder restore is not currently supported on FlexGroup volume backups.
- Directory/folder restore is not currently supported from replications and/or local snapshots.
- Restoring from FlexGroup volumes to FlexVol volumes, or FlexVol volumes to FlexGroup volumes is not supported.
- The file being restored must be using the same language as the language on the destination volume. You will receive an error message if the languages are not the same.

- The *High* restore priority is not supported when restoring data from Azure archival storage to StorageGRID systems.
- Quick restore limitations:
 - The destination location must be a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater.
 - It is not supported with backups located in archived storage.
 - FlexGroup volumes are supported only if the source system from which the cloud backup was created was running ONTAP 9.12.1 or greater.
 - SnapLock volumes are supported only if the source system from which the cloud backup was created was running ONTAP 9.11.0 or greater.

Get started

Learn about BlueXP backup and recovery

The BlueXP backup and recovery service provides efficient, secure, and cost-effective data protection for NetApp ONTAP data, Kubernetes persistent volumes, databases, and virtual machines, both on premises and in the cloud. Backups are automatically generated and stored in an object store in your public or private cloud account.

The service performs block-level, incremental-forever replication and preserves all storage efficiencies, which significantly reduces the amount of data that's replicated and stored. Additionally, you pay only for what's protected and use the lowest-cost storage tiers available, which makes BlueXP backup and recovery very cost effective.

When necessary, you can restore an entire *volume* from a backup to the same or different working environment. When backing up ONTAP data, you can also choose to restore a folder or one or more *files* from a backup to the same or different working environment.

[Learn more about BlueXP backup and recovery.](#)

Backup and recovery can be used to:

- Back up and restore ONTAP volume data from Cloud Volumes ONTAP and on-premises ONTAP systems. [See detailed features here.](#)
- Back up and restore Kubernetes persistent volumes. [See detailed features here.](#)
- Back up the application-consistent Snapshots from on-premises ONTAP systems using BlueXP backup and recovery for applications. [See detailed features here.](#)
- Back up datastores to the cloud and restore virtual machines back to the on-premises vCenter using BlueXP backup and recovery for VMware. [See detailed features here.](#)

[Watch a quick demo](#)




When the BlueXP Connector is deployed in a government region in the cloud, or in a site without internet access (a dark site), BlueXP backup and recovery only supports backup and restore operations from ONTAP systems. When using these types of deployment methods, BlueXP backup and recovery does not support backup and restore operations from Kubernetes clusters, applications, or virtual machines.

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. Volume snapshots are not included in the backup image. After the initial backup, all additional backups are incremental, which means that only changed blocks and new blocks are backed up. This keeps network traffic to a minimum.

The following image shows the relationship among components:



Where backups reside

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-clusteruuid`. 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 storage account for the object store bucket.
- In ONTAP S3, BlueXP uses an existing user account for the S3 bucket.

When are backups taken

- Hourly backups start 5 minutes past the hour, every hour.
- Daily backups start just after midnight each day.
- Weekly backups start just after midnight on Sunday mornings.
- Monthly backups start just after midnight on the first day of each month.

- Yearly backups start just after midnight on the first day of the year.

The start time is based on the time zone set on each source ONTAP system. You can't schedule backup operations at a user-specified time from the UI. For more information, contact your System Engineer.

Backup copies are associated with your NetApp account

Backup copies are associated with the [NetApp account](#) in which the BlueXP Connector resides.

If you have multiple Connectors in the same NetApp account, each Connector displays the same list of backups. That includes the backups associated with Cloud Volumes ONTAP and on-premises ONTAP instances from other Connectors.

Set up licensing for BlueXP backup and recovery

You can license BlueXP backup and recovery by purchasing a pay-as-you-go (PAYGO) or annual marketplace subscription from your cloud provider, or by purchasing a bring-your-own-license (BYOL) from NetApp. A valid license is required to activate BlueXP backup and recovery on a working environment, to create backups of your production data, and to restore backup data to a production system.

A few notes before you read any further:

- If you've already subscribed to the pay-as-you-go (PAYGO) subscription in your cloud provider's marketplace for a Cloud Volumes ONTAP system, then you're automatically subscribed to BlueXP backup and recovery as well. You won't need to subscribe again.
- The BlueXP backup and recovery bring-your-own-license (BYOL) is a floating license that you can use across all systems associated with your BlueXP account. So if you have sufficient backup capacity available from an existing BYOL license, you won't need to purchase another BYOL license.
- If you are using a BYOL license, it is recommended that you subscribe to a PAYGO subscription as well. If you back up more data than allowed by your BYOL license, or if the term of your license expires, then backup continues through your pay-as-you-go subscription - there is no disruption of service.
- When backing up on-prem ONTAP data to StorageGRID, you need a BYOL license, but there's no cost for cloud provider storage space.

[Learn more about the costs related to using BlueXP backup and recovery.](#)

30-day free trial

A BlueXP backup and recovery 30-day free trial is available if you sign up for a pay-as-you-go subscription in your cloud provider's marketplace. The free trial starts at the time that you subscribe to the marketplace listing. Note that if you pay for the marketplace subscription when deploying a Cloud Volumes ONTAP system, and then start your BlueXP backup and recovery free trial 10 days later, you'll have 20 days remaining to use the free trial.

When the free trial ends, you'll be switched over automatically to the PAYGO subscription without interruption. If you decide not to continue using BlueXP backup and recovery, just [unregister BlueXP backup and recovery from the working environment](#) before the trial ends and you won't be charged.

Use a BlueXP backup and recovery PAYGO subscription

For pay-as-you-go, you'll pay your cloud provider for object storage costs and for NetApp backup licensing costs on an hourly basis in a single subscription. You should subscribe even if you have a free trial or if you bring your own license (BYOL):

- Subscribing ensures that there's no disruption of service after your free trial ends. When the trial ends, you'll be charged hourly according to the amount of data that you back up.
- If you back up more data than allowed by your BYOL license, then data backup and restore operations continue through your pay-as-you-go subscription. For example, if you have a 10 TiB BYOL license, all capacity beyond the 10 TiB is charged through the PAYGO subscription.

You won't be charged from your pay-as-you-go subscription during your free trial or if you haven't exceeded your BYOL license.

There are a few PAYGO plans for BlueXP backup and recovery:

- A "Cloud Backup" package that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP data.
- A "CVO Professional" package that enables you to bundle Cloud Volumes ONTAP and BlueXP backup and recovery. This includes unlimited backups for the Cloud Volumes ONTAP system using the license (backup capacity is not counted against the licensed capacity). This option doesn't enable you to back up on-premises ONTAP data.

Note that this option also requires a Backup and recovery PAYGO subscription, but no charges will be incurred for eligible Cloud Volumes ONTAP systems.

- A "CVO Edge Cache" package has the same capabilities as the "CVO Professional" package, but it also includes support for the [BlueXP edge caching](#) service. You are entitled to deploy one BlueXP edge caching Edge system for each 3 TiB of provisioned capacity on the Cloud Volumes ONTAP system. This option is available through the Azure and Google Marketplaces, and it doesn't enable you to back up on-premises ONTAP data.

[Learn more about these capacity-based license packages.](#)

Use these links to subscribe to BlueXP backup and recovery from your cloud provider marketplace:

- AWS: [Go to the BlueXP Marketplace offering for pricing details.](#)
- Azure: [Go to the BlueXP Marketplace offering for pricing details.](#)
- Google Cloud: [Go to the BlueXP Marketplace offering for pricing details.](#)

Use an annual contract

Pay for BlueXP backup and recovery annually by purchasing an annual contract. They're available in 1-, 2-, or 3-year terms.

If you have an annual contract from a marketplace, all BlueXP backup and recovery consumption is charged against that contract. You can't mix and match an annual marketplace contract with a BYOL.

When using AWS, there are two annual contracts available from the [AWS Marketplace page](#) for Cloud Volumes ONTAP and on-premises ONTAP systems:

- A "Cloud Backup" plan that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP

data.

If you want to use this option, set up your subscription from the Marketplace page and then [associate the subscription with your AWS credentials](#). Note that you'll also need to pay for your Cloud Volumes ONTAP systems using this annual contract subscription since you can assign only one active subscription to your AWS credentials in BlueXP.

- A "CVO Professional" plan that enables you to bundle Cloud Volumes ONTAP and BlueXP backup and recovery. This includes unlimited backups for the Cloud Volumes ONTAP system using the license (backup capacity is not counted against the licensed capacity). This option doesn't enable you to back up on-premises ONTAP data.

See the [Cloud Volumes ONTAP licensing topic](#) to learn more about this licensing option.

If you want to use this option, you can set up the annual contract when you create a Cloud Volumes ONTAP working environment and BlueXP prompts you to subscribe to the AWS Marketplace.

When using Azure there are two annual contracts available from the [Azure Marketplace page](#) for Cloud Volumes ONTAP and on-premises ONTAP systems:

- A "Cloud Backup" plan that enables you to back up Cloud Volumes ONTAP data and on-premises ONTAP data.

If you want to use this option, set up your subscription from the Marketplace page and then [associate the subscription with your Azure credentials](#). Note that you'll also need to pay for your Cloud Volumes ONTAP systems using this annual contract subscription since you can assign only one active subscription to your Azure credentials in BlueXP.

- A "CVO Professional" plan that enables you to bundle Cloud Volumes ONTAP and BlueXP backup and recovery. This includes unlimited backups for the Cloud Volumes ONTAP system using the license (backup capacity is not counted against the licensed capacity). This option doesn't enable you to back up on-premises ONTAP data.

See the [Cloud Volumes ONTAP licensing topic](#) to learn more about this licensing option.

If you want to use this option, you can set up the annual contract when you create a Cloud Volumes ONTAP working environment and BlueXP prompts you to subscribe to the Azure Marketplace.

When using GCP, contact your NetApp sales representative to purchase an annual contract. The contract is available as a private offer in the Google Cloud Marketplace.

After NetApp shares the private offer with you, you can select the annual plan when you subscribe from the Google Cloud Marketplace during BlueXP backup and recovery activation.

Use a BlueXP backup and recovery BYOL license

Bring-your-own licenses from NetApp provide 1-, 2-, or 3-year terms. You pay only for the data that you protect, calculated by the logical used capacity (*before* any efficiencies) of the source ONTAP volumes which are being backed up. This capacity is also known as Front-End Terabytes (FETB).

The BYOL BlueXP backup and recovery license is a floating license where the total capacity is shared across all systems associated with your BlueXP account. For ONTAP systems, you can get a rough estimate of the capacity you'll need by running the CLI command `volume show -fields logical-used-by-afs` for the volumes you plan to back up.

If you don't have a BlueXP backup and recovery BYOL license, click the chat icon in the lower-right of BlueXP to purchase one.

Optionally, if you have an unassigned node-based license for Cloud Volumes ONTAP that you won't be using, you can convert it to a BlueXP backup and recovery license with the same dollar-equivalence and the same expiration date. [Go here for details](#).

You use the BlueXP digital wallet to manage BYOL licenses. You can add new licenses, update existing licenses, and view license status from the BlueXP digital wallet.

Obtain your BlueXP backup and recovery license file

After you've purchased your BlueXP backup and recovery (Cloud Backup) license, you activate the license in BlueXP either by entering the BlueXP backup and recovery serial number and NetApp Support Site (NSS) account, or by uploading the NetApp License File (NLF). The steps below show how to get the NLF license file if you plan to use that method.

If you're running BlueXP backup and recovery in an on-premises site that doesn't have internet access, meaning that you've deployed the BlueXP Connector in [private mode](#), you'll need to obtain the license file from an internet-connected system. Activating the license using the serial number and NetApp Support Site account is not available for private mode installations.

Before you begin

You'll need to have the following information before you start:

- BlueXP backup and recovery serial number

Locate this number from your Sales Order, or contact the account team for this information.

- BlueXP Account ID

You can find your BlueXP Account ID by selecting the **Account** drop-down from the top of BlueXP, and then clicking **Manage Account** next to your account. Your Account ID is in the Overview tab. For private mode site without internet access, use **account-DARKSITE1**.

Steps

1. Sign in to the [NetApp Support Site](#) and click **Systems > Software Licenses**.
2. Enter your BlueXP backup and recovery license serial number.

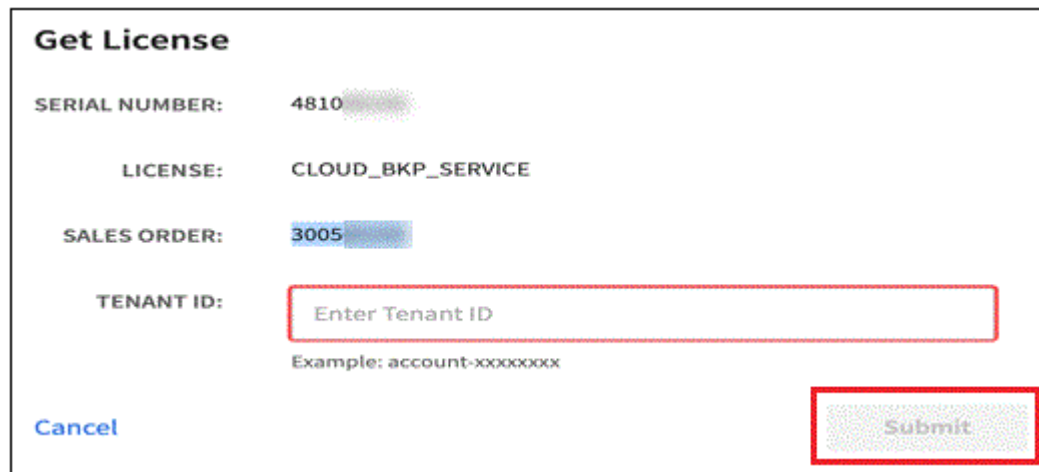
Software Licenses

Serial Number

481*

Serial #	Cluster SN	License Name	License Key	Host ID	Value	End Date
4810		CLOUD_BKP_SERVICE	Get NetApp License File		100	12/31/9998

3. In the **License Key** column, click **Get NetApp License File**.
4. Enter your BlueXP Account ID (this is called a Tenant ID on the support site) and click **Submit** to download the license file.



Get License

SERIAL NUMBER: 4810

LICENSE: CLOUD_BKP_SERVICE

SALES ORDER: 3005

TENANT ID:

Example: account-xxxxxxx

[Cancel](#) [Submit](#)

Add BlueXP backup and recovery BYOL licenses to your account

After you purchase a BlueXP backup and recovery license for your NetApp account, you need to add the license to BlueXP.

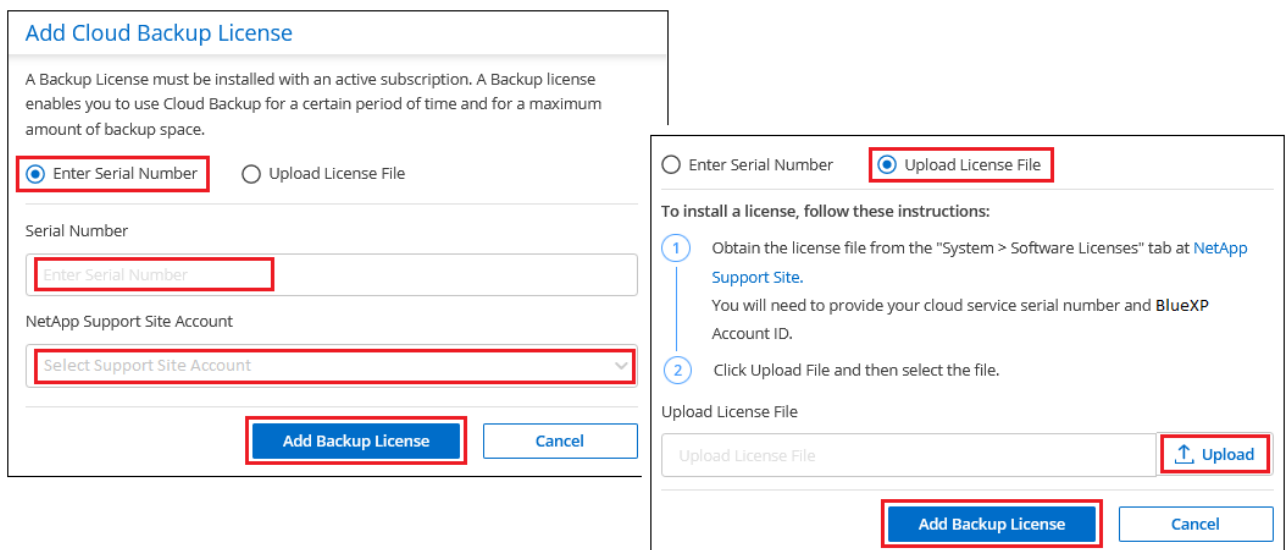
Steps

1. From the BlueXP menu, click **Governance > Digital wallet** and then select the **Data Services Licenses** tab.
2. Click **Add License**.
3. In the *Add License* dialog, enter the license information and click **Add License**:

- If you have the backup license serial number and know your NSS account, select the **Enter Serial Number** option and enter that information.

If your NetApp Support Site account isn't available from the drop-down list, [add the NSS account to BlueXP](#).

- If you have the backup license file (required when installed in a dark site), select the **Upload License File** option and follow the prompts to attach the file.



Add Cloud Backup License

A Backup License must be installed with an active subscription. A Backup license enables you to use Cloud Backup for a certain period of time and for a maximum amount of backup space.

☒ Enter Serial Number ☐ Upload License File

Serial Number

NetApp Support Site Account

[Add Backup License](#) [Cancel](#)

☐ Enter Serial Number ☒ Upload License File

To install a license, follow these instructions:

- 1 Obtain the license file from the "System > Software Licenses" tab at [NetApp Support Site](#). You will need to provide your cloud service serial number and BlueXP Account ID.
- 2 Click Upload File and then select the file.

Upload License File

[Upload](#)


[Add Backup License](#) [Cancel](#)

Result

BlueXP adds the license so that BlueXP backup and recovery is active.

Update a BlueXP backup and recovery BYOL license

If your licensed term is nearing the expiration date, or if your licensed capacity is reaching the limit, you'll be notified in the Backup UI. This status also appears in the BlueXP digital wallet page and in [Notifications](#).




You can update your BlueXP backup and recovery license before it expires so that there is no interruption in your ability to back up and restore your data.

Steps

1. Click the chat icon in the lower-right of BlueXP, or contact Support, to request an extension to your term or additional capacity to your BlueXP backup and recovery license for the particular serial number.

After you pay for the license and it is registered with the NetApp Support Site, BlueXP automatically updates the license in the BlueXP digital wallet and the Data Services Licenses page will reflect the change in 5 to 10 minutes.

2. If BlueXP can't automatically update the license (for example, when installed in a dark site), then you'll need to manually upload the license file.
 - a. You can [obtain the license file from the NetApp Support Site](#).
 - b. On the BlueXP digital wallet page *Data Services Licenses* tab, click ... for the service serial number you are updating, and click **Update License**.



- c. In the *Update License* page, upload the license file and click **Update License**.

Result

BlueXP updates the license so that BlueXP backup and recovery continues to be active.

BYOL license considerations

When using a BlueXP backup and recovery BYOL license, BlueXP displays a warning in the user interface when the size of all the data you are backing up is nearing the capacity limit or nearing the license expiration date. You'll receive these warnings:

- When backups have reached 80% of licensed capacity, and again when you have reached the limit
- 30 days before a license is due to expire, and again when the license expires

Use the chat icon in the lower right of the BlueXP interface to renew your license when you see these warnings.

Two things can happen when your BYOL license expires:

- If the account you are using has a marketplace PAYGO account, the backup service continues to run, but you are shifted over to a PAYGO licensing model. You are charged for the capacity that your backups are using.
- If the account you are using doesn't have a marketplace account, the backup service continues to run, but you will continue to see the warnings.

Once you renew your BYOL subscription, BlueXP automatically updates the license. If BlueXP can't access the license file over the secure internet connection (for example, when installed in a dark site), you can obtain the file yourself and manually upload it to BlueXP. For instructions, see [how to update a BlueXP backup and recovery license](#).

Systems that were shifted over to a PAYGO license are returned to the BYOL license automatically. And systems that were running without a license will stop seeing the warnings.

Monitor data protection

Report on data protection coverage

With BlueXP backup and recovery reports, you can ensure that critical data is protected according to your organization's defined policies and provide audits for compliance needs.

BlueXP backup and recovery reports help you accomplish the following:

- **Operations visibility:** Monitor your service level agreements regarding data protection, backup success rate, and backup window alignment to business needs.
- **Compliance and auditing:** Use operational and inventory reports in your internal and external audit processes for ongoing monitoring of compliance.



Report activities are monitored in the Job Monitoring log so that you can audit all activities.
[Learn about Job Monitoring.](#)

Reports scope

The BlueXP backup and recovery reports provide information about the following aspects:

- **Connector location:** On-premises or the cloud
- **Source:** Cloud Volumes ONTAP volumes, on-premises ONTAP volumes, applications, or Kubernetes persistent volumes
- **Destination:** Any of the cloud providers, NetApp StorageGRID, or ONTAP S3
- **ONTAP versions:** 9.13.0

Create a Backup Inventory report

From the BlueXP backup and recovery Reports tab, you can create the Backup Inventory report and filter its contents. With the Backup Inventory report, you can see all your backups for a specific account, working environment, or SVM inventory.

The Backup Inventory report shows the following information and more:

- Account, working environment, and SVM
- Protected and non-protected volumes
- Backup target
- Applied backup policy
- Encryption style (provider-managed key or user-managed key)
- DataLock and Ransomware protection status (governance, compliance, or none)
- Archive enabled status
- Count of backup copies
- Backup type (scheduled or user-initiated ad-hoc backup)

- Storage class
- Snapshot label



The Backup Inventory report doesn't include expired or failed backup information.

The top of the report includes a graph that shows the following information:

- Count of volumes in scope with at least one backup
- Total of inactive volumes plus active volumes

The Backup Inventory report shows the following charts:

- **Volume backup status:** Shows protected compared to non-protected volumes for the selected scope.
- **Volumes by backup count:** Groups volumes by the number of available backup copies for this volume.

Steps

1. From the top menu, select **Reports**.
2. Select **Backup inventory**.
3. Select **Create report**.
4. Select the account, working environment, and SVM.



You can select multiple working environments and SVMs.

5. Select the timeframe: last 24 hours, week, or month.
6. Review the report sections (Snapshot Policies, Replication Policies, or Backup Policies), depending on your report selections.
7. (Optional) Filter the results by job status.
8. (Optional) Export the report contents in .CSV format by selecting **Download CSV**.

Create a Data Protection Job Activity report

Proactive monitoring can reduce effort required to monitor all resources in your ecosystem. Beginning with ONTAP 9.13.0, the Data Protection Job Activity report provides information about snapshot, backup, clone, and restore operations that you can use with your SLA monitoring and track backup and recovery rates.

The report applies to BlueXP backup and recovery operations for Cloud Volumes ONTAP, on-premises, applications, and Kubernetes data.

The Data Protection Job Activity report shows the following information and more:

- Account, working environment, and SVM
- Job type (backup or restore)
- Resource name (volume or application)
- Job status
- Start and end times and duration
- Policy name for backup jobs

- Snapshot label for backup jobs

The charts at the top of the page show the following information:

- Jobs by type
 - Count of ONTAP volumes backup and restore jobs
 - Count of application backup and restore jobs
 - Count of virtual machine backup and restore jobs
 - Count of Kubernetes backup and restore jobs
- Daily job activity

Steps

1. From the top menu, select **Reports**.
2. Select **Data protection job activity**.
3. Select **Create report**.
4. Select the account, working environment, and SVM.
5. Select the timeframe: last 24 hours, week, or month.
6. (Optional) Filter the results by job status, job types (backup or restore), and resource.
7. (Optional) Export the report contents in .CSV format by selecting **Download CSV**.

Monitor the status of backup and restore jobs

You can monitor the status of local Snapshots, replications, and backup to object storage jobs that you initiated, and restore jobs that you initiated. You can see the jobs that have completed, are in progress, or failed so you can diagnose and fix problems. Using the BlueXP Notification Center, you can enable notifications to be sent by email so you can be informed of important system activity even when you're not logged into the system. Using the BlueXP Timeline, you can see details of all actions initiated via the UI or API.

View job status on the Job Monitor

You can view a list of all the Snapshot, replication, backup to object storage, and restore operations and their current status in the **Job Monitoring** tab. This includes operations from your Cloud Volumes ONTAP, on-premises ONTAP, applications, virtual machines, and Kubernetes systems. Each operation, or job, has a unique ID and a status.

The status can be:

- Success
- In Progress
- Queued
- Warning
- Failed

Snapshots, replications, backups to object storage and restore operations that you initiated from the BlueXP

backup and recovery UI and API are available in the Job Monitoring tab.



If you've upgraded your ONTAP systems to 9.13.x and you don't see ongoing scheduled backup operations in the Job Monitor, then you'll need to restart the BlueXP backup and recovery service. [Learn how to restart BlueXP backup and recovery.](#)

Steps

1. From the BlueXP menu, select **Protection > Backup and recovery**.
2. Select the **Job Monitoring** tab.

The screenshot shows the 'Job Monitoring' interface. At the top, there's a header with 'Job Monitoring' and a 'Last Updated' timestamp of 'July 27, 2023, 09:28:18'. Below the header, there's a section for 'Advanced Search & Filtering' and a 'Timeframe: Last Month' filter. The main content area displays a table of jobs, with a 'Jobs(31)' label and a download icon. The table has columns for Job ID, Type, Protection Type, Resource Name, Status, Job Name, and Start Time. Two jobs are visible: one for 'production_kafka1' with status 'Success' and another for 'production_kafka1' with status 'Success'.

Job ID	Type	Protection Type	Resource Name	Status	Job Name	Start Time
2639e43c-3b44-4297...	Protection	Replication	production_kafka1	Success	Replicate production_kafka1 to...	Jul 25 2023, 11:30
409e9010-fba1-4371...	Protection	Backup to Cloud	production_kafka1	Success	Initialize backup for cb53ded0...	Jul 25 2023, 11:30

This screenshot shows the default column headings.

3. To show additional columns (Working Environment, SVM, User Name, Workload, Policy Name, Snapshot Label), select **+**.

Search and filter the list of jobs

You can filter the operations on the Job Monitoring page using several filters, such as policy, Snapshot label, type of operation (protection, restore, retention, or other) and protection type (local Snapshot, replication, or backup to the cloud).

By default, the Job Monitoring page shows protection and recovery jobs from the last 24 hours. You can change the timeframe using the Timeframe filter.

Steps

1. Select the **Job Monitoring** tab.
2. To sort the results differently, select each column heading to sort by Status, Start Time, Resource Name, and more.
3. If you're looking for specific jobs, select the **Advanced Search & Filtering** area to open the Search panel.

Use this panel to enter a free text search for any resource; for example "volume 1" or "application 3". You can also filter the jobs list according to the items in the drop-down menus.

Advanced Search & Filtering

Clear All Filters

Volume_1 + Add

Timeframe Type Protection type Workload Status Working Environment

SVM

Timeframe: Last Week Free Text: Volume_1 Type: Protection Workload: Volumes

Search

This screenshot shows how you would search for all "Volume" "Backup" jobs for volumes named "Volume_1" in the "past week".


Most of the filters are self-explanatory. The filter for "Workload" enables you to view jobs in the following categories:

- Volumes (Cloud Volumes ONTAP and on-premises ONTAP volumes)
- Applications
- Virtual Machines
- Kubernetes



- You can search for data within a specific "SVM" only if you have first selected a Working Environment.
- You can search using the "Protection type" filter only when you have selected the "Type" of "Protection".

4.

To update the page immediately, select the  button. Otherwise, this page refreshes every 15 minutes so that you'll always see the most recent job status results.


View job details

You can view details corresponding to a specific completed job. You can export details for a particular job in a JSON format.

You can view details such as job type (scheduled or on-demand), SnapMirror backup type (initial or periodic) start and end times, duration, amount of transferred data from working environment to object storage, average transfer rate, policy name, retention lock enabled, ransomware scan performed, protection source details, and protection target details.

Restore jobs show details such as backup target provider (Amazon Web Services, Microsoft Azure, Google Cloud, on-premises), S3 bucket name, SVM name, source volume name, destination volume, Snapshot label, recovered objects count, file names, file sizes, last modification date, and full file path.

Steps

1. Select the **Job Monitoring** tab.
2. Select the name of the job.
3. Select the Actions menu  and select **View Details**.

Job Monitoring > <Job Name: Backup "Volume_Name_1">

Job Name: Backup "Volume_Name_1"

Job ID: e2d802f2-dc5ce2d802f2-dc5ce2d802f2-dc5c

Backup
Job Type

Source Volume Name
Backup from

AWS Bucket
Backup to

Success
Job Status

Close All

Backup from

aws	Working Environment Working Environment Name	SVM Name SVM Name	Volume Name Volume Name	FlexVol Volume Type	Snapshot Label Name Snapshot Label
-----	---	----------------------	----------------------------	------------------------	---------------------------------------

Backup to

aws	AWS Provider	N.Virginia Region	01234567890123456789 Account ID	Target Bucket Name Bucket Name
-----	-----------------	----------------------	------------------------------------	-----------------------------------

Backup Details

Success Job Status	Scheduled Backup Job Type	Snapmirror Initialize Scheduled Backup	Backup Policy Name Policy Name	Disabled Ransomware Protection
-----------------------	------------------------------	---	-----------------------------------	-----------------------------------



4. Expand each section to see details.

Download Job Monitoring results as a report

You can download the contents of the main Job Monitoring page as a report after you've refined it. BlueXP backup and recovery generates and downloads a .CSV file that you can review and send to other groups as needed. The .CSV file includes up to 10,000 rows of data.

From the Job Monitoring Details information, you can download a JSON file containing details for a single job.

Steps

1. Select the **Job Monitoring** tab.
2. To download a CSV file for all jobs, select the  button and locate the file in your download directory.
3. To download a JSON file for a single job, select the Actions menu  for the job, select **Download JSON File**, and locate the file in your download directory.

Review retention (backup lifecycle) jobs

Monitoring of retention (or *backup lifecycle*) flows helps you with audit completeness, accountability, and backup safety. To help you track the backup lifecycle, you might want to identify the expiration of all backup copies.

A backup lifecycle job tracks all Snapshot copies that are deleted or in the queue to be deleted. Beginning with ONTAP 9.13, you can look at all job types called "Retention" on the Job Monitoring page.

The "Retention" job type captures all Snapshot deletion jobs initiated on a volume that is protected by BlueXP backup and recovery.

Steps

1. Select the **Job Monitoring** tab.

2. Select the **Advanced Search & Filtering** area to open the Search panel.
3. Select "Retention" as the job type.

Review backup and restore alerts in the BlueXP Notification Center

The BlueXP Notification Center tracks the progress of backup and restore jobs that you've initiated so you can verify whether the operation was successful or not.

In addition to viewing the alerts in the Notification Center, you can configure BlueXP to send certain types of notifications by email as alerts so you can be informed of important system activity even when you're not logged into the system. [Learn more about the Notification Center and how to send alert emails for backup and restore jobs.](#)

The Notification Center displays numerous Snapshot, replication, backup to cloud, and restore events, but only certain events trigger email alerts:

Operation type	Event	Alert level	Email sent
Activation	Backup and recovery activation failed for working environment	Error	Yes
Activation	Backup and recovery edit failed for working environment	Error	Yes
Local Snapshot	BlueXP backup and recovery ad-hoc Snapshot creation job failure	Error	Yes
Replication	BlueXP backup and recovery ad-hoc replication job failure	Error	Yes
Replication	BlueXP backup and recovery replication pause job failure	Error	No
Replication	BlueXP backup and recovery replication brake job failure	Error	No
Replication	BlueXP backup and recovery replication resync job failure	Error	No
Replication	BlueXP backup and recovery replication stop job failure	Error	No
Replication	BlueXP backup and recovery replication reverse resync job failure	Error	Yes
Replication	BlueXP backup and recovery replication delete job failure	Error	Yes




Beginning with ONTAP 9.13.0, all alerts appear for Cloud Volumes ONTAP and on-premises ONTAP systems. For systems with Cloud Volumes ONTAP 9.13.0 and on-premises ONTAP, only the alert related to "Restore job completed, but with warnings" appears.

By default, BlueXP Account Admins receive emails for all "Critical" and "Recommendation" alerts. All other users and recipients are set up, by default, not to receive any notification emails. Emails can be sent to any BlueXP users who are part of your NetApp Cloud Account, or to any other recipients who need to be aware of backup and restore activity.

To receive the BlueXP backup and recovery email alerts, you'll need to select the notification severity types "Critical", "Warning", and "Error" in the Alerts and Notifications Settings page.

[Learn how to send alert emails for backup and restore jobs.](#)

Steps

1. From the BlueXP menu bar, select the .
2. Review the notifications.

Review operation activity in the BlueXP Timeline

You can view details of backup and restore operations for further investigation in the BlueXP Timeline. The BlueXP Timeline provides details of each event, whether user-initiated or system-initiated and shows actions initiated in the UI or via the API.

[Learn about the differences between the Timeline and the Notification Center.](#)


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.
- Ability to perform a "quick restore" of a volume from cloud storage to a Cloud Volumes ONTAP system; perfect for disaster recovery situations where you need to provide access to a volume as soon as possible.
- Data is restored on a block level, placing the data directly in the location you specify, all while preserving the original ACLs.
- Browsable and searchable 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 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
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

Backup File Location		Destination Working Environment
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)
- SnapLock Compliance volumes (requires ONTAP 9.14 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 using 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 using 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 using 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.](#)

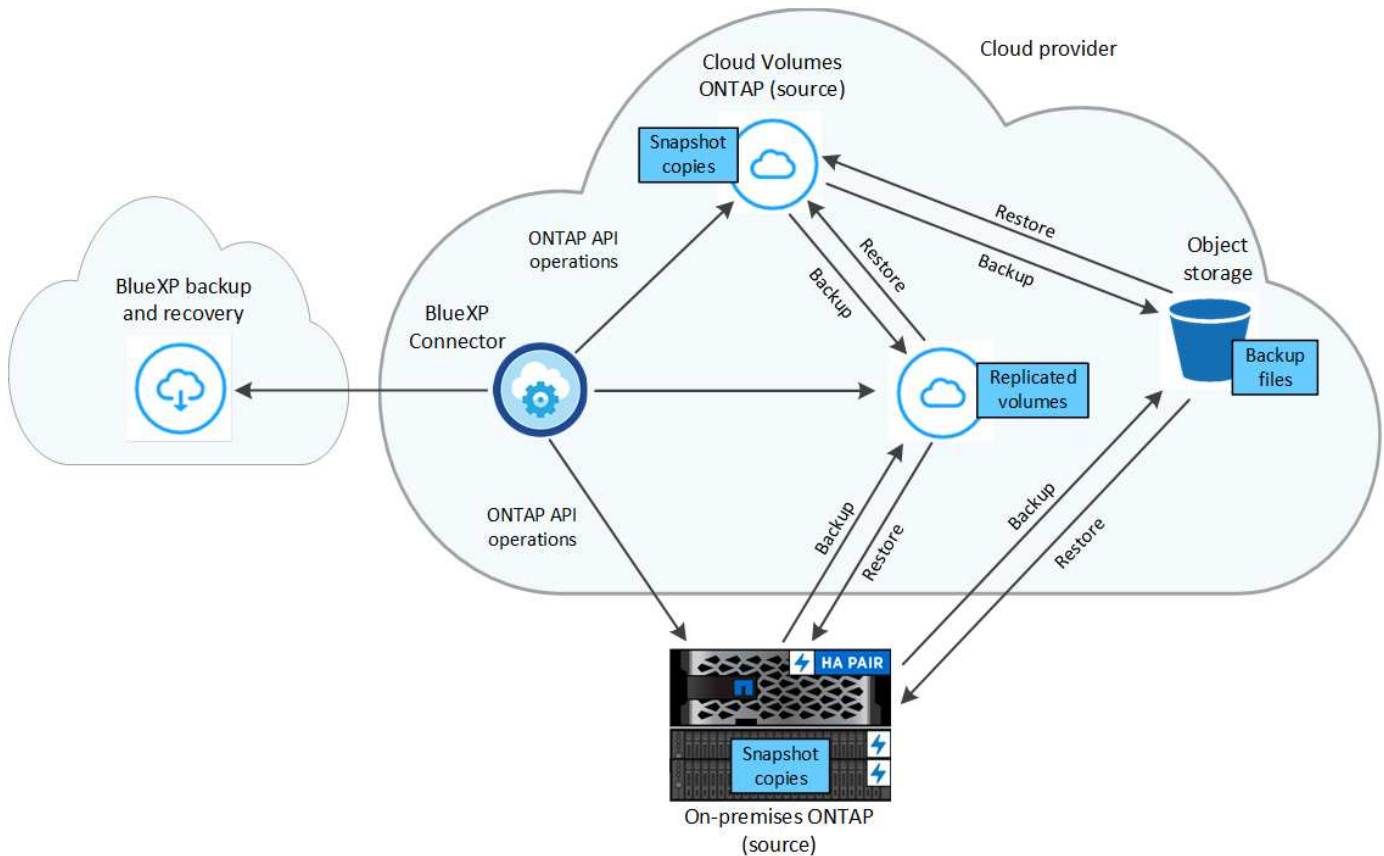
How BlueXP backup and recovery works

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.

The backup retention period is the same as the backup schedule retention period; plus 14 days. 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.

See [DataLock and Ransomware protection](#) for more details about how DataLock and Ransomware protection works.



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)

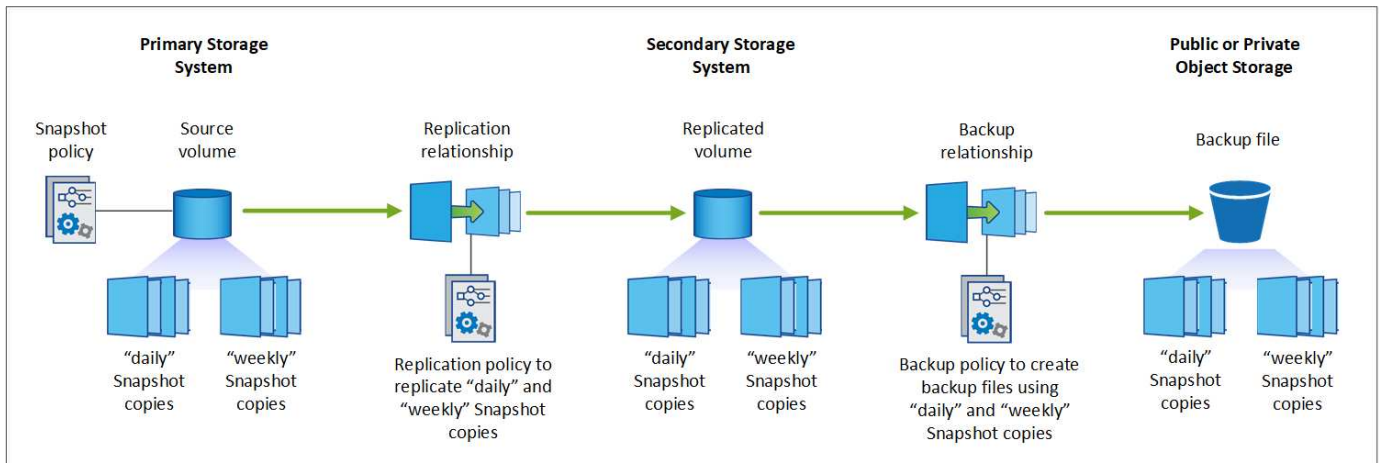
Which backup architecture will you use

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 backup 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 custom policies using BlueXP backup recovery, 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.

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.

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. See the topics for creating backups to AWS S3 for details.

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.](#)

[Watch this video 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.](#)

See the backup topic for the cloud provider where you plan to create backup files for the list of valid Connector locations. 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. See how to [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. See how to [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. See

how to [Back up on-premises ONTAP data to StorageGRID](#) for details.

- 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. See how to [Back up on-premises ONTAP data to ONTAP S3](#) for details.

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 due to 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 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.
- If you choose to create a schedule:
 - Select the frequency of hourly, daily, weekly, monthly, or yearly.
 - Specify the Snapshot labels denoting the schedule and retention.
 - Enter when and how often the Snapshot will be taken.
 - Retention: Enter the number of Snapshots to keep.
- 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](#)
[Create a backup to object storage policy using the ONTAP CLI](#)

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").

[Prepare your BlueXP Connector](#)

3

Verify license requirements

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

[Verify license requirements.](#)

4

Verify ONTAP networking requirements for replicating volumes

Ensure that the primary and secondary storage 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.

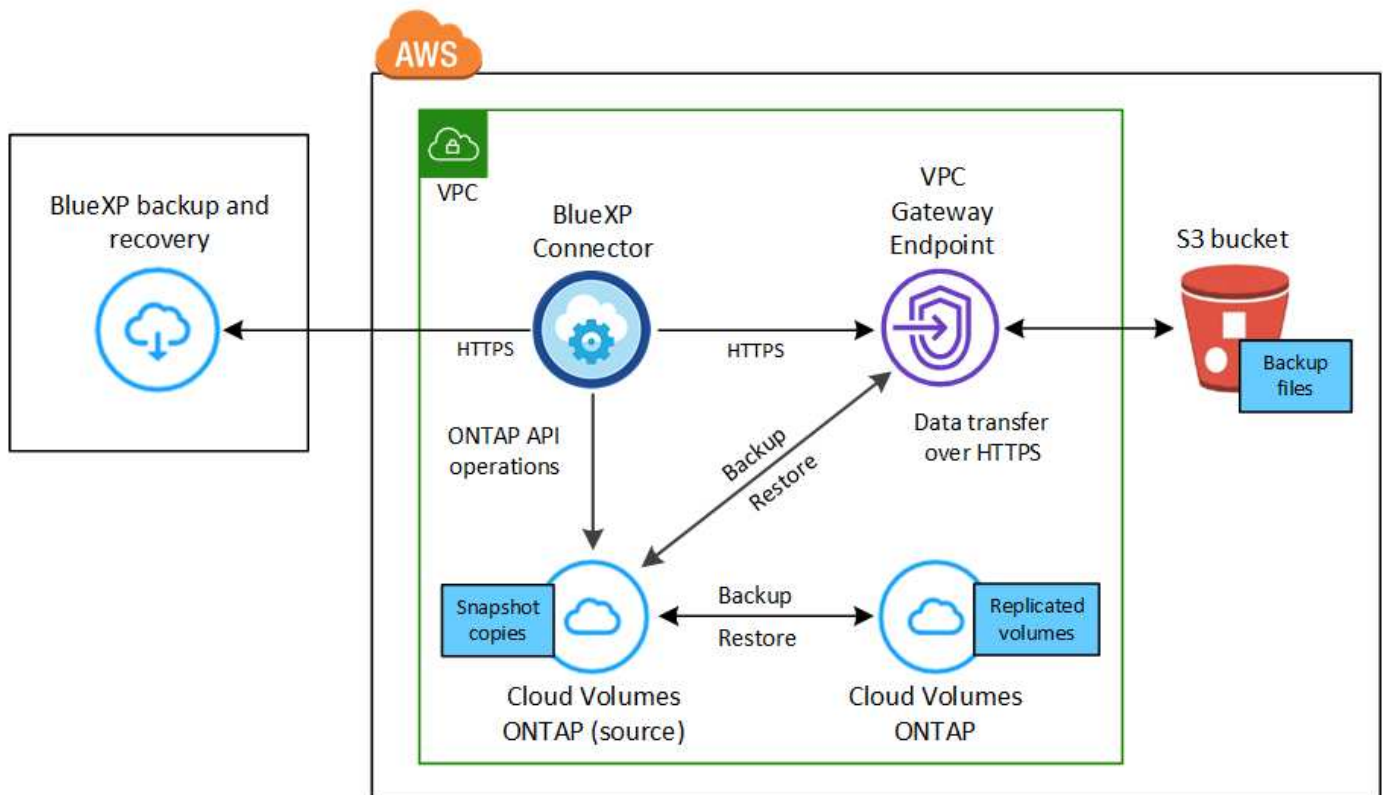
[Activate backups on your ONTAP volumes.](#)

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 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. (Volumes with SnapLock Compliance mode require ONTAP 9.14 or later.)

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.
- 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 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.

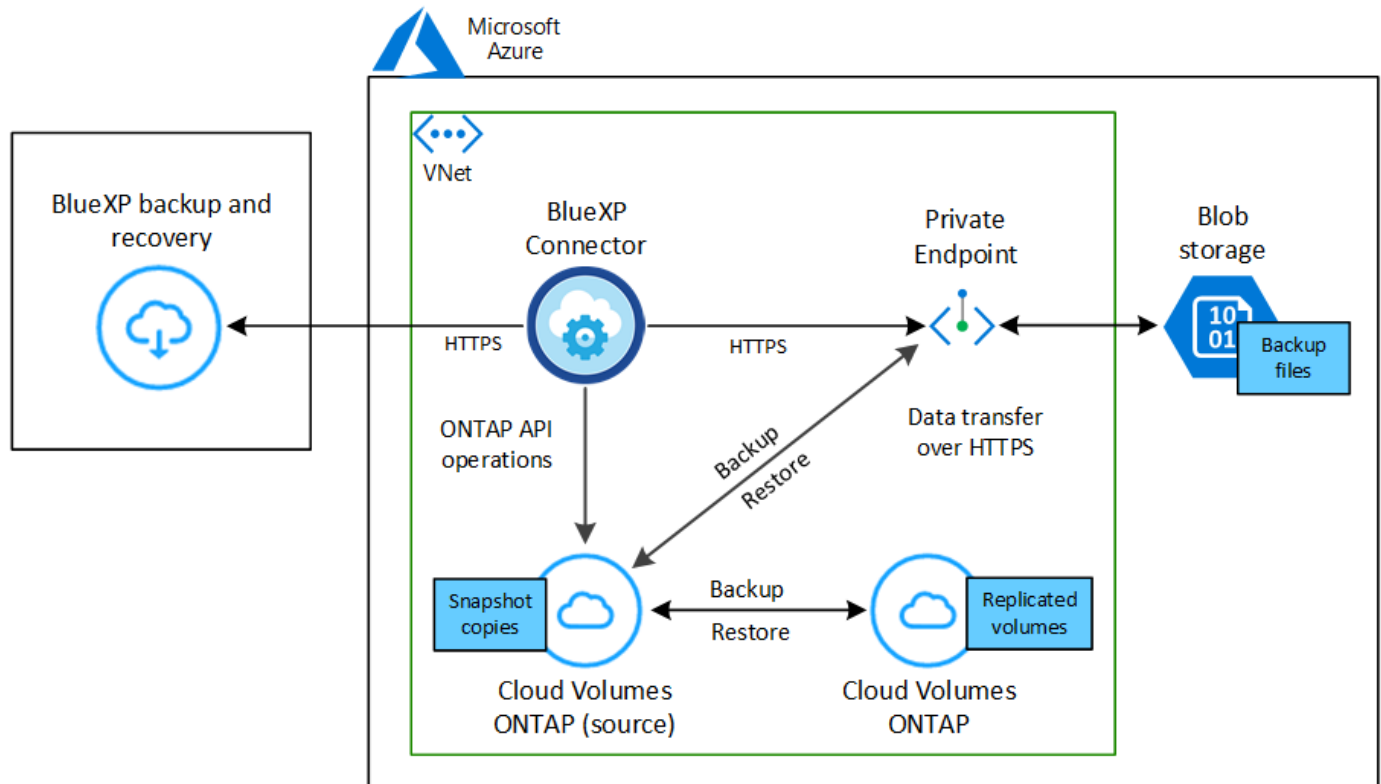
[Activate backups on your ONTAP volumes](#).

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* as the permission model. The *Azure role-based access control* (Azure RBAC) permission model is not currently supported.

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. (Volumes with SnapLock Compliance mode require ONTAP 9.14 or later.)

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.
- 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 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 the predefined Storage Admin role.
- 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").

[Prepare your BlueXP Connector](#)

3

Verify license requirements

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

[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.

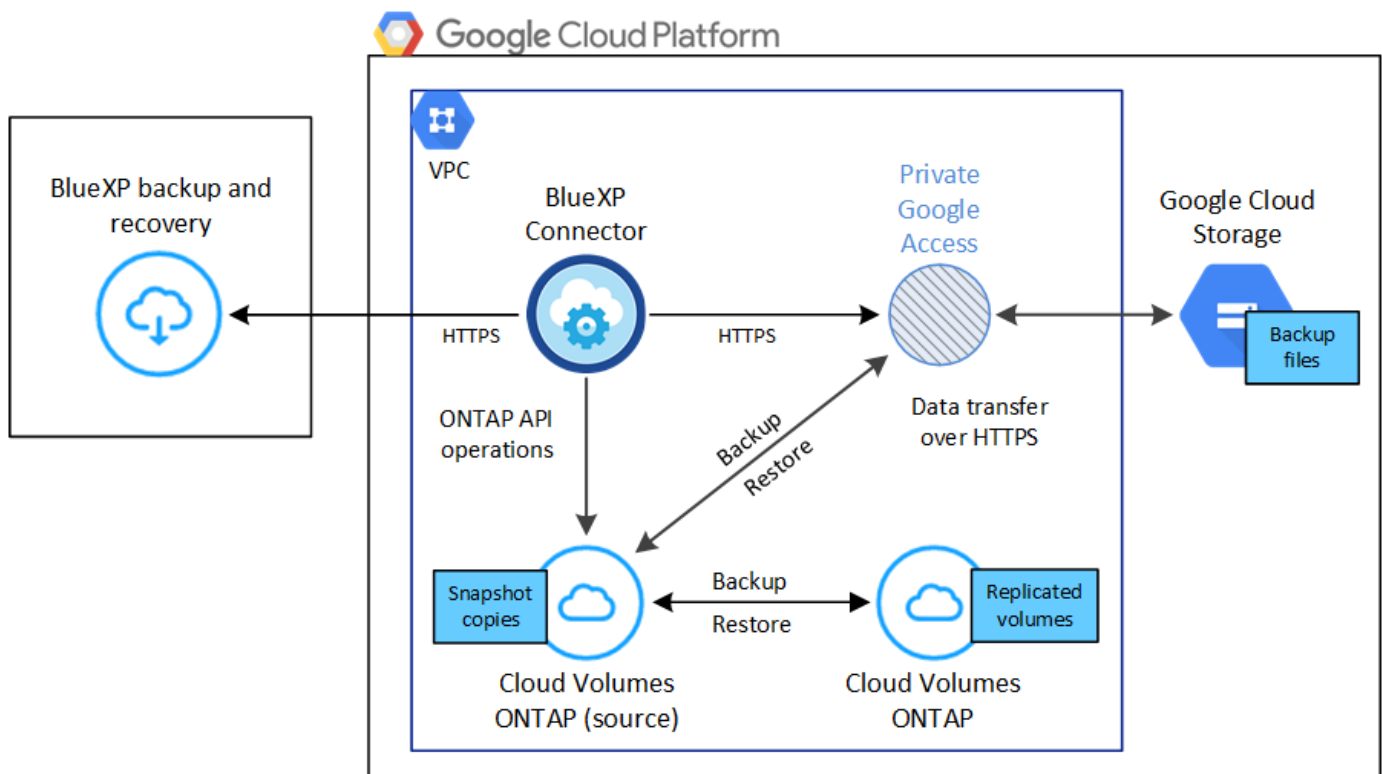
[Activate backups on your ONTAP volumes.](#)

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 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 predefined Storage Admin role. [Learn how to create a service account](#).

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.

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

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](#).

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. (Volumes with SnapLock Compliance mode require ONTAP 9.14 or later.)

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.

1

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. (Volumes with SnapLock Compliance mode require ONTAP 9.14 or later.)

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](#).

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 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**.

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 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.

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 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. (Volumes with SnapLock Compliance mode aren't currently supported require ONTAP 9.14 or later.)

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.
- 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.

[Identify the connection method.](#)

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.

[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 Google Cloud 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 Google Cloud.

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

5

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. [Learn how to get your Google Cloud environment ready to receive ONTAP backups.](#)

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 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

When you set up backup, you need to provide storage access keys for a service account that has specific permissions. 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](#)

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. (Volumes with SnapLock Compliance mode require ONTAP 9.14 or later.)

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.

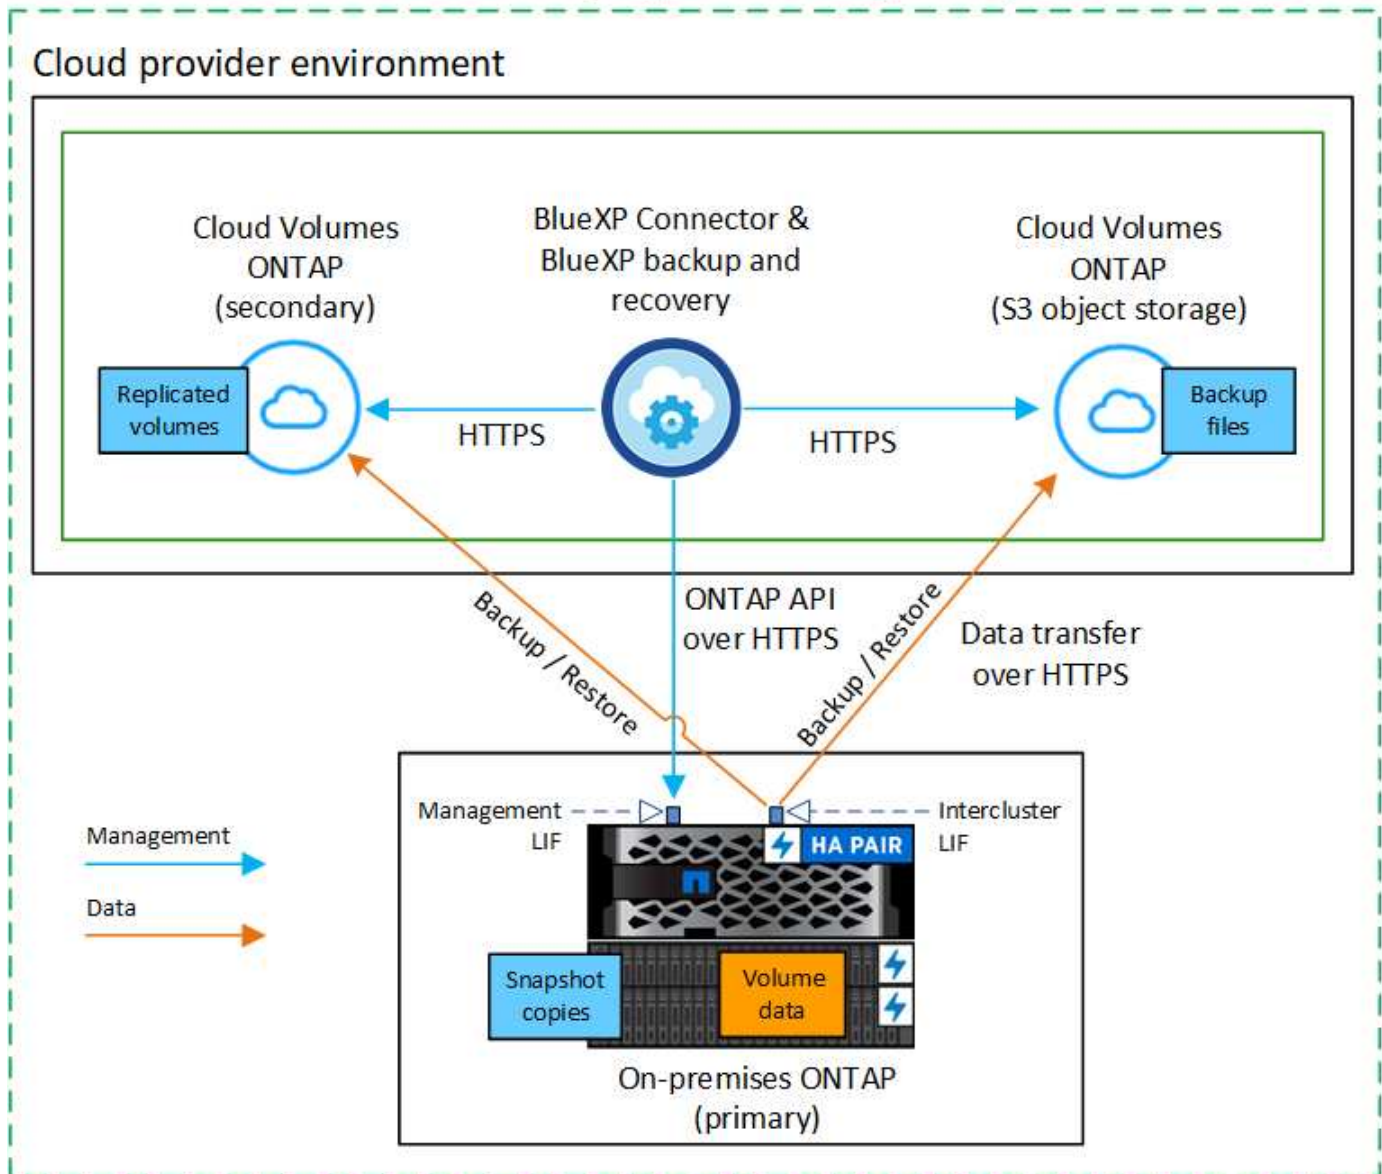
Connector installed on-premises (Public)



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. (Volumes with SnapLock Compliance mode require ONTAP 9.14 or later.)

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.
 - 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

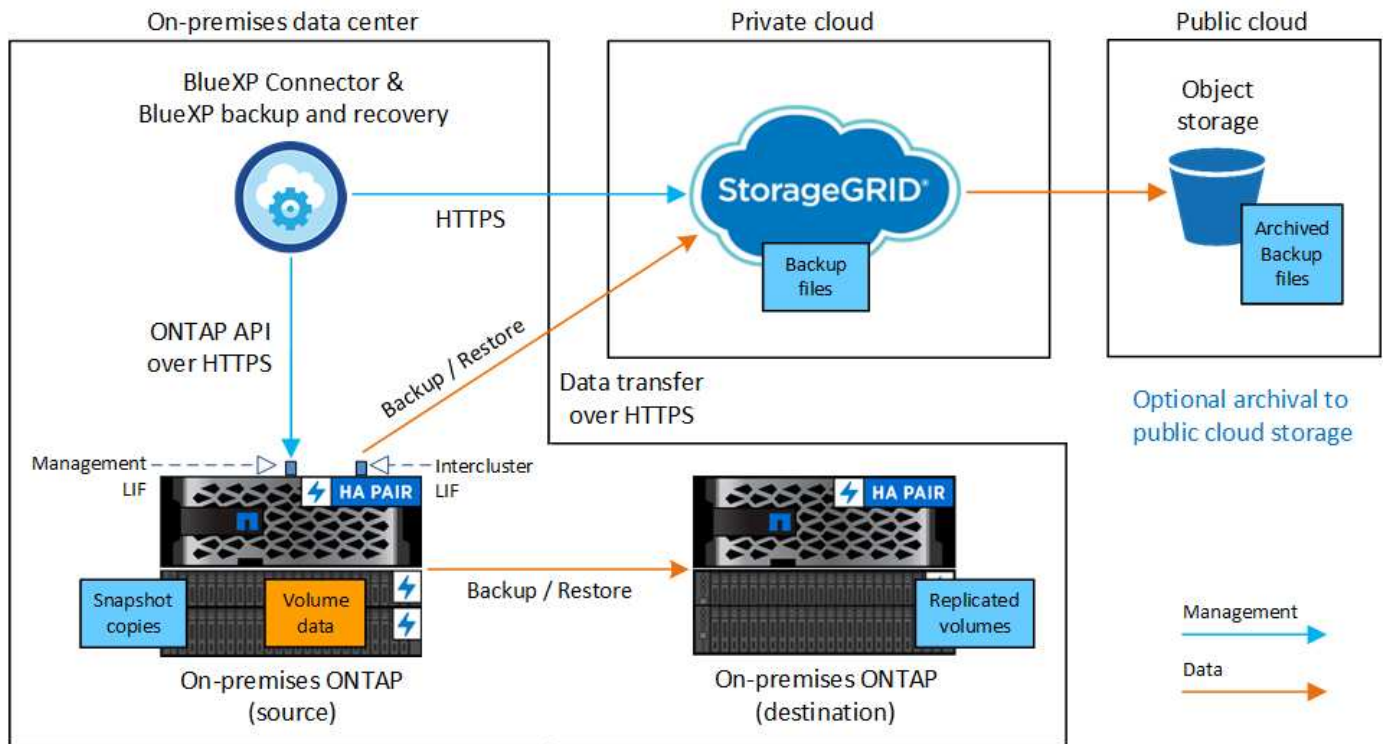
Select the working environment and click **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel. Then follow the setup wizard to select the replication and backup policies that you'll use and the volumes you want to back up.

[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. (Volumes with SnapLock Compliance mode require ONTAP 9.14 or later.)

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.
- Select **Create**.

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

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.



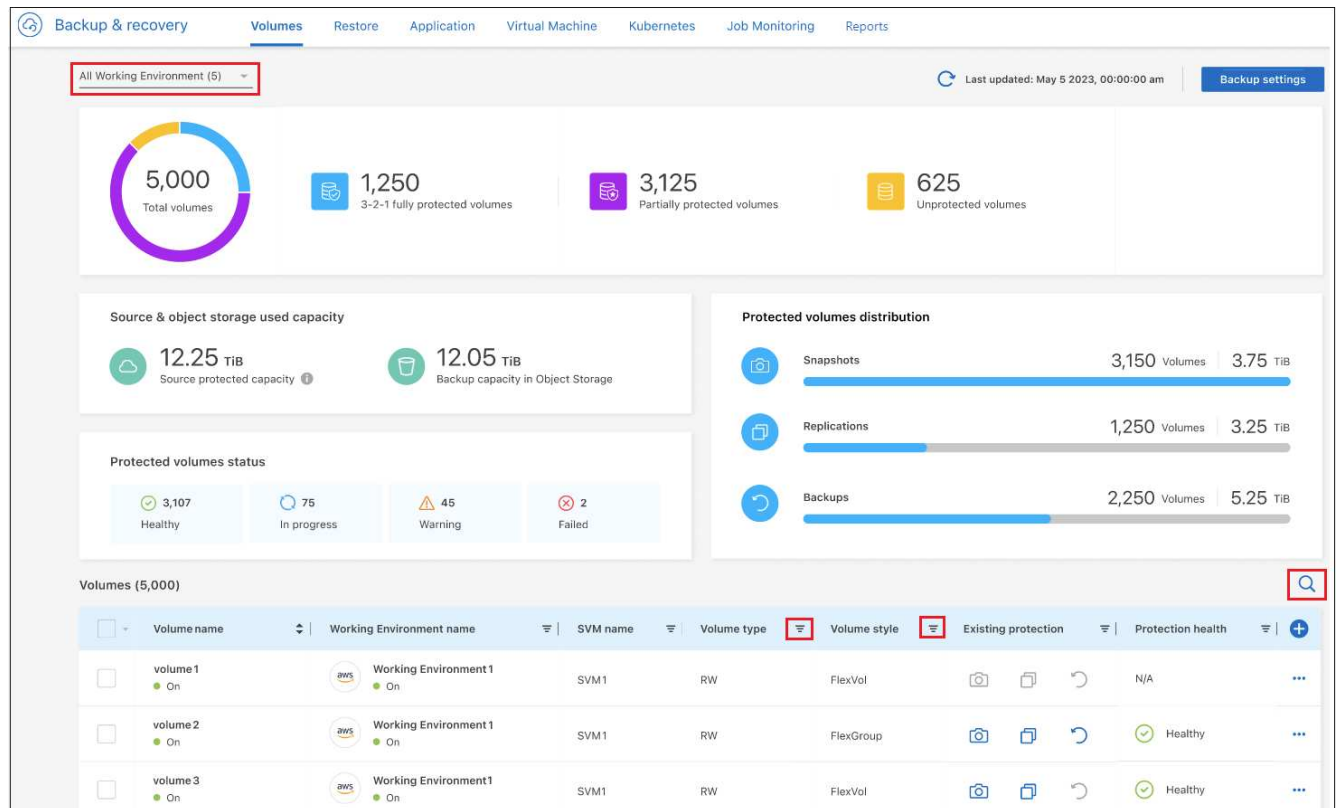
Do not manage or change backup files directly on your storage systems or from your cloud provider environment. This may corrupt the files and will result in an unsupported configuration.

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

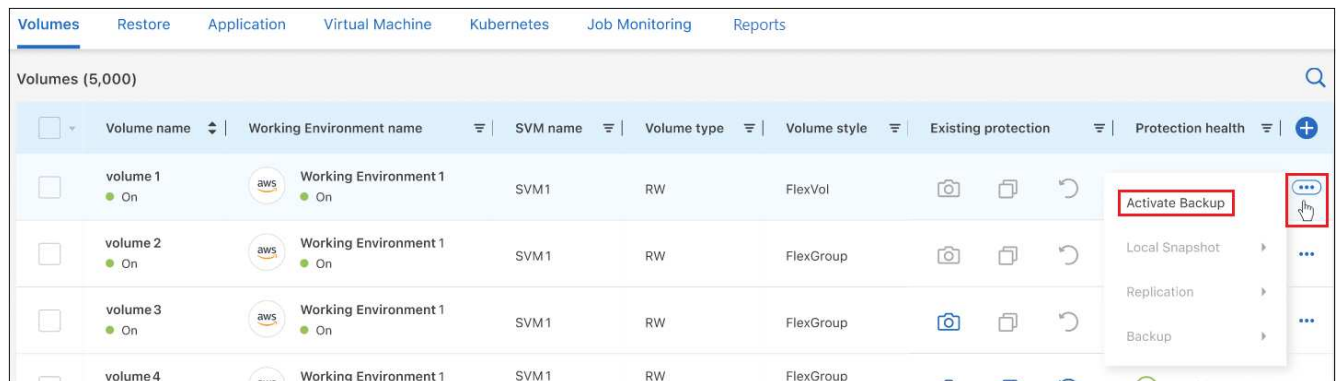
information for each type of backup.

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

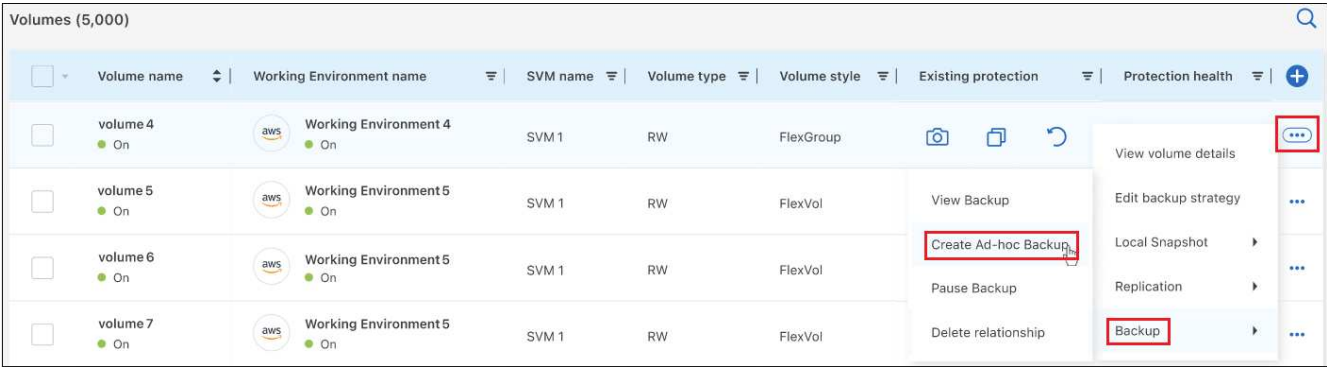
up. The name of the Snapshot will begin with `cbs-snapshot-adhoc-`. [See how to delete a Snapshot using the ONTAP CLI.](#)



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**.



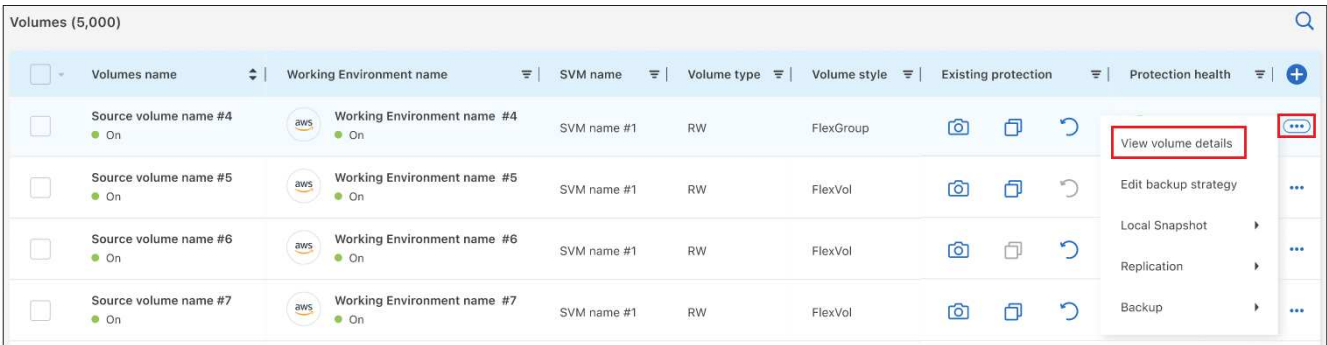
The Backup Status column for that volume displays "In Progress" until the backup is created.

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

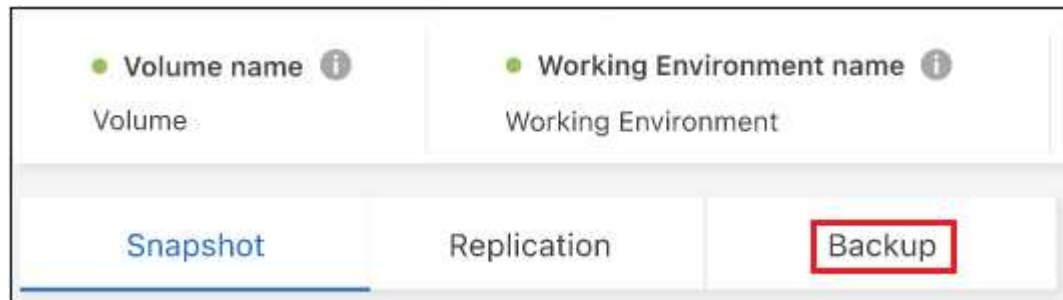
1. From the **Volumes** tab, click ... for the source volume and select **View volume details**.




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**.



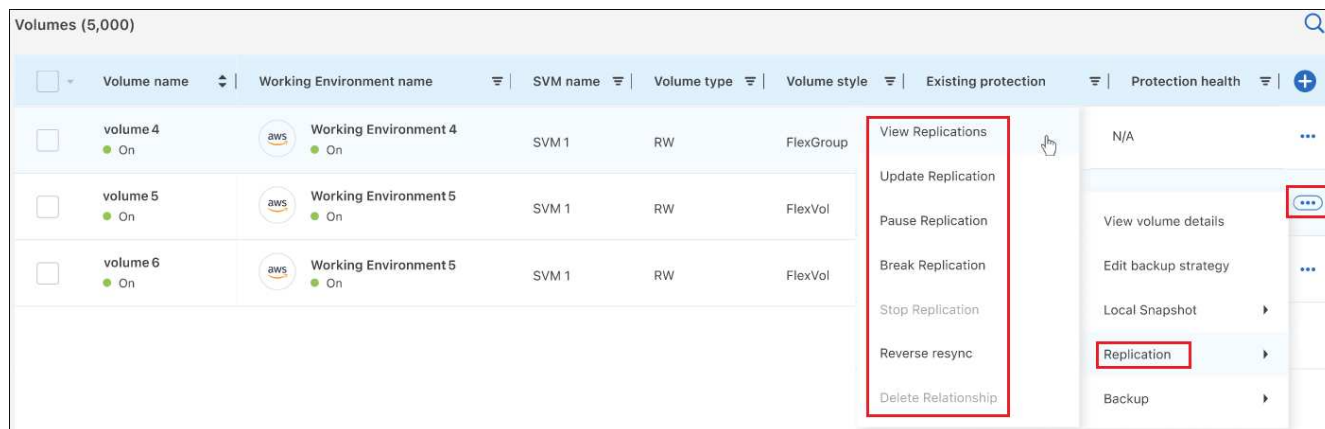
The Ransomware Protection column will show that the scan is In Progress.

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.



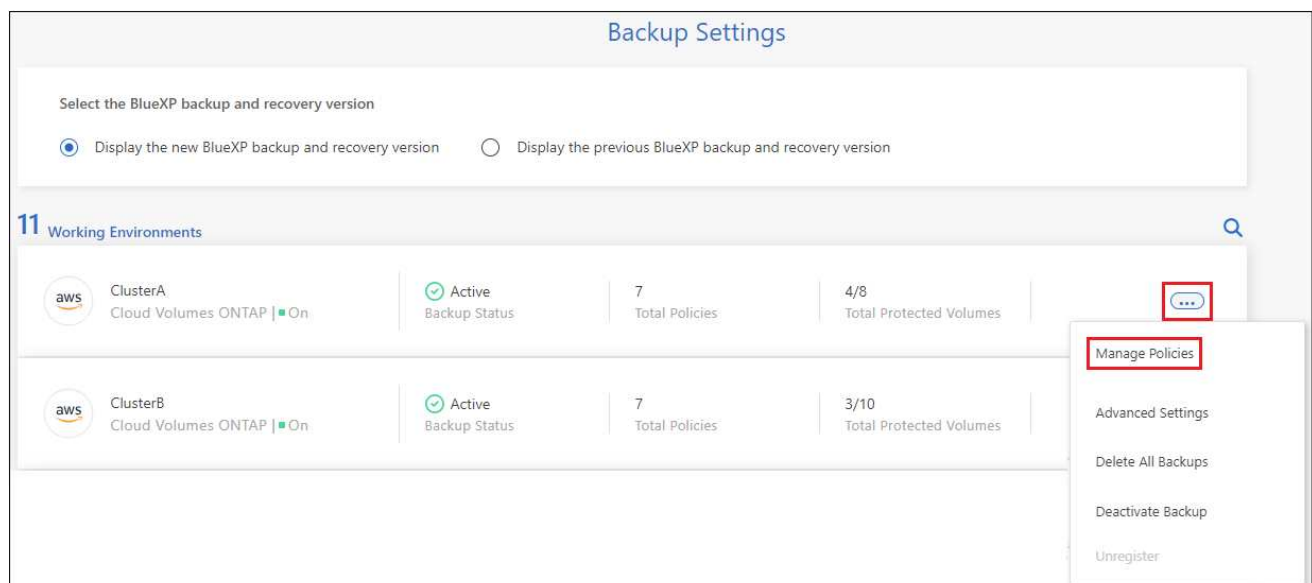
- If you enabled *DataLock and Ransomware Protection* in the initial policy when activating BlueXP backup and recovery for this cluster, any policies that you edit must be configured with the same DataLock setting (Governance or Compliance). And if you did not enable *DataLock and Ransomware Protection* 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**.



3. 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

4. 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.

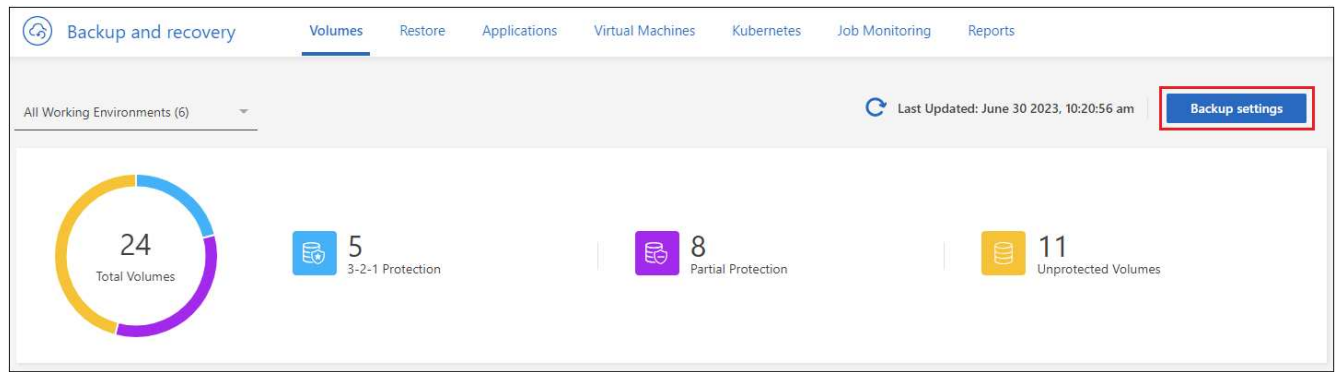
If you want to apply a new backup policy to certain volumes in a working environment, you first need to add the backup policy to the working environment. Then you can [apply the policy to volumes in that working environment](#).




- 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

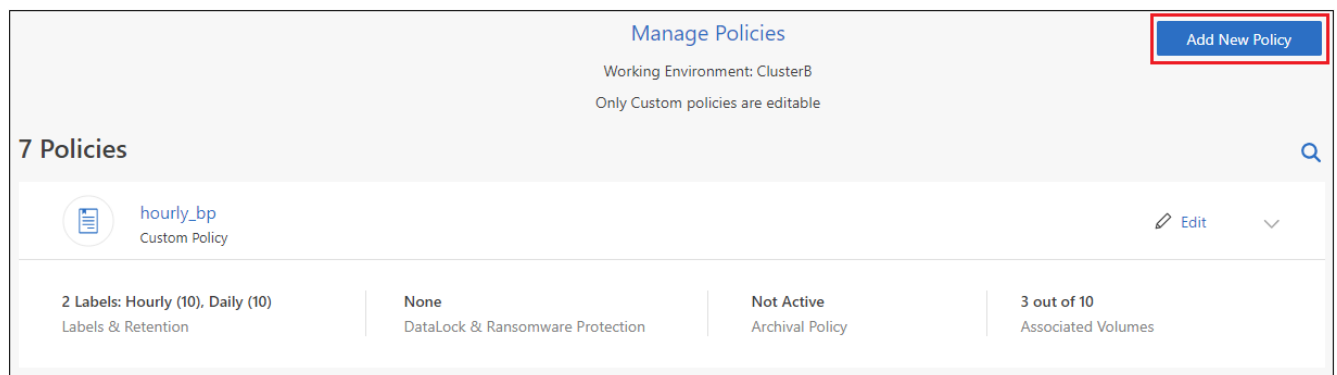
1. From the **Volumes** tab, select **Backup Settings**.



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.

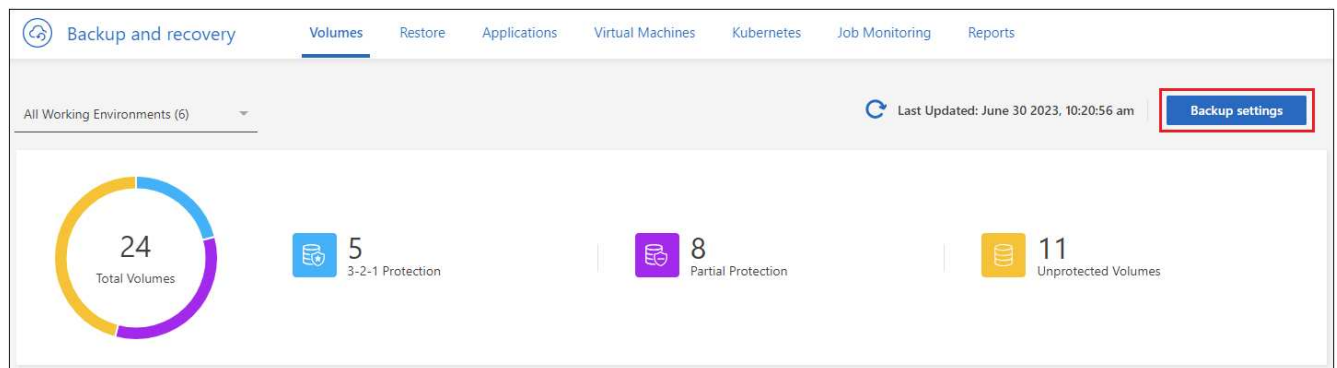
Delete all backup files for a working environment

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](#).

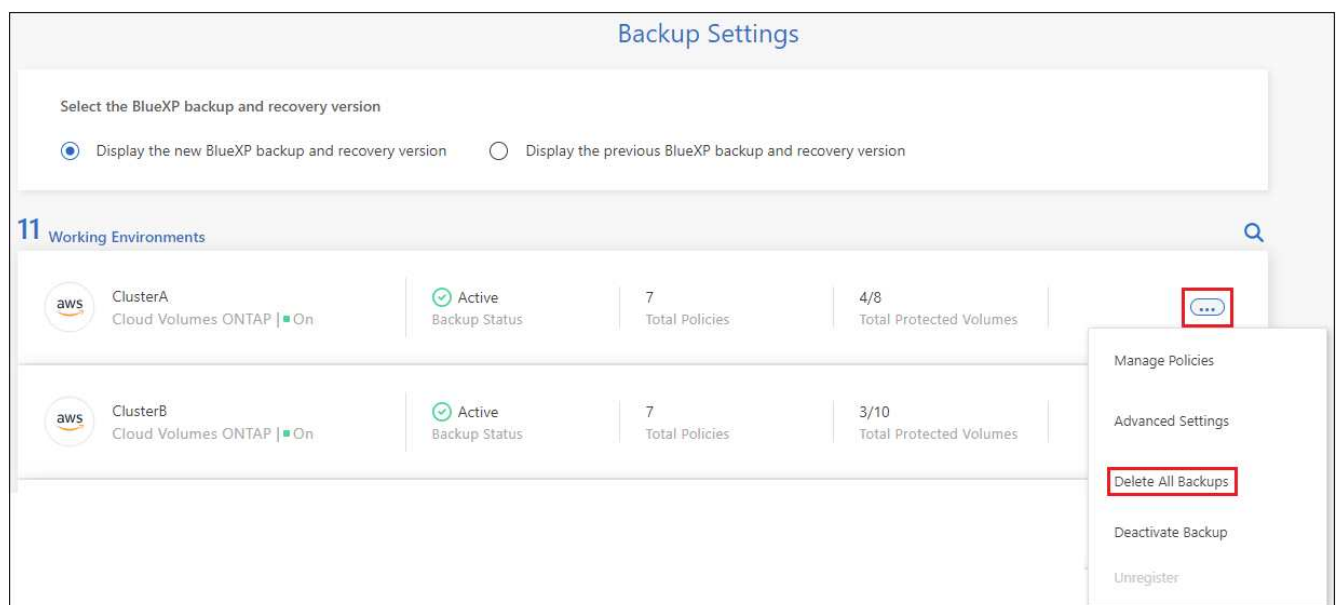
Note that this action does not affect Snapshot copies or replicated volumes - these types of backup files are not deleted.

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**.

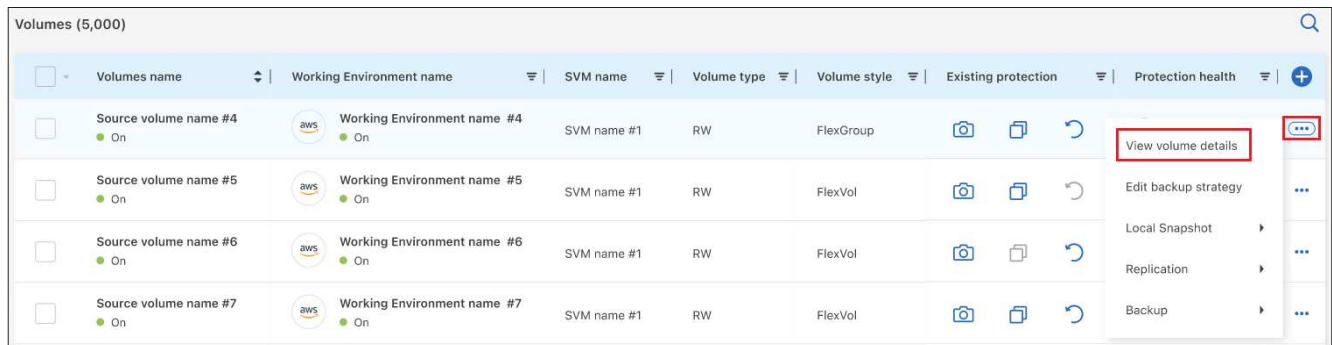
Delete a single backup file for a volume

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).

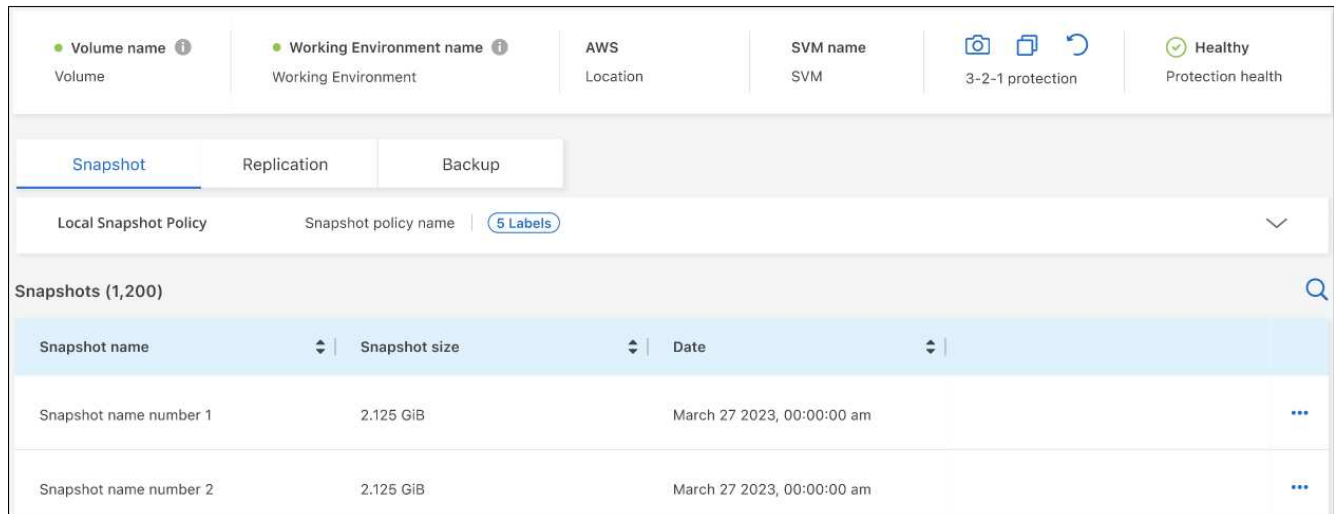
Steps

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
<input type="checkbox"/>	Source volume name #4 On	Working Environment name #4 On	SVM name #1	RW	FlexGroup		<div>View volume details Edit backup strategy Local Snapshot Replication Backup</div>
<input type="checkbox"/>	Source volume name #5 On	Working Environment name #5 On	SVM name #1	RW	FlexVol		
<input type="checkbox"/>	Source volume name #6 On	Working Environment name #6 On	SVM name #1	RW	FlexVol		
<input type="checkbox"/>	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	SVM name	3-2-1 protection	Healthy
Volume	Working Environment	Location	SVM	3-2-1 protection	Protection health

Snapshot

Replication

Backup

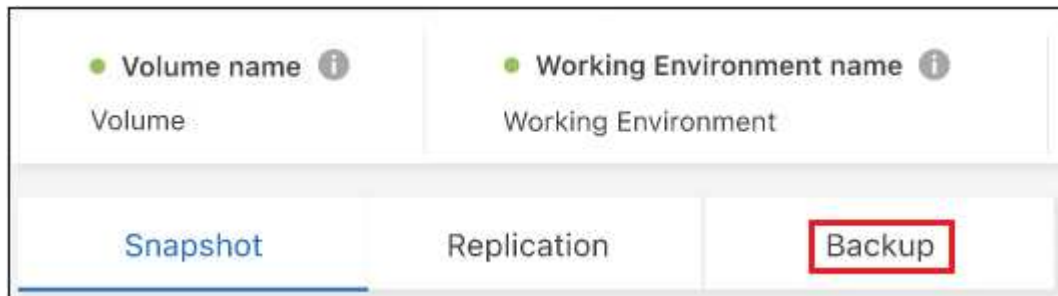
Local Snapshot Policy

Snapshot policy name

5 Labels

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, [delete all backups for the volume](#) if you want to remove all backup files.

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

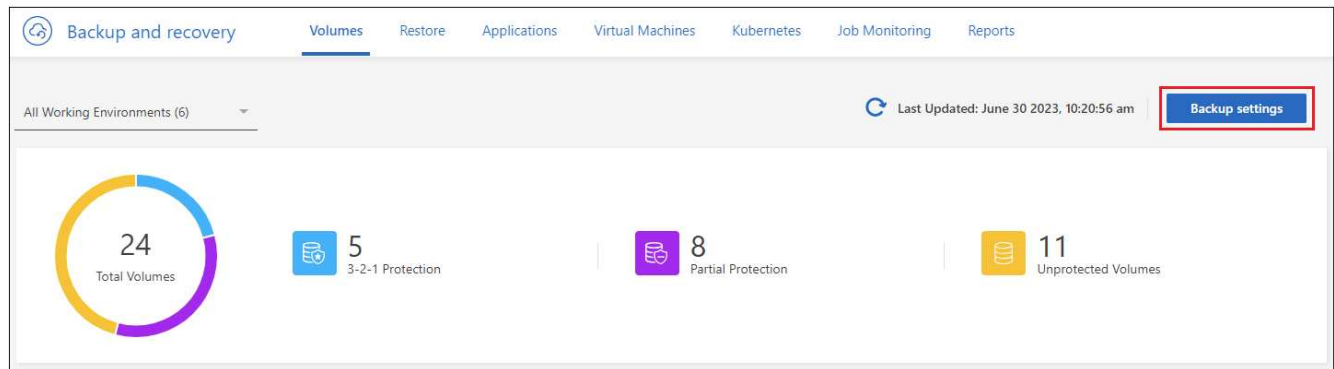
Deactivate BlueXP backup and recovery for a working environment

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.

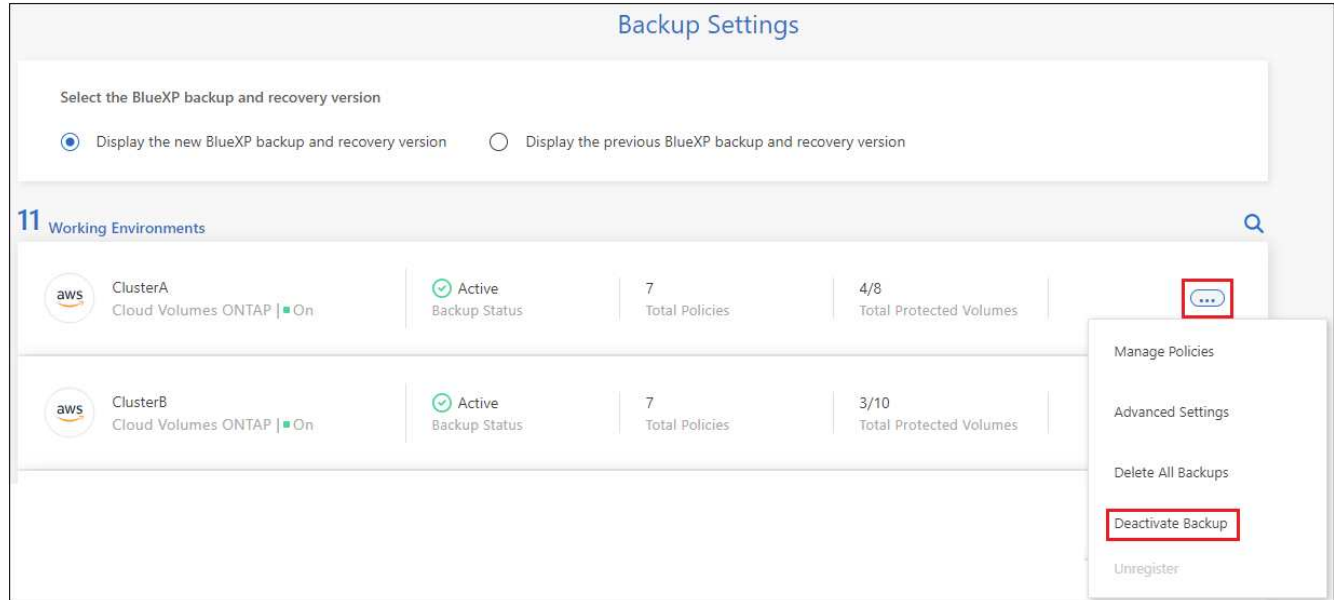
Note that you'll continue to be charged by your cloud provider for object storage costs for the capacity that your backups use unless you [delete the backups](#).

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

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 restoring from the cloud, additional bucket and public cloud provider resources required per account
No additional cloud provider costs required	When restoring from the cloud, additional costs are required when scanning your backups and volumes for search results

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	No
Replicated volume	Yes	No	No	Yes	Yes	No
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.

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 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

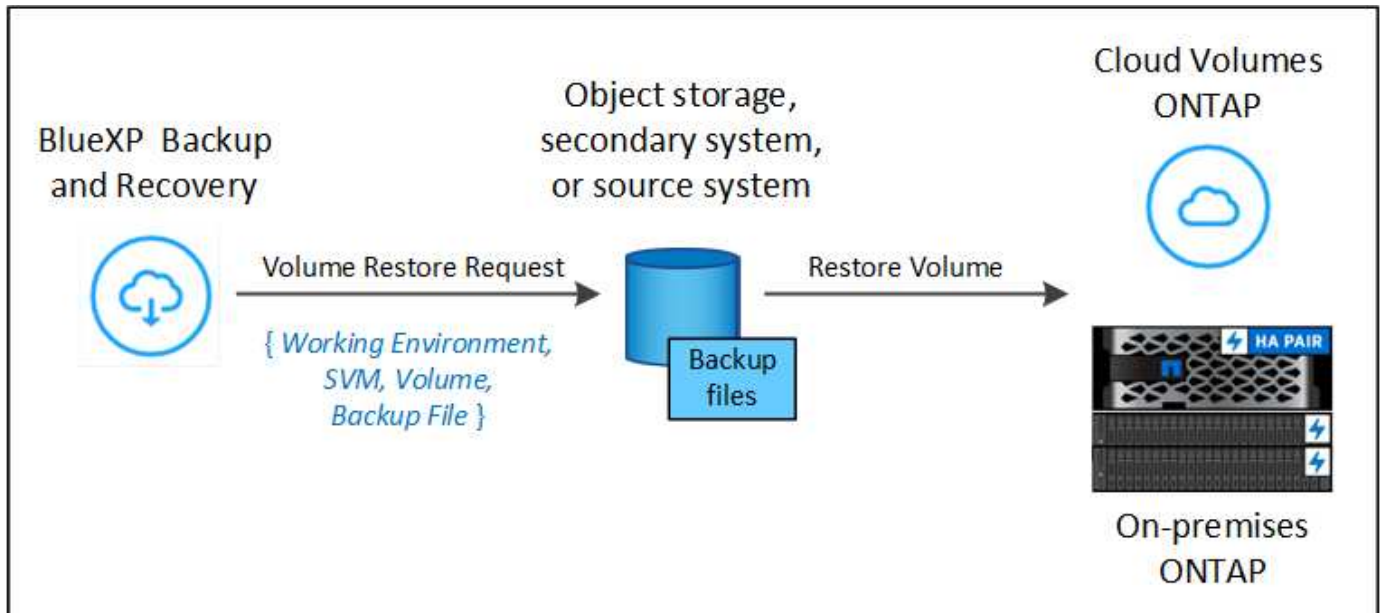
When you restore a volume from a backup file, BlueXP backup and recovery creates a *new* volume using the data from the backup. When using a backup from object storage, you can restore the data to a volume in the original working environment, to a different working environment that's located in the same 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, 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.



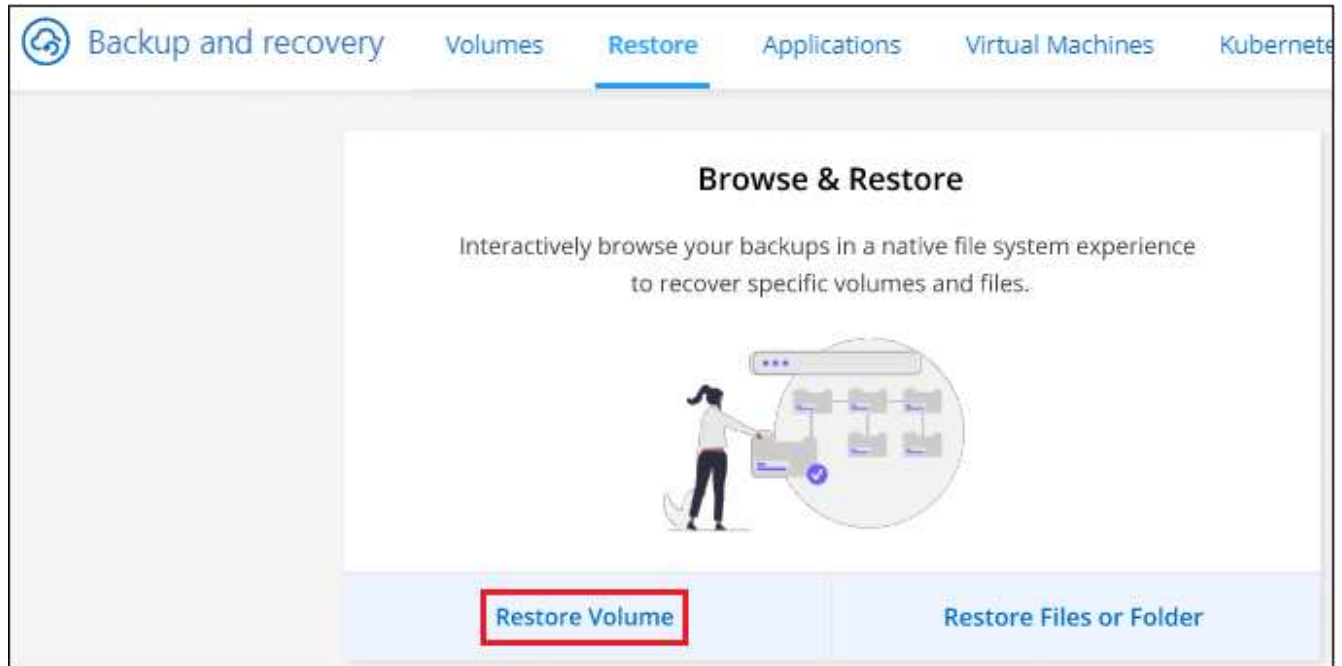
As you can see, you'll need to know the source working environment name, storage VM, volume name, and backup file date to perform a volume restore.

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.

When restoring a backup from object storage to a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater, 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 **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 only need to restore 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.

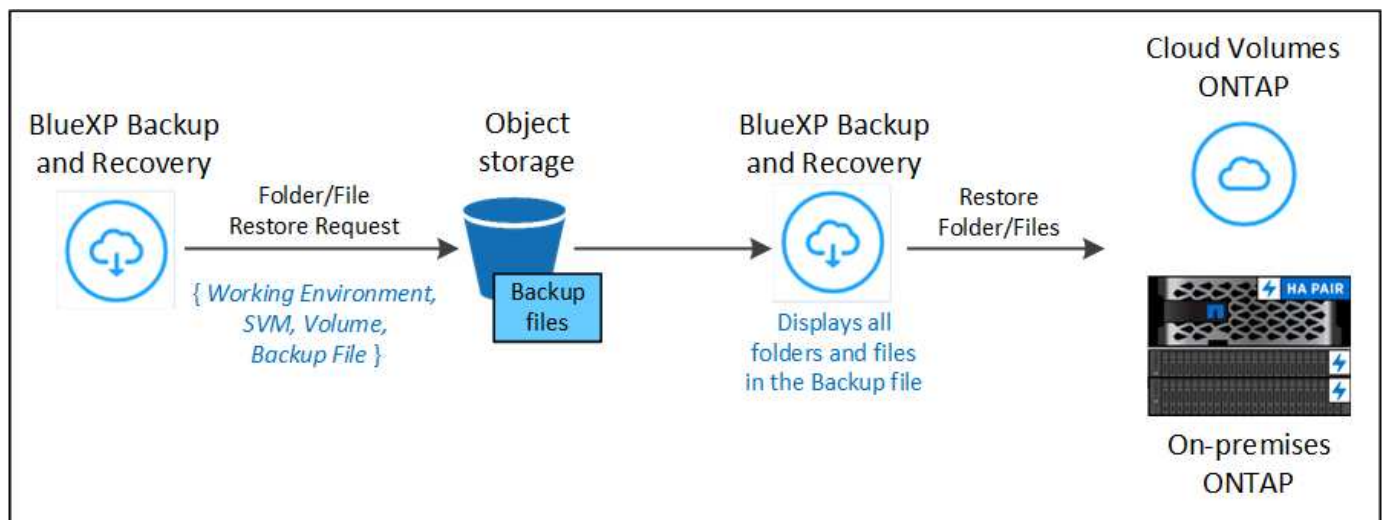
Prerequisites

- 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.

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

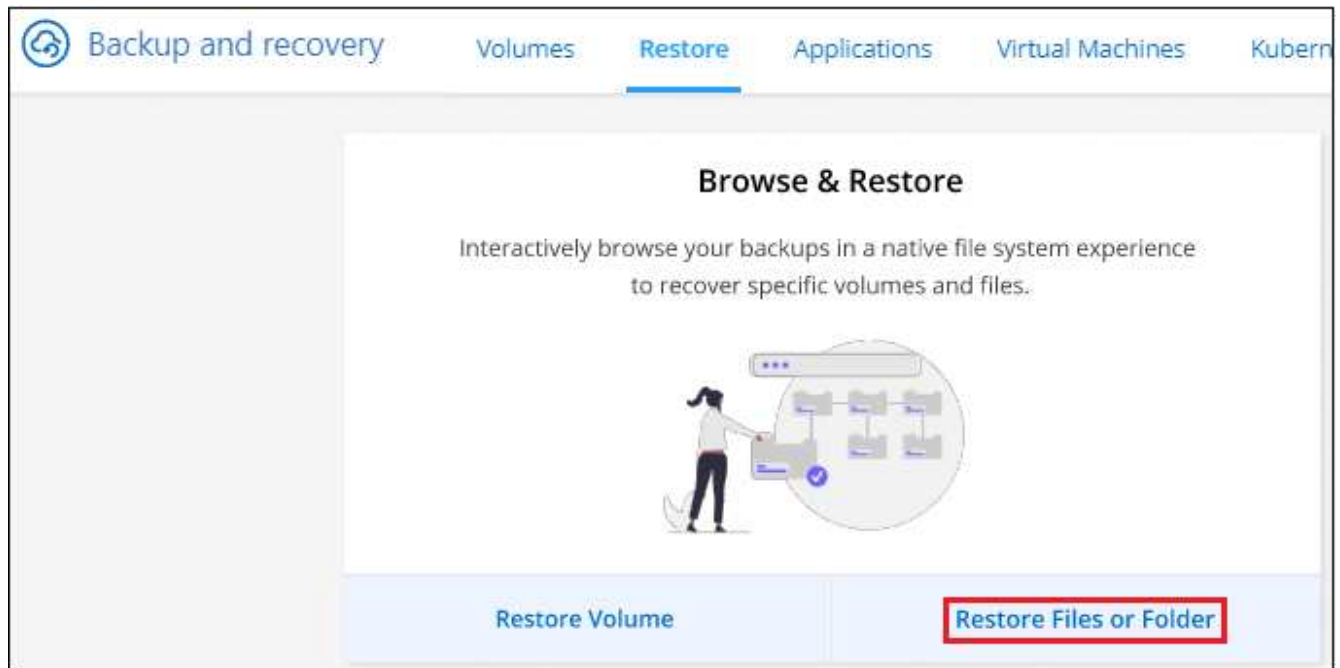
Follow these steps to restore folders or files to a volume from an ONTAP volume backup. You should know the name of the volume and the date of the backup file that you want to use to restore the folder or file(s). This functionality uses Live Browsing so that you can view the list of directories and 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.

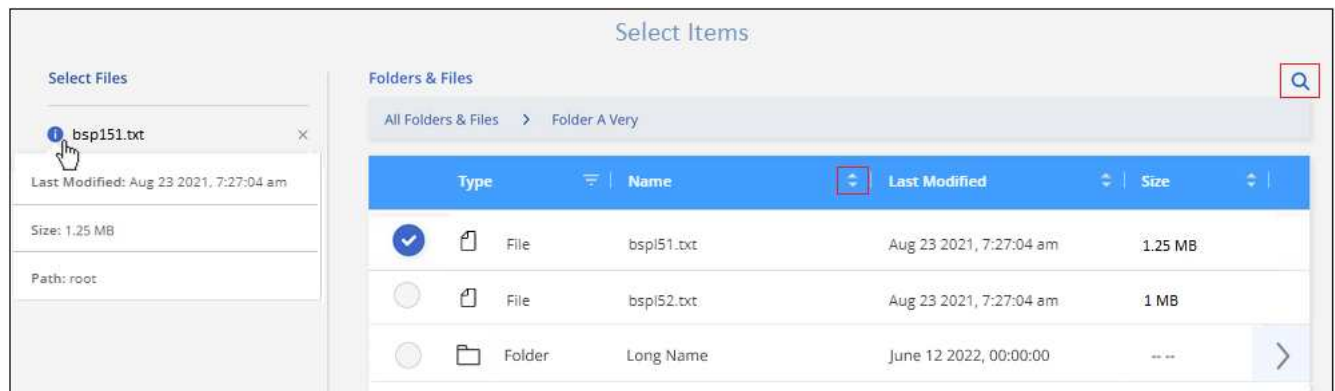
If you are restoring folders or files from a backup file that resides in an archival storage tier, then you can select the Restore Priority.

[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.)

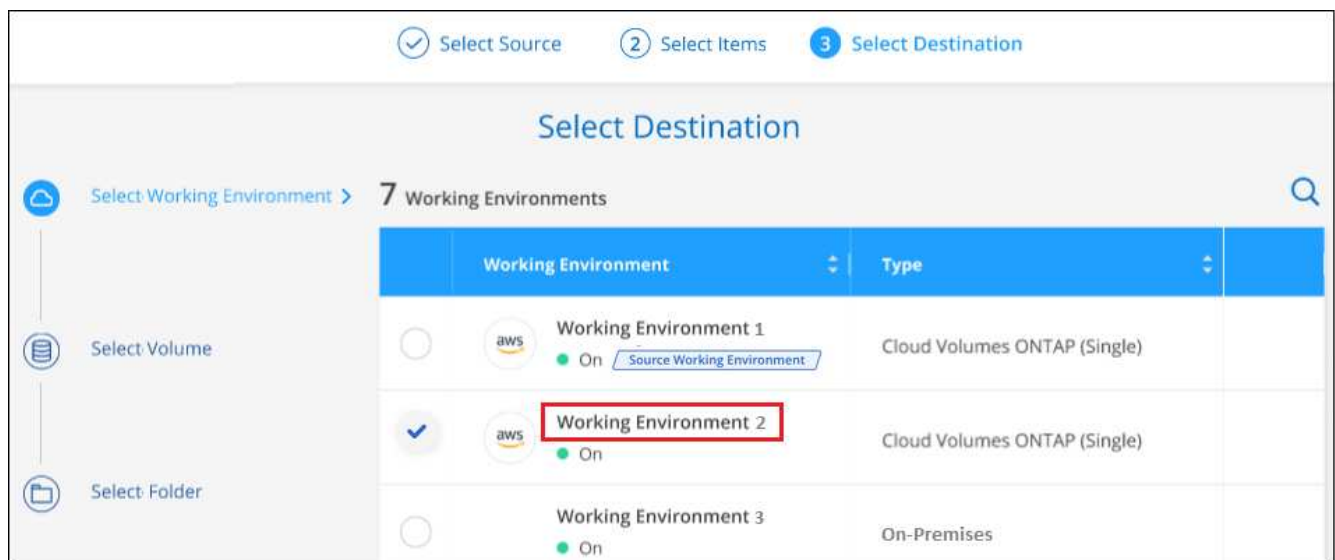


6. In the *Select Items* page, select the folder or file(s) that you want to restore and click **Continue**. To assist you in finding the item:

- 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.

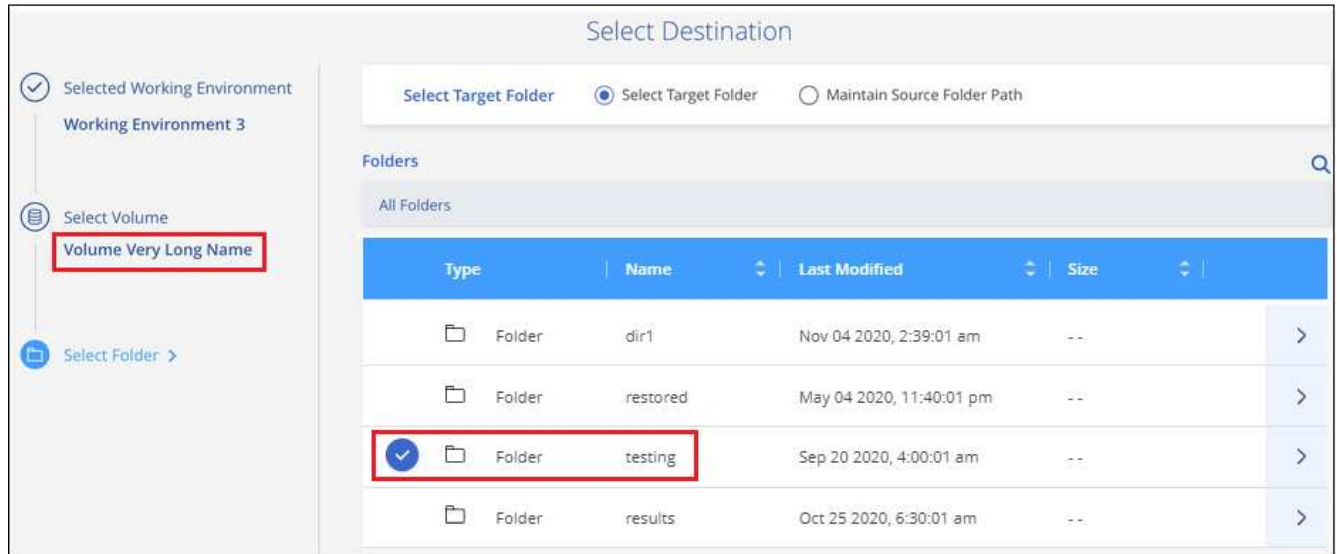


If you select an on-premises cluster and you haven't already configured the cluster connection to the object storage, you are prompted for additional information:


- 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.

Restoring 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.

When restoring a cloud backup to a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater, 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.

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.
- 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. Quick restore for SnapLock volumes is supported only if the source system from which the cloud backup was created was running ONTAP 9.11.0 or greater. There is no minimum requirement for FlexVol volumes.
- 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

mysearch 1a

All Resources 1b

Filter by: 1c

Working Environment File Type Size Creation Time Backup Location


Working Environment 2 Working Environment 4 JPEG August 10, 2021 - September 10, 2021 Cloud Primary (Local)

2 Search

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 also identify local volume Snapshot copies that contain the file in your search. You can choose to restore from the cloud backup file or from the Snapshot copy.

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.

When restoring a backup from object storage to a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater, you'll have the option to perform a "quick restore" operation.

- 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.

Manage cluster-level backup settings

You can change many cluster-level backup settings that you set when activating BlueXP backup and recovery for each ONTAP system. 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 more.



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.

Cluster-level backup settings are available in the *Advanced Settings* page.

The full set of backup settings that you can change includes:

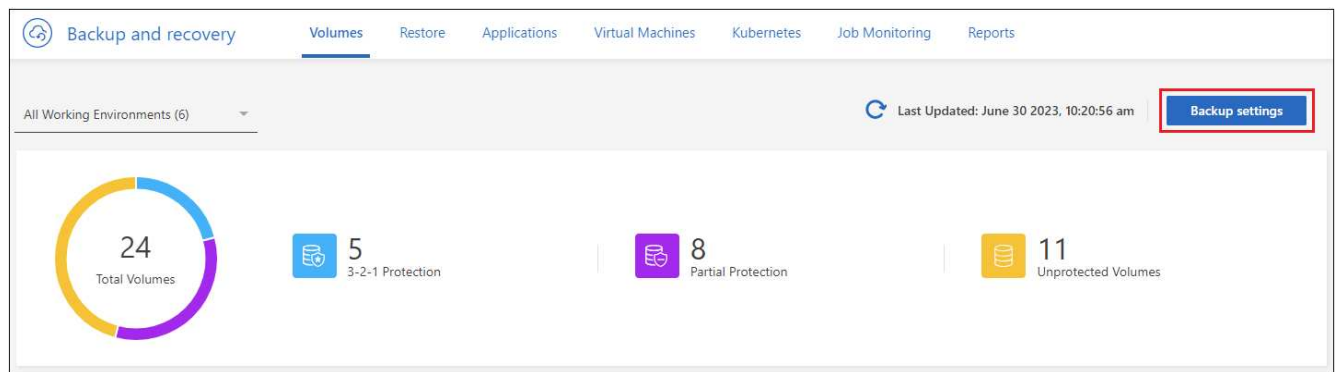
- Changing the network bandwidth allocated to upload backups to object storage
- Changing whether historical Snapshot copies are included in your initial baseline backup files for future volumes
- Changing whether "yearly" snapshots are removed from the source system

View cluster-level backup settings

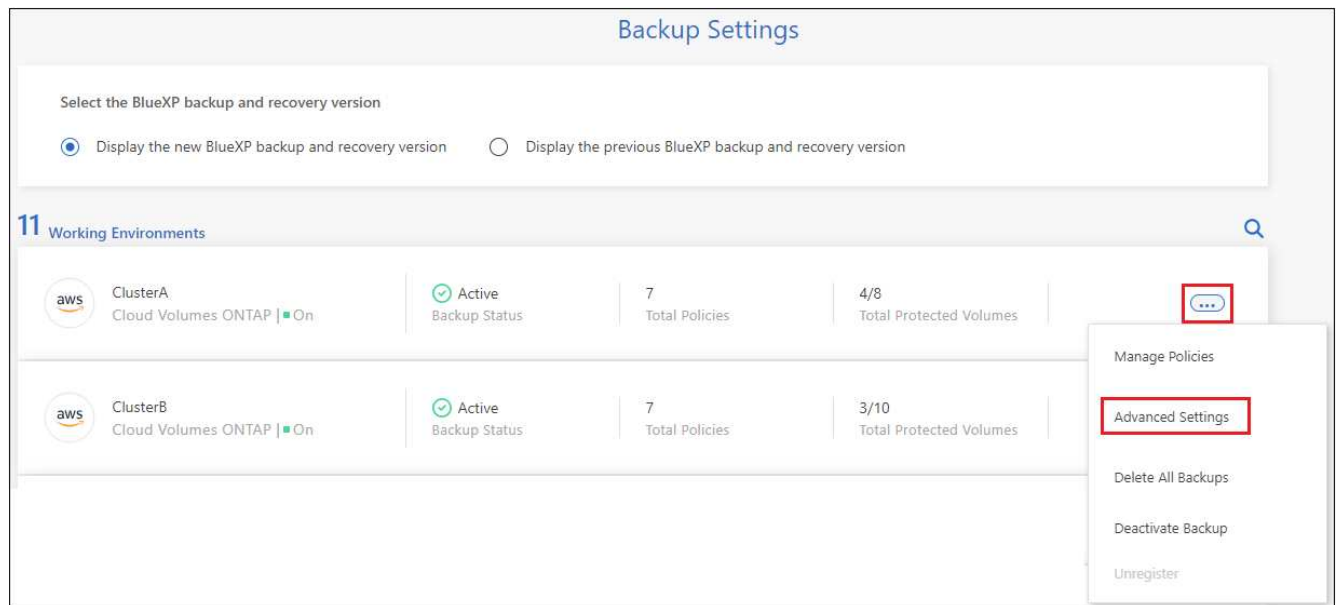
You can view the cluster-level backup settings for each working environment.

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.

Advanced Settings		
Working Environment: ClusterA		
Max Transfer Rate	Unlimited	▼
Export existing Snapshot copies	Disabled	▼
Yearly Snapshot Deletion	Enabled	▼

If you need to make any changes, just 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. You can choose a value between 1 and 1,000 Mbps as the maximum transfer rate.

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 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

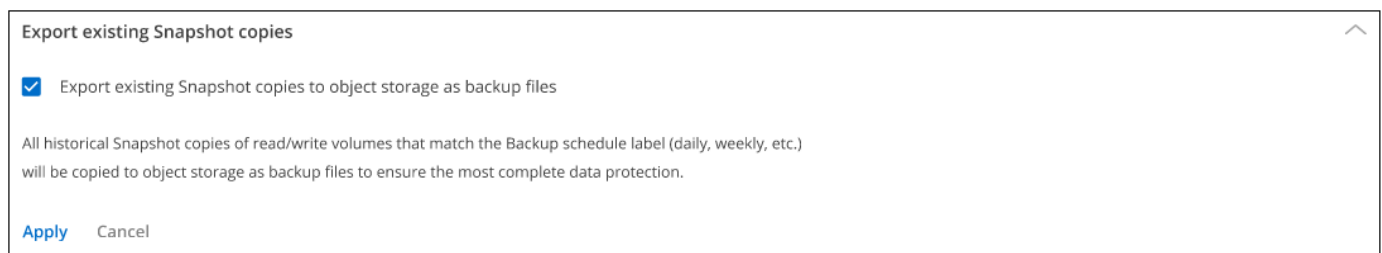
Select the **Limited** radio button and enter the maximum bandwidth that can be used, or select **Unlimited** to indicate that there is no limit.

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.

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

Just select whether you want existing Snapshot copies to be exported, and click **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.

Yearly Snapshot Deletion

Enabled

☒ Enabled

Yearly Snapshot copies are deleted from the source system after being transferred to object storage as backups.

☐ Disabled

Yearly Snapshot copies are retained on the source system. Note that these snapshots can be large.

Apply

Cancel

Select **Disabled** and click **Apply** if you want to retain the yearly snapshots on the source system.

Back up and restore on-premises applications data

Protect your on-premises applications data

BlueXP backup and recovery for applications provides data protection capabilities for application consistent Snapshots from on-premises ONTAP primary to cloud provider.

You can back up Oracle, Microsoft SQL, SAP HANA, MongoDB, MySQL, and PostgreSQL applications data from on-premises ONTAP systems to Amazon Web Services, Microsoft Azure, Google Cloud Platform, and StorageGRID.

For more information about BlueXP backup and recovery for applications, refer to:


- [Application aware backup with BlueXP backup and recovery and SnapCenter](#)
- [BlueXP backup and recovery for applications podcast](#)

Requirements

Read the following requirements to make sure that you have a supported configuration before you start backing up application data to cloud provider.

- ONTAP 9.8 or later
- BlueXP
- SnapCenter Server 4.6 or later
 - You should be using SnapCenter Server 4.7 or later if you want to use the following features:
 - Protect backups from on-premises secondary storage
 - Protect SAP HANA applications
 - Protect Oracle and SQL applications that are on VMware environment
 - Storage export of a backup
 - Deactivate backups
 - Unregister SnapCenter Server
 - You should be using SnapCenter Server 4.9 or later if you want to use the following features:
 - Mount Oracle database backups
 - Restore to the alternate storage
 - You should be using SnapCenter Server 4.9P1 if you want to protect MongoDB, MySQL, and PostgreSQL applications
- At least one backup per application should be available in SnapCenter Server
- At least one daily, weekly, or monthly policy in SnapCenter with no label or same label as that of the policy in BlueXP

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



Register SnapCenter Server

Only a user with SnapCenterAdmin role can register the host on which SnapCenter Server 4.6 or later is running. You can register multiple SnapCenter Server hosts in BlueXP.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. From the **Settings** drop-down, click **SnapCenter Servers**.
3. Click **Register SnapCenter Server**.
4. Specify the following details:
 - a. In the SnapCenter Server field, specify the FQDN or IP address of the SnapCenter Server host.
 - b. In the Port field, specify the port number on which the SnapCenter Server host is running.

You should ensure that the port is open for communication to happen between SnapCenter Server and BlueXP.

- c. In the Tags field, specify a site name, city name, or any custom name with which you want to tag the SnapCenter Server.

The tags are comma separated.

- d. In the Username and Password field, specify the credentials of the user with SnapCenterAdmin role.
5. Select the Connector from the **Connector** drop-down.
6. Click **Register**.

After you finish

Click **Backup & Restore > Applications** to view all the applications that are protected using the registered SnapCenter Server host. By default, the applications are automatically discovered every day midnight.

The supported applications and their configurations are:

- Oracle database:
 - Full backups (data + log) created with at least one daily, weekly, or monthly schedules
 - SAN, NFS, VMDK-SAN, VMDK-NFS, and RDM
- Microsoft SQL Server database:
 - Standalone, failover cluster instances, and availability groups
 - Full backups created with at least one daily, weekly, or monthly schedules
 - SAN, VMDK-SAN, VMDK-NFS, and RDM
- SAP HANA database:
 - Single Container 1.x
 - Multiple Database Container 2.x
 - HANA System Replication (HSR)

You should have at least one backup on both primary and secondary sites. You can decide to do a proactive failure or a deferred failover to the secondary.

- Non-data Volumes (NDV) resources such as HANA binaries, HANA archive log volume, HANA shared volume, and so on
- MongoDB
- MySQL
- PostgreSQL

The following databases are not displayed:

- Databases that have no backups
- Databases that have only on-demand or hourly policy
- Oracle databases residing on NVMe

Create a policy to back up applications

You should create a policy to back up the application data to cloud.

Before you begin

- If you want to move backups from object store to archival storage, ensure that you are using the required ONTAP version.
 - If you are using Amazon Web Services, you should be using ONTAP 9.10.1 or later

- If you are using Microsoft Azure, you should be using ONTAP 9.10.1 or later
- If you are using Google Cloud, you should be using ONTAP 9.12.1 or later
- If you are using StorageGrid, you should be using ONTAP 9.12.1 or later
- You should configure the archive access tier for each cloud provider.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. From the Settings drop-down, click **Policies > Create Policy**.
3. In the Policy Details section, specify the policy name.
4. In the Retention section, select one of the retention type and specify the number of backups to retain.
5. Select Primary or Secondary as the backup storage source.
6. (Optional) If you want to move backups from object store to archival storage after a certain number of days for cost optimization, select the **Tier Backups to Archival** checkbox.
7. Click **Create**.



You cannot edit or delete a policy, which is associated with an application.

Back up on-premises applications data to Amazon Web Services

Complete a few steps to back up the applications data from ONTAP to Amazon Web Services.

BlueXP supports data lock and ransomware protection. If ONTAP cluster is running on ONTAP 9.11.1 or later and if you have not configured archival storage, you can protect the backups from overwriting, deletion, and ransomware threats.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. Click **...** corresponding to the application and click **Activate Backup**.
3. In the Assign Policy page, select the policy and click **Next**.
4. Add the working environment.

Configure the cluster management LIF that you want BlueXP to discover. After adding the working environment for one of the applications, it can be reused for all the other applications residing on the same ONTAP cluster.

- a. Select the SVM and click **Add Working Environment**.
- b. In the Add Working Environment wizard:
 - i. Specify the IP address of the cluster management LIF.
 - ii. Specify the credentials of the ONTAP cluster user.

BlueXP backup and recovery for applications only support cluster admin.

c. Click **Add Working Environment**.

5. Select **Amazon Web Services** as the cloud provider.

a. Specify the AWS account.

b. In the AWS Access Key field, specify the key.

c. In the AWS Secret Key field, specify the password.

d. Select the region where you want to create the backups.

e. Specify the IP space.

f. Select the archival tier if you have configured archival storage in the policy.

It is recommended to set the archival tier because this is an one time activity and you will not be allowed to set it up later.

6. Configure data lock and ransomware protection.

7. Review the details and click **Activate Backup**.

Back up on-premises applications data to Microsoft Azure

Complete a few steps to back up the applications data from ONTAP to Microsoft Azure.

BlueXP supports data lock and ransomware protection. If ONTAP cluster is running on ONTAP 9.12.1 or later and if you have not configured archival storage, you can protect the backups from overwriting, deletion, and ransomware threats.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.

2. Click  corresponding to the application and click **Activate Backup**.

3. In the Assign Policy page, select the policy and click **Next**.

4. Add the working environment.

Configure the cluster management LIF that you want BlueXP to discover. After adding the working environment for one of the applications, it can be reused for all the other applications residing on the same ONTAP cluster.

a. Select the SVM and click **Add Working Environment**.

b. In the Add Working Environment wizard:

i. Specify the IP address of the cluster management LIF.

ii. Specify the credentials of the ONTAP cluster user.

BlueXP backup and recovery for applications only support cluster admin.

c. Click **Add Working Environment**.

5. Select **Microsoft Azure** as the cloud provider.

a. Specify the Azure subscription ID.

b. Select the region where you want to create the backups.

c. Either create a new resource group or use an existing resource group.

- d. Specify the IP space.
- e. Select the archival tier if you have configured archival storage in the policy.


It is recommended to set the archival tier because this is an one time activity and you will not be allowed to set it up later.

- 6. Configure data lock and ransomware protection.
- 7. Review the details and click **Activate Backup**.

Back up on-premises applications data to Google Cloud Platform

Complete a few steps to back up the applications data from ONTAP to Google Cloud Platform.

Steps

- 1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
- 2. Click  corresponding to the application and click **Activate Backup**.
- 3. In the Assign Policy page, select the policy and click **Next**.
- 4. Add the working environment.

Configure the cluster management LIF that you want BlueXP to discover. After adding the working environment for one of the applications, it can be reused for all the other applications residing on the same ONTAP cluster.

- a. Select the SVM and click **Add Working Environment**.
- b. In the Add Working Environment wizard:
 - i. Specify the IP address of the cluster management LIF.
 - ii. Specify the credentials of the ONTAP cluster user.

BlueXP backup and recovery for applications only support cluster admin.

- c. Click **Add Working Environment**.
- 5. Select **Google Cloud Platform** as the cloud provider.
 - a. Select the Google Cloud Project where you want the Google Cloud Storage bucket to be created for backups.
 - b. In the Google Cloud Access Key field, specify the key.
 - c. In the Google Cloud Secret Key field, specify the password.
 - d. Select the region where you want to create the backups.
 - e. Specify the IP space.
 - f. Select the archival tier.

It is recommended to set the archival tier because this is an one time activity and you will not be allowed to set it up later.

- 6. Review the details and click **Activate Backup**.

Back up on-premises applications data to StorageGRID

Complete a few steps to back up the applications data from ONTAP to StorageGRID.

BlueXP supports data lock and ransomware protection. If ONTAP cluster is running on ONTAP 9.11.1 or later, StorageGRID systems are 11.6.0.3 or later, and if you have not configured archival storage, you can protect the backups from overwriting, deletion, and ransomware threats.

Before you begin

When backing up data to StorageGRID, a Connector must be available on your premises. You will either need to install a new Connector or make sure that the currently selected Connector resides on-prem. The Connector can be installed in a site with or without internet access.

For information, see [Create Connectors for StorageGRID](#).

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. Click **...** corresponding to the application and click **Activate Backup**.
3. In the Assign Policy page, select the policy and click **Next**.
4. Add the working environment.

Configure the cluster management LIF that you want BlueXP to discover. After adding the working environment for one of the applications, it can be reused for all the other applications residing on the same ONTAP cluster.

- a. Select the SVM and click **Add Working Environment**.
- b. In the Add Working Environment wizard:
 - i. Specify the IP address of the cluster Management LIF.
 - ii. Specify the credentials of the ONTAP cluster user.

BlueXP backup and recovery for applications only support cluster admin.

- c. Click **Add Working Environment**.

5. Select **StorageGRID**.

- a. Specify the FQDN of the StorageGRID Server and the port on which the StorageGRID server is running.

Enter the details in the format FQDN:PORT.

- b. In the Access Key field, specify the key.
- c. In the Secret Key field, specify the password.
- d. Specify the IP space.
- e. Specify the archival tier if you have configured archival storage in the policy.

If you select...	Perform the following...
AWS	<ol style="list-style-type: none"> 1. Either select the StorageGrid from the drop-down or add the StorageGrid cluster. 2. Specify the AWS account. 3. In the AWS Access Key field, specify the key. 4. In the AWS Secret Key field, specify the password. 5. Select the region where you want to create the backups. 6. Click Save.
Azure	<ol style="list-style-type: none"> 1. Either select the StorageGrid cluster from the drop-down or add the cluster. 2. Specify the Azure subscription ID. 3. Select the region where you want to create the backups. 4. Either create a new resource group or use an existing resource group. 5. Click Save.

It is recommended to set the archival tier because this is an one time activity and you will not be allowed to set it up later.

6. Configure data lock and ransomware protection.
7. Review the details and click **Activate Backup**.

Manage protection of applications

You can manage protection of applications by viewing policies, viewing backups, viewing the changes to database layout, policies, and resource group, and monitoring all the operations from the BlueXP UI.

View policies

You can view all the policies. For each of these policies, when you view the details all the associated applications are listed.

Steps

1. Click **Backup and recovery > Applications**.
2. From the **Settings** drop-down, click **Policies**.
3. Click **View Details** corresponding to policy whose details you want to view.

The associated applications are listed.



You cannot edit or delete a policy, which is associated with an application.

You can also view cloud extended SnapCenter policies, by running the `Get-SmResources` cmdlet in SnapCenter.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running `Get-Help` command name.

View backups on cloud

You can view the backups on cloud in the BlueXP UI.

Steps

1. Click **Backup and recovery > Applications**.
2. Click **...** corresponding to the application and click **View Details**.



The time taken for the backups to be listed depends on ONTAP's default replication schedule.

- For Oracle databases, both data and log backups, system change number (SCN) for each backup, end date for each backup are listed. You can select only the data backup and restore the database to original location. You can mount the data backup and log backup to alternate location.
- For Microsoft SQL Server databases, only the full backups and the end date for each backup are listed. You can select the backup and restore the database to original or alternate location.
- For Microsoft SQL Server instance, backups of the databases under that instance is listed.
- For SAP HANA databases, only the data backups and the end date for each backup are listed. You can select the backup and perform storage export on a given host.
- For MongoDB, MySQL, and PostgreSQL, only the data backups and the end date for each backup are listed. You can select the backup and perform storage export on a given host.



The backups created before enabling cloud protection are not listed for restore.

You can also view these backups by running the `Get-SmBackup` cmdlet in SnapCenter.

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running `Get-Help` command name.

Deactivate backup

You can delete all the backups that are moved to object store from both SnapCenter and the object store.

Steps

1. Click **Backup and recovery > Applications**.
2. Click **...** corresponding to the application and click **Deactivate Backup**.

By default the check box is selected and it deletes all the backups that are moved to object store from both SnapCenter and the object store.

If you clear the checkbox, the backups are retained only in the object store but these backups cannot be used for application level restore. Later when you activate the backup for this application, the backups retained in object store are not listed for restore.

3. Click **Deactivate Backup**.

Database layout change

When volumes are added to the database, SnapCenter Server labels the snapshots on the new volumes automatically as per the policy and the schedule. These new volumes will not have the object store endpoint and you should refresh the volumes by executing the following steps:

Steps

1. Click **Backup and recovery > Applications**.
2. From the **Settings** drop-down, click **SnapCenter Servers**.
3. Click **...** corresponding to the SnapCenter Server hosting the application and click **Refresh**.

The new volumes are discovered.

4. Click **...** corresponding to the application and click **Refresh Protection** to enable cloud protection for the new volume.
 - If a storage volume is added to the application after configuring the cloud provider, SnapCenter Server labels the snapshots for new backups on which the application is residing.
 - You should manually delete the object store relationship if the removed volume is not used by any other applications.
 - If you update the application inventory, it will contain the current storage layout of the application.

Policy or resource group change

If there is a change to the SnapCenter policy or resource group, you should refresh the protection relationship.

Steps

1. Click **Backup and recovery > Applications**.
2. Click **...** corresponding to the application and click **Refresh Protection**.

Unregister SnapCenter Server

Steps

1. Click **Backup and recovery > Applications**.
2. From the **Settings** drop-down, click **SnapCenter Servers**.
3. Click **...** corresponding to the SnapCenter Server and click **Unregister**.

By default the check box is selected and it deletes all the backups that are moved to object store from both SnapCenter and the object store.

If you clear the checkbox, the backups are retained only in the object store but these backups cannot be used for application level restore. Later when you activate the backup for this application, the backups retained in object store are not listed for restore.

Monitor Jobs

Jobs are created for all the Cloud Backup operations. You can monitor all the jobs and all the sub tasks that are performed as part of each task.

Steps

1. Click **Backup and recovery > Job Monitoring**.

When you initiate an operation, a window appears stating that the job is initiated. You can click the link to monitor the job.

2. Click the primary task to view the sub tasks and status of each of these sub tasks.

Configure CA Certificates

You can configure CA signed certificate if you want to include additional security to your environment.

Configure SnapCenter CA signed certificate in BlueXP Connector

You should configure SnapCenter CA signed certificate in BlueXP Connector so that the Connector can verify the SnapCenter's certificate.

Before you begin

You should run the following command in the BlueXP Connector to get the `<base_mount_path>`:

```
sudo docker volume ls | grep snapcenter_volume | awk {'print $2'} | xargs sudo
docker volume inspect | grep Mountpoint
```

Steps

1. Log in to the Connector.
`cd <base_mount_path> mkdir -p server/certificate`
2. Copy the root CA and intermediate CA files to the `<base_mount_path>/server/certificate` directory.

The CA files should be in .pem format.

3. If you have CRL files, perform the following steps:
 - a. `cd <base_mount_path> mkdir -p server/crl`
 - b. Copy the CRL files to the `<base_mount_path>/server/crl` directory.
4. Connect to the `cloudmanager_snapcenter` and modify the `enableCACert` in `config.yml` to true.
`sudo docker exec -t cloudmanager_snapcenter sed -i 's/enableCACert: false/enableCACert: true/g' /opt/netapp/cloudmanager-snapcenter/config/config.yml`
5. Restart `cloudmanager_snapcenter` container.
`sudo docker restart cloudmanager_snapcenter`

Configure CA signed certificate for BlueXP Connector

If 2way SSL is enabled in SnapCenter, you should perform the following steps on the Connector to use the CA certificate as the client certificate when the Connector is connecting with the SnapCenter.

Before you begin

You should run the following command to get the `<base_mount_path>`:

```
sudo docker volume ls | grep snapcenter_volume | awk {'print $2'} | xargs sudo
docker volume inspect | grep Mountpoint
```

Steps

1. Log in to the Connector.

```
cd <base_mount_path> mkdir -p client/certificate
```

2. Copy the CA signed certificate and key file to the `<base_mount_path>/client/certificate` in the Connector.

The file name should be `certificate.pem` and `key.pem`. The `certificate.pem` should have the entire chain of the certificates like intermediate CA and root CA.

3. Create the PKCS12 format of the certificate with the name `certificate.p12` and keep at `<base_mount_path>/client/certificate`.

Example: `openssl pkcs12 -inkey key.pem -in certificate.pem -export -out certificate.p12`

4. Connect to the `cloudmanager_snapcenter` and modify the `sendCACert` in `config.yml` to true.

```
sudo docker exec -t cloudmanager_snapcenter sed -i 's/sendCACert:
false/sendCACert: true/g' /opt/netapp/cloudmanager-
snapcenter/config/config.yml
```

5. Restart `cloudmanager_snapcenter` container.

```
sudo docker restart cloudmanager_snapcenter
```

6. Perform the following steps on the SnapCenter to validate the certificate sent by the Connector.

- a. Login to the SnapCenter Sever host.
- b. Click **Start > Start Search**.
- c. Type `mmc` and press **Enter**.
- d. Click **Yes**.
- e. In File menu, click **Add/Remove Snap-in**.
- f. Click **Certificates > Add > Computer account > Next**.
- g. Click **Local computer > Finish**.
- h. If you have no more snap-ins to add to the console, click **OK**.
- i. In the console tree, double-click **Certificates**.
- j. Right-click the **Trusted Root Certification Authorities store**.
- k. Click **Import** to import the certificates and follow the steps in the **Certificate Import Wizard**.

Restore on-premises applications data

Restore Oracle database

You can restore Oracle database either to the original location or to the alternate location. For a RAC database, the data is restored to the on-premises node where the backup was created.

Only full database with control file restore is supported. If the archive logs are not present in the AFS, you should specify the location that contains the archive logs required for recovery.



Single File Restore (SFR) is not supported.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. In the **Filter By** field, select the filter **Type** and from the drop-down select **Oracle**.
3. Click **View Details** corresponding to the database that you want to restore and click **Restore**.
4. In the Restore options page, specify the location where you want to restore the database files.

If you...	Do this...
Want to restore to the original location	<ol style="list-style-type: none"> 1. Select Restore to original location. 2. If the snapshot is in archival storage, select the priority to restore your data from the archival storage. 3. Click Next. 4. Select Database State if you want to change the state of the database to the state required to perform restore and recovery operations. <p>The various states of a database from higher to lower are open, mounted, started, and shutdown.</p> <ul style="list-style-type: none"> ◦ If the database is in a higher state but the state must be changed to a lower state to perform a restore operation, you must select this check box. ◦ If the database is in a lower state but the state must be changed to a higher state to perform the restore operation, the database state is changed automatically even if you do not select the check box. ◦ If a database is in the open state, and for restore the database needs to be in the mounted state, then the database state is changed only if you select this check box. 5. Specify the recovery scope. <ul style="list-style-type: none"> ◦ Select All Logs if you want to recover to the last transaction. ◦ Select Until SCN (System Change Number) if you want to recover to a specific SCN. ◦ Select Date and Time if you want to recover to a specific data and time. <p>You must specify the date and time of the database host's time zone.</p> ◦ Select No recovery if you do not want to recover. 6. If the archive logs are not present in the active file system, you should specify the location that contains the archive logs required for recovery. 7. Select the check box if you want to open the database after recovery. <p>In a RAC setup, only the RAC instance that is used for recovery is opened after recovery.</p>

If you...	Do this...
<p>Want to temporarily restore to another storage and then copy the restored files to the original location</p>	<ol style="list-style-type: none"> 1. Select Restore to original location. 2. If the snapshot is in archival storage, select the priority to restore your data from the archival storage. 3. Select Change storage location. <p>If you select Change storage location, you can append a suffix to the destination volume. If you have not selected the checkbox, then by default _restore is appended to the destination volume.</p> 4. Click Next. 5. In the Storage mapping page, specify the alternate storage location details where the data restored from the object store will be stored temporarily. <p>If you select an on-premises ONTAP system and if you haven't configured the cluster connection to the object storage, you are prompted for additional information regarding the object store.</p> 6. Click Next. 7. Select Database State if you want to change the state of the database to the state required to perform restore and recovery operations. <p>The various states of a database from higher to lower are open, mounted, started, and shutdown.</p> <ul style="list-style-type: none"> ◦ If the database is in a higher state but the state must be changed to a lower state to perform a restore operation, you must select this check box. ◦ If the database is in a lower state but the state must be changed to a higher state to perform the restore operation, the database state is changed automatically even if you do not select the check box. ◦ If a database is in the open state, and for restore the database needs to be in the mounted state, then the database state is changed only if you select this check box. 8. Specify the recovery scope. <ul style="list-style-type: none"> ◦ Select All Logs if you want to recover to the last transaction. ◦ Select Until SCN (System Change Number) if you want to recover to a specific SCN.

If you...	Do this...
Want to restore to an alternate location	<ol style="list-style-type: none"> 1. Select Restore to alternate location. 2. If the snapshot is in archival storage, select the priority to restore your data from the archival storage. 3. If you want to restore to alternate storage, perform the following: <ol style="list-style-type: none"> a. Select Change storage location. If you select Change storage location, you can append a suffix to the destination volume. If you have not selected the checkbox, then by default _restore is appended to the destination volume. b. Click Next. c. In the Storage mapping page, specify the alternate storage location details where the data from the object store needs to be restored. 4. Click Next. 5. In the Destination host page, select the host on which the database will be mounted. <ol style="list-style-type: none"> a. (Optional) For NAS environment, specify the FQDN or IP address of the host to which the volumes restored from object store are to be exported. b. (Optional) For SAN environment, specify the initiators of the host to which LUNs of the volumes restored from object store are to be mapped. 6. Click Next.

5. Review the details and click **Restore**.

The **Restore to alternate location** option mounts the selected backup on the given host. You should manually bring up the database.

After mounting the backup, you cannot mount it again until it is unmounted. You can use the **Unmount** option from the UI to unmount the backup.

For information on how to bring up the Oracle database see, [Knowledge base article](#).

Restore SQL Server database

You can restore SQL Server database either to the original location or to the alternate location.





Single File Restore (SFR), Recovery of log backups, and reseed of availability groups are not supported.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. In the **Filter By** field, select the filter **Type** and from the drop-down select **SQL**.
3. Click **View Details** to view all the available backups.
4. Select the backup and click **Restore**.
5. In the Restore options page, specify the location where you want to restore the database files.

If you...	Do this...
Want to restore to the original location	<ol style="list-style-type: none">1. Select Restore to original location.2. If the snapshot is in archival storage, select the priority to restore your data from the archival storage.3. Click Next.
Want to temporarily restore to another storage and then copy the restored files to the original location	<ol style="list-style-type: none">1. Select Restore to original location.2. If the snapshot is in archival storage, select the priority to restore your data from the archival storage.3. Select Change storage location. If you select Change storage location, you can append a suffix to the destination volume. If you have not selected the checkbox, then by default _restore is appended to the destination volume.4. Click Next.5. In the Storage mapping page, specify the alternate storage location details where the data restored from the object store will be stored temporarily.6. Click Next.

If you...	Do this...
<p>Want to restore to an alternate location</p>	<ol style="list-style-type: none"> 1. Select Restore to alternate location. 2. If the snapshot is in archival storage, select the priority to restore your data from the archival storage. 3. Click Next. 4. In the Destination host page, select a host name, provide a database name (optional), select an instance, and specify the restore paths. <div data-bbox="922 596 976 653">  </div> <div data-bbox="1037 558 1448 688"> <p>The file extension provided in the alternate path must be same as the file extension of the original database file.</p> </div> <ol style="list-style-type: none"> 5. Click Next.
<p>Want to temporarily restore to another storage and then copy the restored files to the alternate location</p>	<ol style="list-style-type: none"> 1. Select Restore to alternate location. 2. If the snapshot is in archival storage, select the priority to restore your data from the archival storage. 3. Select Change storage location. <p>If you select Change storage location, you can append a suffix to the destination volume. If you have not selected the checkbox, then by default _restore is appended to the destination volume.</p> 4. Click Next. 5. In the Storage mapping page, specify the alternate storage location details where the data restored from the object store will be stored temporarily. 6. Click Next. 7. In the Destination host page, select a host name, provide a database name (optional), select an instance, and specify the restore paths. <div data-bbox="922 1692 976 1749">  </div> <div data-bbox="1037 1654 1448 1785"> <p>The file extension provided in the alternate path must be same as the file extension of the original database file.</p> </div> <ol style="list-style-type: none"> 8. Click Next.

6. In the **Pre-operations** select, select one of the following options:
 - Select **Overwrite the database with same name during restore** to restore the database with the same name.
 - Select **Retain SQL database replication settings** to restore the database and retain the existing replication settings.
7. In the **Post-operations** section, to specify the database state for restoring additional transactional logs, select one of the following options:
 - Select **Operational, but unavailable** if you are restoring all of the necessary backups now.

This is the default behavior, which leaves the database ready for use by rolling back the uncommitted transactions. You cannot restore additional transaction logs until you create a backup.
 - Select **Non-operational, but available** to leave the database non-operational without rolling back the uncommitted transactions.

Additional transaction logs can be restored. You cannot use the database until it is recovered.
 - Select **Read-only mode, and available** to leave the database in read-only mode.

This option undoes uncommitted transactions, but saves the undone actions in a standby file so that recovery effects can be reverted.

If the Undo directory option is enabled, more transaction logs are restored. If the restore operation for the transaction log is unsuccessful, the changes can be rolled back. The SQL Server documentation contains more information.
8. Click **Next**.
9. Review the details and click **Restore**.

Restore SAP HANA database

You can restore SAP HANA database to any host.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. In the **Filter By** field, select the filter **Type** and from the drop-down select **HANA**.
3. Click **View Details** corresponding to the database that you want to restore and click **Restore**.
4. In the Restore options page, specify one of the following:
 - a. For NAS environment, specify the FQDN or IP address of the host to which the volumes restored from object store are to be exported.
 - b. For SAN environment, specify the initiators of the host to which LUNs of the volumes restored from object store are to be mapped.
5. If the snapshot is in archival storage, select the priority to restore your data from the archival storage.
6. If there is not enough space on the source storage or the source storage is down, select **Change storage location**.

If you select **Change storage location**, you can append a suffix to the destination volume. If you have not selected the checkbox, then by default **_restore** is appended to the destination volume.

7. Click **Next**.
8. In the Storage mapping page, specify the alternate storage location details where the data restored from the object store will be stored.
9. Click **Next**.
10. Review the details and click **Restore**.

This operation does only the storage export of the selected backup on the given host. You should manually mount the filesystem and bring up the database. After utilizing the volume, the storage Administrator can delete the volume from the ONTAP cluster.

For information on how to bring up the SAP HANA database see, [TR-4667: Overview of SAP system copy workflow with SnapCenter](#) and [TR-4667: Overview of SAP system clone workflow with SnapCenter](#).

Restore MongoDB, MySQL, and PostgreSQL databases

You can restore MongoDB, MySQL, and PostgreSQL databases to any host.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. In the **Filter By** field, select the filter **Type** and from the drop-down select **MongoDB, MySQL, or PostgreSQL**.
3. Click **View Details** corresponding to the database that you want to restore and click **Restore**.
4. In the Restore options page, specify one of the following:
 - a. For NAS environment, specify the FQDN or IP address of the host to which the volumes restored from object store are to be exported.
 - b. For SAN environment, specify the initiators of the host to which LUNs of the volumes restored from object store are to be mapped.
5. If the snapshot is in archival storage, select the priority to restore your data from the archival storage.
6. If there is not enough space on the source storage or the source storage is down, select **Change storage location**.

If you select **Change storage location**, you can append a suffix to the destination volume. If you have not selected the checkbox, then by default **_restore** is appended to the destination volume.

7. Click **Next**.
8. In the Storage mapping page, specify the alternate storage location details where the data restored from the object store will be stored.
9. Click **Next**.
10. Review the details and click **Restore**.

This operation does only the storage export of the selected backup on the given host. You should manually mount the filesystem and bring up the database. After utilizing the volume, the storage Administrator can delete the volume from the ONTAP cluster.

Back up and restore cloud-native applications data

Protect your cloud-native applications data

BlueXP backup and recovery for applications provides application consistent data protection capabilities for applications running on NetApp Cloud Storage. BlueXP backup and recovery offers efficient, application consistent, policy-based protection of the following applications:

- Oracle databases residing on Amazon FSx for NetApp ONTAP, Cloud Volumes ONTAP, and Azure NetApp Files
- SAP HANA systems residing on Azure NetApp Files
- Microsoft SQL Server databases residing on Amazon FSx for NetApp ONTAP

Architecture

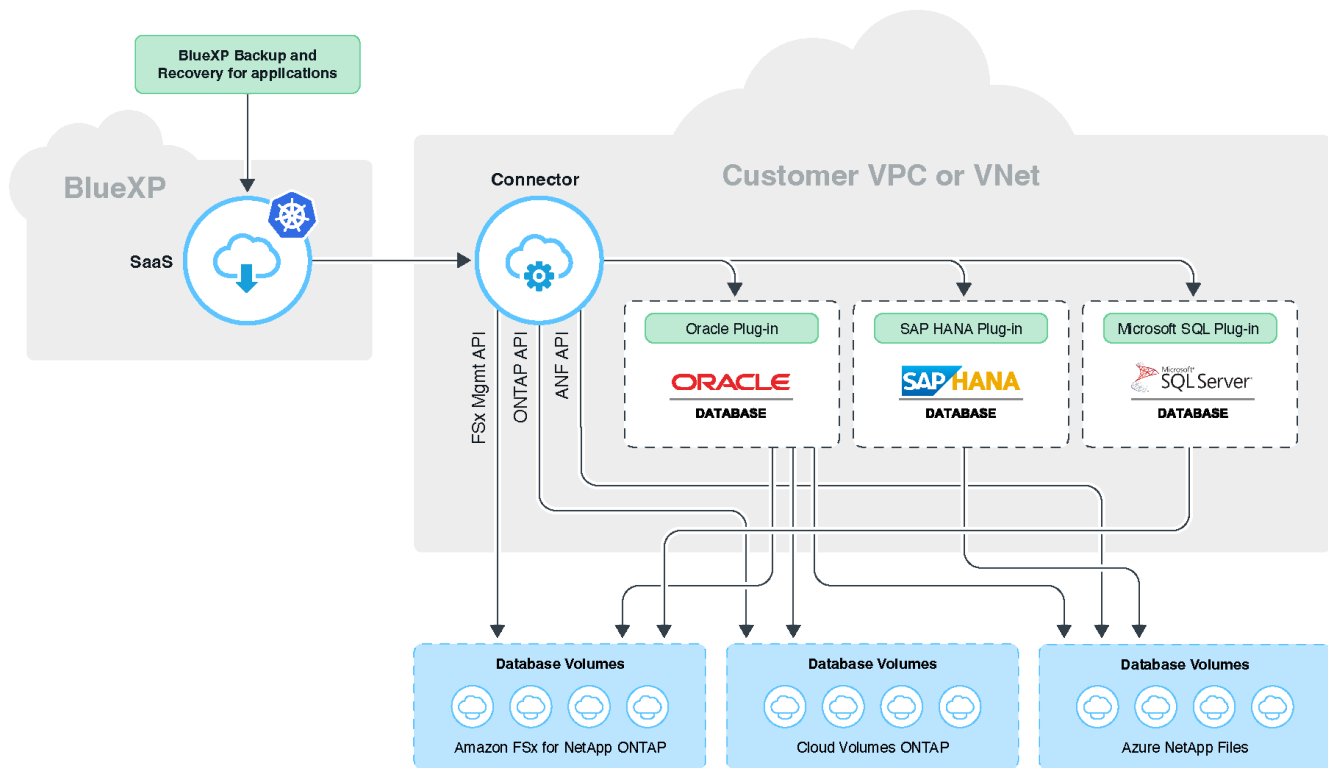
The BlueXP backup and recovery for applications architecture includes the following components.

- The BlueXP backup and recovery is a set of data protection services hosted as a SaaS service by NetApp and is based on the BlueXP SaaS platform.

It orchestrates the data protection workflows for applications residing on NetApp Cloud Storage.

- BlueXP UI offers data protection capabilities for applications and can be accessed from the BlueXP UI.
- BlueXP Connector is a component that runs in your cloud network and interacts with storage systems and application specific plug-ins.
- The application specific plug-in is a component that runs on each application host and interacts with the databases running on the host while performing data protection operations.

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



For any user-initiated request, the BlueXP UI communicates with the BlueXP SaaS which upon validating the request processes the same. If the request is to run a workflow such as a backup, restore, or clone, the SaaS service initiates the workflow and where required, forwards the call to the BlueXP Connector. The Connector then communicates with storage system and application specific plug-in as part of running the workflow tasks.

The Connector can be deployed in the same VPC or VNet as that of the applications, or in a different one. If the Connector and applications are on different network, you should establish a network connectivity between them.



A single BlueXP Connector can communicate with multiple storage systems and multiple application plug-ins. You will need a single Connector to manage your applications as long as there is connectivity between the Connector and application hosts.



The BlueXP SaaS infrastructure is resilient to availability zone failures within a region. It supports regional failures by failing over to a new region and this failover involves a downtime of around 2 hours.

Protect Oracle databases

Features

- Add host and deploy plug-in

You can deploy plugin using UI, script, or manually.

- Auto-discovery of Oracle databases
- Backing up Oracle databases residing on Amazon FSx for NetApp ONTAP, Cloud Volumes ONTAP, and Azure NetApp Files

- Full (data + control + archive log files) backup
- On-demand backup
- Scheduled backup based on the system-defined or custom policies

You can specify different scheduling frequencies such as hourly, daily, weekly, and monthly in the policy. You can also specify the post-scripts that will be executed after successful backup to copy the snapshot to secondary storage.

- Backups of Oracle databases on Azure NetApp Files can be cataloged using Oracle RMAN
- Retaining backups based on the policy
- Restoring Oracle databases residing on Amazon FSx for NetApp ONTAP, Cloud Volumes ONTAP, and Azure NetApp Files
 - Restoring complete Oracle database (data files + control file) from the specified backup
 - Recovering Oracle database with until SCN, until time, all available logs, and no recovery options
- Restoring Oracle databases on Azure NetApp Files to alternate location
- Cloning of Oracle Databases residing on Amazon FSx for NetApp ONTAP and Cloud Volumes ONTAP to source or alternate target hosts
 - Basic one-click clone
 - Advanced cloning using custom clone specification file
 - Clone entities name can be auto-generated or be identical to the source
 - Viewing clone hierarchy
 - Deleting cloned databases
- Monitoring backups, restore, clone, and other jobs
- Displaying the protection summary on the dashboard
- Sending alerts through email
- Upgrade the host plug-in

Limitations

- Does not support Oracle 11g
- Does not support mount, catalog, and verification operations on backups
- Does not support Oracle on RAC and Data Guard
- For Cloud Volumes ONTAP HA, only one of the network interface IPs are used. If the connectivity of the IP goes down or if you cannot access the IP, data protection operations fail.
- The network interface IP addresses of Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP must be unique in the BlueXP account and region.

Protect SAP HANA databases

Features

- Manually add SAP HANA systems
- Backing up SAP HANA databases

- On-demand backup (File-based and Snapshot copy based)
- Scheduled backup based on the system-defined or custom policies

You can specify different scheduling frequencies such as hourly, daily, weekly, and monthly in the policy.

- HANA System Replication (HSR) aware
- Retaining backups based on the policy
- Restoring complete SAP HANA database from the specified backup
- Backing up and restoring HANA Non-Data Volumes and global Non-Data Volumes
- Prescript and postscript support using environmental variables for backup and restore operations
- Creating action plan for failure scenarios using pre-exit option

Limitations

- For HSR configuration, only 2-node HSR is supported (1 primary and 1 secondary)
- Retention will not be triggered if the postscript fails during restore operation

Protect Microsoft SQL Server database

Features

- Manually add host and deploy the plug-in
- Discover the databases manually
- Back up SQL Server instances residing on Amazon FSx for NetApp ONTAP
 - On-demand backup
 - Scheduled backup based on the policy
 - Log backup of Microsoft SQL Server instance
- Restore the database to original location

Limitations

- Backup is supported only for SQL Server instances
- Failover Cluster Instance (FCI) configuration is not supported
- BlueXP UI does not support SQL database specific operations

All Microsoft SQL Server database specific operations are performed by running REST APIs.

- Restore to alternate location is not supported

Back up cloud-native Oracle databases

Quick start

Get started quickly by following these steps.

1

Verify support for your configuration

- Operating System:
 - RHEL 7.5 or later and 8.x
 - OL 7.5 or later and 8.x
 - SLES 15 SP4
- NetApp Cloud Storage:
 - Amazon FSx for NetApp ONTAP
 - Cloud Volumes ONTAP
 - Azure NetApp Files
- Storage layouts:
 - NFS v3 and v4.1 (including dNFS)
 - iSCSI with ASM (ASMFD, ASMLib and ASMUdev)



Azure NetApp Files does not support SAN environment.

- Database layouts: Oracle Standard and Oracle Enterprise Standalone (legacy and multitenant CDB and PDB)
- Database versions: 19c and 21c

2

Sign up to BlueXP

BlueXP is accessible from a web-based console. When you get started with BlueXP, your first step is to sign up using your existing NetApp Support Site credentials or by creating a NetApp cloud login. For information, refer to [Sign up to BlueXP](#).

3

Log into BlueXP

After you sign up to BlueXP, you can log in from the web-based console. For information, refer to [Log into BlueXP](#).

4

Manage your BlueXP account

You can administer your account by managing users, service accounts, workspaces, and Connectors. For information, refer to [Manage your BlueXP account](#).

Configure FSx for ONTAP

Using BlueXP you should create an FSx for ONTAP working environment to add and manage volumes and additional data services. You should also create a Connector in AWS that enables BlueXP to manage resources and processes within your public cloud environment.

Create FSx for ONTAP working environment

You should create the FSx for ONTAP working environments where your databases are hosted. For information, refer to [Get started with Amazon FSx for ONTAP](#) and [Create and manage an Amazon FSx for ONTAP working environment](#).

You can create the FSx for ONTAP working environment either using BlueXP or AWS. If you have created using AWS, then you should discover the FSx for ONTAP systems in BlueXP.

Create a Connector

An Account Admin needs to create a Connector in AWS that enables BlueXP to manage resources and processes within your public cloud environment.

For information, refer to [Creating a Connector in AWS from BlueXP](#).

- You should use the same connector to manage both FSx for ONTAP working environment and databases.
- If you have the FSx for ONTAP working environment and databases in the same virtual private cloud (VPC), you can deploy the connector in the same VPC.
- If you have the FSx for ONTAP working environment and databases in different VPCs:
 - If you have NAS (NFS) workloads configured on FSx for ONTAP, then you can create the connector on either of the VPCs.
 - If you have only SAN workloads configured and not planning to use any NAS (NFS) workloads, then you should create the connector in the VPC where the FSx for ONTAP system is created.



For using NAS (NFS) workloads, you should have transit gateway between the database VPC and Amazon VPC. The NFS IP address which is a floating IP address can be accessed from another VPC only through transit gateway. We cannot access the floating IP addresses by peering the VPCs.

After creating the Connector, click **Storage > Canvas > My Working Environments > Add Working Environment** and follow the prompts to add the working environment.

Ensure that there is connectivity from the Connector to the Oracle database hosts and FSx working environment. The Connector should be able to connect to the cluster management IP address of the FSx working environment.

- Add the working environment by clicking **Storage > Canvas > My Working Environments > Add Working Environment**.

Ensure that there is connectivity from the connector to the database hosts and FSx for ONTAP working environment. The connector should connect to the cluster management IP address of the FSx for ONTAP working environment.

- Copy the Connector ID by clicking **Connector > Manage Connectors** and selecting the Connector name.

Configure Cloud Volumes ONTAP

Using BlueXP you should create a Cloud Volumes ONTAP working environment to add and manage volumes and additional data services. You should also create a Connector for your cloud environment that enables BlueXP to manage resources and processes within your public cloud environment.

Create Cloud Volumes ONTAP working environment

You can discover and add existing Cloud Volumes ONTAP systems to BlueXP. For information, refer to [Adding existing Cloud Volumes ONTAP systems to BlueXP](#).

Create a Connector

You can get started with Cloud Volumes ONTAP for your cloud environment in a few steps. For more information, refer one of the following:

- [Quick start for Cloud Volumes ONTAP in AWS](#)
- [Quick start for Cloud Volumes ONTAP in Azure](#)
- [Quick start for Cloud Volumes ONTAP in Google Cloud](#)

You should use the same connector to manage both Cloud Volumes ONTAP working environment and databases.

- If you have the Cloud Volumes ONTAP working environment and databases in the same virtual private cloud (VPC) or VNet, you can deploy the connector in the same VPC or VNet.
- If you have the Cloud Volumes ONTAP working environment and databases in different VPCs or VNets, ensure that the VPCs or VNets are peered.

Configure Azure NetApp Files

Using BlueXP you should create a Azure NetApp Files working environment to add and manage volumes and additional data services. You should also create a Connector in Azure that enables BlueXP to manage resources and processes within your public cloud environment.

Create Azure NetApp Files working environment

You should create Azure NetApp Files working environments where your databases are hosted. For more information, refer to [Learn about Azure NetApp Files](#) and [Create an Azure NetApp Files working environment](#).

Create a connector

A BlueXP account admin should deploy a Connector in Azure that enables BlueXP to manage resources and processes within your public cloud environment.

For information, refer to [Create a Connector in Azure from BlueXP](#).

- Ensure that there is connectivity from the connector to the database hosts.
- If you have the Azure NetApp Files working environment and databases in the same Virtual Network (VNet), you can deploy the connector in the same VNet.
- If you have the Azure NetApp Files working environment and databases in different VNets and have NAS (NFS) workloads configured on Azure NetApp Files, then you can create the connector on either of the VNets.

After creating the connector, add the working environment by clicking **Storage > Canvas > My Working Environments > Add Working Environment**.

Install SnapCenter Plug-in for Oracle and add database hosts

You should install the SnapCenter Plug-in for Oracle on each of the Oracle database hosts, add the database hosts, and discover the databases on the host to assign policies and create backups.

- If SSH is enabled for the database host, you can install the plug-in using one of the methods:
 - Install the plug-in and add host from the UI using SSH option. [Learn more](#).
 - Install the plug-in using script and add host from the UI using manual option. [Learn more](#).
- If SSH is disabled, install the plug-in manually and add host from the UI using manual option. [Learn more](#).

Prerequisites

Before adding the host, you should ensure that the prerequisites are met.

- You should have created the working environment and the Connector.
- Ensure that the Connector has connectivity to the Oracle database hosts.

For information on how to resolve the connectivity issue, refer to [Failed to validate connectivity from BlueXP connector host to application database host](#).

When the connector is lost or if you have created a new connector, you should associate the connector with the existing application resources. For instructions to update the Connector, see [Update the Connector Details](#).

- Ensure that the BlueXP user has the “Account Admin” role.
- Ensure that non root (sudo) account is present on the application host for data protection operations.
- Ensure that either Java 11 (64-bit) Oracle Java or OpenJDK is installed on each of the Oracle database hosts and the JAVA_HOME variable is set appropriately.
- Ensure that the Connector has the communication enabled to the SSH port (default: 22) if SSH based installation is performed.
- Ensure that the Connector has the communication enabled to plug-in port (default: 8145) for the data protection operations to work.
- Ensure that the you have the latest version of plug-in is installed. To upgrade the plug-in, refer to [Upgrade SnapCenter Plug-in for Oracle Database](#).

Add host from UI using SSH option

Steps

1. In the BlueXP UI, click **Protection > Backup and recovery > Applications**.

If you have already added a host and want to add another host, click **Applications > Manage Databases > Add** and then proceed with step 5.

2. Click **Discover Applications**.
3. Select **Cloud Native** and click **Next**.

A service account (*SnapCenter-account-`<accountid>`*) with *SnapCenter System* role is created to perform scheduled data protection operations for all the users in this account.

The service account (*SnapCenter-account-**<accountid>***) is used to run the scheduled backup operations. You should never delete the service account.

You can view the service account by clicking **Account > Manage Account > Members**.

4. Select Oracle as the application type.
5. In the Host details page, perform the following:

- a. Select **Using SSH**.
- b. Specify the FQDN or IP address of the host where you want to install the plug-in.

Ensure that the Connector can communicate with the database host using the FQDN or IP address.

- c. Specify the non-root(sudo) user using which the plug-in package will be copied to the host.

Root user is not supported.

- d. Specify the SSH and plug-in port.

Default SSH port is 22 and the plug-in port is 8145.

You can close the SSH port on the application host after installing the plug-in. The SSH port is not required for any data protection operations.

- e. Select the Connector.
- f. (Optional) If key less authentication is not enabled between the Connector and the host, you should specify the SSH private key that will be used to communicate with the host.



The SSH private key is not stored anywhere in the application and is not used for any other operations.

- g. Click **Next**.
6. In the Configuration page, perform the following:
 - a. Configure sudo access for the SnapCenter user in the Oracle database host by logging in to the Linux machine running Oracle database.
 - b. Copy the text displayed in BlueXP UI.
 - c. Create the */etc/sudoers.d/snapcenter* file on the Linux machine and paste the copied text.
 - d. In the BlueXP UI, select the checkbox and click **Next**.
7. Review the details and click **Discover Applications**.

- After the plug-in is installed, the discovery operation starts.
- After completing the discovery operation, all the databases on the host are displayed. If OS authentication is disabled for the database, click **Configure** to enable database authentication. For more information, refer to [Configure Oracle database credentials](#).
- Click **Settings** and select **Hosts** to view all the hosts.
- Click **Settings** and select **Policies** to view the pre-canned policies. Review the pre-canned policies and you can either edit them to meet your requirement or create a new policy.

Add host from UI using manual option and install the plug-in using script

Configure SSH key based authentication for the Oracle host non-root user account and perform the following

steps to install the plug-in.

Before you begin

Ensure that the SSH connection to the Connector is enabled.

Steps

1. In the BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. Click **Discover Applications**.
3. Select **Cloud Native** and click **Next**.

A service account (*SnapCenter-account-**<accountid>***) with *SnapCenter System* role is created to perform scheduled data protection operations for all the users in this account.

The service account (*SnapCenter-account-**<accountid>***) is used to run the scheduled backup operations.

You should never delete the service account.

You can view the service account by clicking **Account > Manage Account > Members**.

4. Select Oracle as the application type.
5. In the Host details page, perform the following:
 - a. Select **Manual**.
 - b. Specify the FQDN or IP address of the host where the plug-in is installed.

Ensure that the Connector can communicate with the database host using the FQDN or IP address.

- c. Specify the plug-in port.

Default port is 8145.

- d. Specify the non-root (sudo) user using which the plug-in package will be copied to the host.
- e. Select the Connector.
- f. Select the check box to confirm that the plug-in is installed on the host.
- g. Click **Next**.

6. In the Configuration page, perform the following:
 - a. Configure sudo access for the SnapCenter user in the Oracle database host by logging in to the Linux machine running Oracle database.
 - b. Copy the text displayed in BlueXP UI.
 - c. Create the */etc/sudoers.d/snapcenter* file on the Linux machine and paste the copied text.
 - d. In the BlueXP UI, select the checkbox and click **Next**.

7. Log into the Connector VM.

8. Install the plug-in using the script provided in the Connector.

```
sudo /var/lib/docker/volumes/service-manager-  
2_cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh  
--host <plugin_host> --username <host_user_name> --sshkey <host_ssh_key>  
--pluginport <plugin_port> --sshport <host_ssh_port>
```

If you are using an older Connector, run the following command to install the plug-in.

```
sudo  
/var/lib/docker/volumes/cloudmanager_scs_cloud_volume/_data/scripts/linux_plug
```

```
in_copy_and_install.sh --host <plugin_host> --username <host_user_name>
--sshkey <host_ssh_key> --pluginport <plugin_port> --sshport <host_ssh_port>
```

Name	Description	Mandatory	Default
plugin_host	Specifies the Oracle host	Yes	-
host_user_name	Specifies the SnapCenter user with SSH privileges on the Oracle host	Yes	-
host_ssh_key	Specifies the SSH key of the SnapCenter user and is used to connect to the Oracle host	Yes	-
plugin_port	Specifies the port used by the plug-in	No	8145
host_ssh_port	Specifies the SSH port on the Oracle host	No	22

For example:

- `sudo /var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh --host 10.0.1.1 --username snapcenter --sshkey /keys/netapp-ssh.ppk`
- `sudo /var/lib/docker/volumes/cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh --host 10.0.1.1 --username snapcenter --sshkey /keys/netapp-ssh.ppk`

9. In the BlueXP UI, review the details and click **Discover Applications**.

- After completing the discovery operation, all the databases on the host are displayed. If OS authentication is disabled for the database, click **Configure** to enable database authentication. For more information, refer to [Configure Oracle database credentials](#).
- Click **Settings** and select **Hosts** to view all the hosts.
- Click **Settings** and select **Policies** to view the pre-canned policies. Review the pre-canned policies and you can either edit them to meet your requirement or create a new policy.

Add host from UI using manual option and install the plug-in manually

If SSH key based authentication is not enabled on the Oracle database host, you should perform the following manual steps to install the plug-in and then add the host from UI using manual option.

Steps

1. In the BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. Click **Discover Applications**.

3. Select **Cloud Native** and click **Next**.

A service account (*SnapCenter-account-**<accountid>***) with *SnapCenter System* role is created to perform scheduled data protection operations for all the users in this account.

The service account (*SnapCenter-account-**<accountid>***) is used to run the scheduled backup operations. You should never delete the service account.

You can view the service account by clicking **Account > Manage Account > Members**.

4. Select Oracle as the application type.
5. In the **Host details** page, perform the following:

- a. Select **Manual**.
- b. Specify the FQDN or IP address of the host where the plug-in is installed.

Ensure that using the FQDN or IP address, the Connector can communicate with the database host.

- c. Specify the plug-in port.

Default port is 8145.

- d. Specify the sudo non-root (sudo) user using which the plug-in package will be copied to the host.
- e. Select the Connector.
- f. Select the check box to confirm that the plug-in is installed on the host.
- g. Click **Next**.

6. In the Configuration page, perform the following:

- a. Configure sudo access for the SnapCenter user in the Oracle database host by logging in to the Linux machine running Oracle database.
- b. Copy the text displayed in BlueXP UI.
- c. Create the */etc/sudoers.d/snapcenter* file on the Linux machine and paste the copied text.
- d. In the BlueXP UI, select the checkbox and click **Next**.

7. Log into the Connector VM.

8. Download the SnapCenter Linux host plug-in binary.

```
sudo docker exec -it cloudmanager_scs_cloud curl -X GET  
'http://127.0.0.1/deploy/downloadLinuxPlugin'
```

The plug-in binary is available at: *cd /var/lib/docker/volumes/service-manager[1]-2_cloudmanager_scs_cloud_volume/_data/\$(sudo docker ps|grep -Po "cloudmanager_scs_cloud:.?*"|sed -e 's/ *\$//'|cut -f2 -d":")/sc-linux-host-plugin*

9. Copy *snapcenter_linux_host_plugin_scs.bin* from the above path to */home/<non root user>/.sc_netapp* path for each of the Oracle database hosts either using scp or other alternate methods.
10. Log into the Oracle database host using the non-root (sudo) account.

11. Change directory to */home/<non root user>/.sc_netapp/* and run the following command to enable execute permissions for the binary.

```
chmod +x snapcenter_linux_host_plugin_scs.bin
```

12. Install the Oracle plug-in as a sudo SnapCenter user.

```
./snapcenter_linux_host_plugin_scs.bin -i silent -DSPL_USER=<non-root>
```

13. Copy *certificate.pem* from `<base_mount_path>/client/certificate/` path of the Connector VM to `/var/opt/snapcenter/spl/etc/` on the plug-in host.
14. Navigate to `/var/opt/snapcenter/spl/etc` and execute the `keytool` command to import the *certificate.pem*.

```
keytool -import -alias agentcert -file certificate.pem -keystore keystore.jks  
-deststorepass snapcenter -noprompt
```
15. Restart SPL: `systemctl restart spl`
16. Validate that the plug-in is reachable from the Connector by running the below command from the Connector.

```
docker exec -it cloudmanager_scs_cloud curl -ik https://<FQDN or IP of the  
plug-in host>:<plug-in port>/PluginService/Version --cert  
/config/client/certificate/certificate.pem --key  
/config/client/certificate/key.pem
```
17. In the BlueXP UI, review the details and click **Discover Applications**.
 - After completing the discovery operation, all the databases on the host are displayed. If OS authentication is disabled for the database, click **Configure** to enable database authentication. For more information, refer to [Configure Oracle database credentials](#).
 - Click **Settings** and select **Hosts** to view all the hosts.
 - Click **Settings** and select **Policies** to view the pre-canned policies. Review the pre-canned policies and you can either edit them to meet your requirement or create a new policy.

Configure Oracle database credentials

You should configure the database credentials that are used to perform data protection operations on Oracle databases.

Steps

1. If OS authentication is disabled for the database, click **Configure** to modify database authentication.
2. Specify the username, password, and the port details.

If the database is residing on ASM, you should also configure the ASM settings.

The Oracle user should have `sysdba` privileges and ASM user should have `sysasm` privileges.

3. Click **Configure**.

Upgrade SnapCenter Plug-in for Oracle Database

You should upgrade the SnapCenter Plug-in for Oracle to gain access to the latest new features and enhancements. You can upgrade from the BlueXP UI or using the command line.

Before you begin

- Ensure that there are no operations running on the host.

Steps

1. Click **Backup and recovery > Applications > Hosts**.
2. Verify if plug-in upgrade is available for any of the hosts by checking the Overall Status column.
3. Upgrade the plug-in from UI or using the command line.

Upgrade using UI	Upgrade using command line
<ol style="list-style-type: none"> 1. Click ... corresponding to the host and click Upgrade Plug-in. 2. In the Configuration page, perform the following: <ol style="list-style-type: none"> a. Configure sudo access for the SnapCenter user in the Oracle database host by logging in to the Linux machine running Oracle database. b. Copy the text displayed in BlueXP UI. c. Edit the <code>/etc/sudoers.d/snapcenter</code> file on the Linux machine and paste the copied text. d. In the BlueXP UI, select the checkbox and click Upgrade. 	<ol style="list-style-type: none"> 1. Log in to Connector VM. 2. Run the following script. <pre> sudo /var/lib/docker/volumes/service- manager- 2_cloudmanager_scs_cloud_volume/_da ta/scripts/linux_plugin_copy_and_in stall.sh --host <plugin_host> --username <host_user_name> --sshkey <host_ssh_key> --pluginport <plugin_port> --sshport <host_ssh_port> --upgrade </pre> <p>If you are using an older Connector, run the following command to upgrade the plug-in.</p> <pre> sudo /var/lib/docker/volumes/cloudmanage r_scs_cloud_volume/_data/scripts/li nux_plugin_copy_and_install.sh --host <plugin_host> --username <host_user_name> --sshkey <host_ssh_key> --pluginport <plugin_port> --sshport <host_ssh_port> --upgrade </pre>

Back up cloud-native Oracle databases

You can create scheduled or on-demand backups by assigning a pre-canned policy or the policy that you created.

You can also catalog the Oracle database backups using Oracle Recovery Manager (RMAN) if you have enabled cataloging while creating a policy. The (RMAN) cataloging is supported only for the databases that are on Azure NetApp Files. The cataloged backups can be used later for block-level restore or tablespace point-in-time recovery operations. The database must be in mounted or higher state for cataloging.

Create policy to protect Oracle database

You can create policies if you do not want to edit the pre-canned policies.

Steps

1. In the Applications page, from the Settings drop-down list, select **Policies**.
2. Click **Create policy**.
3. Specify a policy name.
4. (Optional) Edit the format of the backup name.
5. Specify the schedule and retention details.

6. If you have selected *daily* and *weekly* as the schedule and you want to enable RMAN cataloging, select **Catalog backup with Oracle Recovery Manager (RMAN)**.
7. (Optional) Enter the post-script path and timeout value for post-script that will be executed after the successful backup such as copying the snapshot to secondary storage.

Optionally, you can also specify the arguments.

You should keep the post-scripts in the path `/var/opt/snapcenter/spl/scripts`.

The post script supports a set of environment variables.

Environmental Variable	Description
SC_ORACLE_SID	Specifies the SID of the Oracle database.
SC_HOST	Specifies the hostname of the database
SC_BACKUP_NAME	Specifies the name of the backup. The data backup name and the log backup name are concatenated using delimiters.
SC_BACKUP_POLICY_NAME	Specifies the name of the policy used to create the backup.
SC_PRIMARY_DATA_VOLUME_FULL_PATH	Specifies the data volume paths concatenated using "," as delimiter. For Azure NetApp Files volumes, the information is concatenated using "/" — /subscriptions/{subscription_id}/resourceGroups/{resource_group}/providers/{provider}/netAppAccounts/{anfaccount}/capacityPools/{capacity_pool}/volumes/{volumename}_
SC_PRIMARY_ARCHIVELOGS_VOLUME_FULL_PATH	Specifies the archive log volume paths concatenated using "," as delimiter. For Azure NetApp Files volumes, the information is concatenated using "/" — /subscriptions/{subscription_id}/resourceGroups/{resource_group}/providers/{provider}/netAppAccounts/{anfaccount}/capacityPools/{capacity_pool}/volumes/{volumename}_

8. Click **Create**.



Configure RMAN catalog repository

You can configure the recovery catalog database as the RMAN catalog repository. If you do not configure the repository, by default, the Control file of the target database becomes the RMAN catalog repository.

Before you begin

You should manually register the target database with RMAN catalog database.

Steps

1. In the Applications page, click  > **View Details**.
2. In the Database details section, click  to configure the RMAN catalog repository.
3. Specify the credentials to catalog backups with RMAN and the Transparent Network Substrate (TNS) name of catalog recovery database.
4. Click **Configure**.

Create a backup of the Oracle Database


You can assign a pre-canned policy or create a policy and then assign it to the database. Once the policy is assigned, the backups are created as per the schedule defined in the policy.



When creating ASM diskgroups on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP, ensure that there are no common volumes across diskgroups. Each diskgroup should have dedicated volumes.

Steps

1. In the Applications page, if the database is not protected using any policy, click **Assign Policy**.

If the database is protected using one or more policies, you can assign more policies by clicking  > **Assign Policy**.

2. Select the policy and click **Assign**.

The backups will be created as per the schedule defined in the policy. If you have enabled RMAN catalog in the policy, the backup at the end of the workflow will launch the cataloging operation as a separate job. The progress of cataloging can be seen from the Job Monitor. Upon successful cataloging, **Backup Details** will show the status of the catalog for each backup.




The service account (*SnapCenter-account-`<account_id>`*) is used to run the scheduled backup operations.

Create on-demand backup of the Oracle database

After assigning the policy, you can create an on-demand backup of the application.

Steps

1. In the Applications page, click  corresponding to the application and click **On-Demand Backup**.
2. If multiple policies are assigned to the application, select the policy, retention tier, and then click **Create Backup**.

If you have enabled RMAN catalog in the policy, the backup at the end of the workflow will launch the cataloging operation as a separate job. The progress of cataloging can be seen from the Job Monitor. Upon successful cataloging, **Backup Details** will show the status of the catalog for each backup.

Limitations

- Does not support consistency group Snapshots for Oracle databases residing on Multiple ASM disk groups with overlap of FSx volumes
- If your Oracle databases are on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP and are configured on ASM, ensure your SVM names are unique across the FSx systems. If you have same SVM name across FSx systems, back up of Oracle databases residing on those SVMs are not supported.
- After restoring a large database (250 GB or more), if you perform a full online backup on the same database the operation might fail with the following error:
failed with status code 500, error
`{"error\":{\"code\":\"app_internal_error\",\"message\":\"Failed to create snapshot. Reason: Snapshot operation not allowed due to clones backed by snapshots. Try again after sometime.\"}}`

For information on how to fix this issue, refer to: [Snapshot operation not allowed due to clones backed by snapshots](#).

Back up cloud-native SAP HANA databases

Quick start

Get started quickly by following these steps.

1

Verify support for your configuration

- Operating System:
 - RHEL 7.6 or later
 - RHEL 8.1 or later for SAP-HANA SPS07
 - SLES 12 SP5 or later and 15 SPX platforms certified by SAP HANA
- NetApp Cloud Storage: Azure NetApp Files
- Storage layouts: For data and log files, Azure supports only NFSv4.1.
- Database layouts:
 - SAP HANA Multitenant Database Container (MDC) 2.0SPS5, 2.0SPS6, 2.0SPS7 with single or multiple tenants
 - SAP HANA single host system, SAP HANA multiple host system, HANA System Replication
- SAP HANA plug-in on the database host

2

Sign up to BlueXP

BlueXP is accessible from a web-based console. When you get started with BlueXP, your first step is to sign up using your existing NetApp Support Site credentials or by creating a NetApp cloud login. For information, refer to [Sign up to BlueXP](#).

3

Log into BlueXP

After you sign up to BlueXP, you can log in from the web-based console. For information, refer to [Log into BlueXP](#).



Manage your BlueXP account

You can administer your account by managing users, service accounts, workspaces, and Connectors. For information, refer to [Manage your BlueXP account](#).

Configure Azure NetApp Files

Using BlueXP you should create a Azure NetApp Files working environment to add and manage volumes and additional data services. You should also create a Connector in Azure that enables BlueXP to manage resources and processes within your public cloud environment.

Create Azure NetApp Files working environment

You should create Azure NetApp Files working environments where your databases are hosted. For more information, refer to [Learn about Azure NetApp Files](#) and [Create an Azure NetApp Files working environment](#).

Create a connector

A BlueXP account admin should deploy a Connector in Azure that enables BlueXP to manage resources and processes within your public cloud environment.

For information, refer to [Create a Connector in Azure from BlueXP](#).

- Ensure that there is connectivity from the connector to the database hosts.
- If you have the Azure NetApp Files working environment and databases in the same Virtual Network (VNet), you can deploy the connector in the same VNet.
- If you have the Azure NetApp Files working environment and databases in different VNets and have NAS (NFS) workloads configured on Azure NetApp Files, then you can create the connector on either of the VNets.

After creating the connector, add the working environment by clicking **Storage > Canvas > My Working Environments > Add Working Environment**.

Install SnapCenter Plug-in for SAP HANA and add database hosts

You should install the SnapCenter Plug-in for SAP HANA on each of the SAP HANA database hosts. Depending on whether the SAP HANA host has an SSH key based authentication enabled, you can follow one of the methods to install the plug-in.

- If SSH is enabled for the database host, you can install the plug-in using SSH option. [Learn more](#).
- If SSH is disabled, install the plug-in manually. [Learn more](#).

Prerequisites

Before adding the host, you should ensure that the prerequisites are met.

- Ensure that either Java 11 (64-bit) Oracle Java or OpenJDK is installed on each of the SAP HANA database hosts.
- You should have added the working environment and created the Connector.
- Ensure that the Connector has connectivity to the SAP HANA database hosts.

For information on how to resolve the connectivity issue, refer to [Failed to validate connectivity from BlueXP connector host to application database host](#).

When the connector is lost or if you have created a new connector, you should associate the connector with the existing application resources. For instructions to update the Connector, see [Update the Connector Details](#).

- Ensure that the BlueXP user has the “Account Admin” role.
- You should have created the SnapCenter user and configured sudo for the non-root (sudo) user. For information, refer to [Configure sudo for SnapCenter user](#).
- You should have installed the SnapCenter Plug-in for SAP HANA before adding the database host.
- While adding the SAP HANA database hosts, you should add the HDB user store keys. The HDB secure user store key is used to store the connection information of SAP HANA database hosts securely on the client and HDBSQL client uses the secure user store key to connect to SAP HANA database host.
- For HANA System Replication (HSR), to protect the HANA systems, you should manually register both primary and secondary HANA systems.



The hostname must be the same as that of the host that is used in the HSR replication.

- Ensure that the Connector has the communication enabled to the SSH port (default: 22) if SSH based installation is performed.
- Ensure that the Connector has the communication enabled to plug-in port (default: 8145) for the data protection operations to work.
- Ensure that the you have the latest version of plug-in is installed. To upgrade the plug-in, refer to [Upgrade SnapCenter Plug-in for SAP HANA Database](#).

Configure sudo for SnapCenter user

Create a non-root (sudo) user to install the plug-in.

Steps

1. Log into the Connector VM.
2. Download the SnapCenter Linux host plug-in binary.

```
sudo docker exec -it cloudmanager_scs_cloud curl -X GET 'http://127.0.0.1/deploy/downloadLinuxPlugin'
```
3. Copy the contents of **sudoeer.txt** located at: `/var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/$(sudo docker ps|grep -Po "cloudmanager_scs_cloud:.?*"|sed -e 's/ *$//'|cut -f2 -d":")/sc-linux-host-plugin`
4. Log into the SAP HANA system host using root user account.
5. Configure sudo access for the non-root user by copying the text copied in the step 3 to `/etc/sudoers.d/snapcenter` file.

In the lines you added to the `/etc/sudoers.d/snapcenter` file, replace the `<LINUXUSER>` with the non-root

user and <USER_HOME_DIRECTORY> with *home/<non-root-user>*.

Install the plug-in using script

Configure SSH key based authentication for the SAP HANA host non-root user account and perform the following steps to install the plug-in.

Before your begin

Ensure that the SSH connection to the Connector is enabled.

Steps

1. Log into Connector VM.
2. Install the plug-in using the script provided in the Connector.

```
sudo bash /var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh  
--host <plugin_host> --username <host_user_name> --sshkey <host_ssh_key>  
--pluginport <plugin_port> --sshport <host_ssh_port>
```

If you are using an older Connector, run the following command to install the plug-in.

```
sudo  
/var/lib/docker/volumes/cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh --host <plugin_host> --username <host_user_name>  
--sshkey <host_ssh_key> --pluginport <plugin_port> --sshport <host_ssh_port>
```

Name	Description	Mandatory	Default
plugin_host	Specifies the SAP HANA host	Yes	-
host_user_name	Specifies the SnapCenter user with SSH privileges on the SAP HANA host	Yes	-
host_ssh_key	Specifies the SSH key of the SnapCenter user and is used to connect to the SAP HANA host	Yes	-
plugin_port	Specifies the port used by the plug-in	No	8145
host_ssh_port	Specifies the SSH port on the SAP HANA host	No	22

For example, ``sudo bash /var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh --host 10.0.1.1 --username snapcenter --sshkey /keys/netapp-ssh.ppk``

After installing the plug-in, you should [Add SAP HANA database hosts](#).

Install the plug-in manually

If SSH key based authentication is not enabled on the HANA host, you should perform the below manual steps to install the plug-in.

Steps

1. Log in to Connector VM.

2. Download the SnapCenter Linux host plug-in binary.

```
sudo docker exec -it cloudmanager_scs_cloud curl -X GET  
'http://127.0.0.1/deploy/downloadLinuxPlugin'
```

The plug-in binary is available at: `cd /var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/$(sudo docker ps|grep -Po "cloudmanager_scs_cloud:.*?"|sed -e 's/*$//'|cut -f2 -d":")/sc-linux-host-plugin`

3. Copy `snapcenter_linux_host_plugin_scs.bin` from the above path to `/home/<non root user>/.sc_netapp` path for each of the SAP HANA database hosts either using scp or other alternate methods.

4. Log into the SAP HANA database host using the non-root (sudo) account.

5. Change directory to `/home/<non root user>/.sc_netapp/` and run the following command to enable execute permissions for the binary.

```
chmod +x snapcenter_linux_host_plugin_scs.bin
```

6. Install the SAP HANA plug-in as a sudo SnapCenter user.

```
./snapcenter_linux_host_plugin_scs.bin -i silent -DSPL_USER=<non-root>
```

7. Copy `certificate.pem` from `<base_mount_path>/client/certificate/` path of the Connector VM to `/var/opt/snapcenter/spl/etc/` on the plug-in host.

8. Navigate to `/var/opt/snapcenter/spl/etc` and execute the keytool command to import the certificate.

```
keytool -import -alias agentcert -file certificate.pem -keystore keystore.jks  
-deststorepass snapcenter -noprompt
```

9. Restart SPL: `systemctl restart spl`

10. Validate that the plug-in is reachable from the Connector by running the below command from the Connector.

```
docker exec -it cloudmanager_scs_cloud curl -ik https://<FQDN or IP of the  
plug-in host>:<plug-in port>/PluginService/Version --cert  
config/client/certificate/certificate.pem --key  
/config/client/certificate/key.pem
```

After installing the plug-in, you should [Add SAP HANA database hosts](#).

Upgrade SnapCenter Plug-in for SAP HANA Database

You should upgrade the SnapCenter Plug-in for SAP HANA database to gain access to the latest new features and enhancements.

Before you begin

- Ensure that there are no operations running on the host.

Steps

1. Configure sudo for SnapCenter user. For information, see [Configure sudo for SnapCenter user](#).

2. Run the following script.

```
/var/lib/docker/volumes/service-manager-  
2_cloudmanager_scs_cloud_volume/_data/scripts/linux_plugin_copy_and_install.sh  
--host <plugin_host> --username <host_user_name> --sshkey <host_ssh_key>  
--pluginport <plugin_port> --sshport <host_ssh_port> --upgrade
```

If you are using an older Connector, run the following command to upgrade the plug-in.

```
/var/lib/docker/volumes/cloudmanager_scs_cloud_volume/_data/scripts/linux_plug  
in_copy_and_install.sh --host <plugin_host> --username <host_user_name>  
--sshkey <host_ssh_key> --pluginport <plugin_port> --sshport <host_ssh_port>  
--upgrade
```

Add SAP HANA database hosts

You should manually add SAP HANA database hosts to assign policies and create backups. Auto discovery of SAP HANA database host is not supported.

Steps

1. In the **BlueXP** UI, select **Protection > Backup and recovery > Applications**.
2. Select **Discover Applications**.
3. Select **Cloud Native > SAP HANA** and select **Next**.
4. In the **Applications** page, select **Add System**.
5. In the **System Details** page, perform the following actions:
 - a. Select the System Type as Multi-tenant database container or Global Non-Data Volumes.
 - b. Enter the SAP HANA system name.
 - c. Specify the SID of the SAP HANA system.
 - d. (Optional) Modify OSDB user.
 - e. If HANA system is configured with HANA System replication, enable **HANA System Replication (HSR) System**.
 - f. Select **HDB Secure User Store Keys** text box to add user store keys details.

Specify the key name, system details, username, and password and click **Add Key**.

You can delete or modify the user store keys.

6. Select **Next**.
7. In the **Host Details** page, perform the following actions:
 - a. Select **Add new host** or **Use existing host**.
 - b. Select **Using SSH** or **Manual**.

For Manual, enter the Host FQDN or IP, Connector, Username, SSH port, Plug-in port, and optionally add and validate the SSH private key.

For SSH, enter the Host FQDN or IP, Connector, Username, and Plug-in port.

c. Select **Next**.

8. In the **Host configuration** page, verify whether the configuration requirements are met.

Select the check boxes to confirm.

9. Select **Next**.

10. In the **Storage Footprint** page, select **Add Storage** and perform the following:

a. Select the working environment and specify the NetApp account.

From the left navigation pane, select BlueXP **Canvas** to add a new working environment.

b. Select the required volumes.

c. Select **Add Storage**.

11. Review all the details and select **Add System**.

You can modify or remove the SAP HANA systems from the UI.

Before removing the SAP HANA system, you should delete all the associated backups and remove the protection.

Add Non-Data Volumes

After adding the multi-tenant database container type SAP HANA system, you can add the Non-Data Volumes of the HANA system.

You can add these resources to resource groups to perform data protection operations after you discover the SAP HANA databases that are available.

Steps

1. In the **BlueXP** UI, click **Protection > Backup and recovery > Applications**.
2. Click **Discover Applications**.
3. Select **Cloud Native > SAP HANA** and click **Next**.
4. In the **Applications** page, click **...** corresponding to the system for which you want to add the Non-Data Volumes and select **Manage System > Non-Data Volume**.

Add Global Non-Data Volumes

After adding the multi-tenant database container type SAP HANA system, you can add the Global Non-Data Volumes of the HANA system.

Steps

1. In the **BlueXP** UI, click **Protection > Backup and recovery > Applications**.
2. Click **Discover Applications**.
3. Select **Cloud Native > SAP HANA** and click **Next**.
4. In the **Applications** page, click **Add System**.
5. In the **System Details** page, perform the following actions:
 - a. From System Type drop-down, select **Global Non-Data Volume**.

- b. Enter the SAP HANA system name.
6. . In the **Host Details** page, perform the following actions:
 - a. Specify the associated SIDs of the SAP HANA system.
 - b. Select the plug-in host
 - c. Click **Next**.
 - d. Review all the details and click **Add System**.

Back up cloud-native SAP HANA databases

You can create a backup by assigning a pre-canned policy or the policy that you created.

Create a policy to protect SAP HANA database

You can create policies if you do not want to use or edit the pre-canned policies.

1. In the **Applications** page, from the Settings drop-down list, select **Policies**.
2. Click **Create policy**.
3. Specify a policy name.
4. (Optional) Edit the format of the Snapshot copy name.
5. Select policy type.
6. Specify the schedule and retention details.
7. (Optional) Specify the scripts. [Prescripts and postscripts](#).
8. Click **Create**.

Prescripts and postscripts

You can provide prescripts, postscripts, and exit scripts while creating a policy. These scripts are run on the HANA host during data protection operation.

The supported format for scripts are .sh, python script, perl script, and so on.

The prescript and the postscript should be registered by the host admin into `/opt/NetApp/snapcenter/scc/etc/allowed_commands.config` file.

```
[root@scspa2622265001 etc]# cat allowed_commands.config
command: mount
command: umount
command: /mnt/scripts/pre_script.sh
command: /mnt/scripts/post_script.sh
```

Environmental variables

For the backup workflow, the following environmental variables are available as part of prescript and postscript.

Environmental Variable	Description
SID	The System Identifier of the HANA Database chosen for restore
BackupName	Backup name chosen for restore operation
UserStoreKeyNames	Configured userstore key for the HANA database
OSDBUser	Configured OSDBUser for the HANA database
PolicyName	Only for scheduled backup
schedule_type	Only for scheduled backup

Create a backup of the SAP HANA Database

You can either assign a pre-canned policy or create a policy and then assign it to the database. Once the policy is assigned, the backups are created as per the schedule defined in the policy.

Before you begin

You should have added the SAP HANA database hosts.

[Add SAP HANA database hosts](#)

About this task

For HANA System Replication (HSR), the scheduled backup job triggers only for the primary HANA system and if the system fails over to the secondary HANA system, the existing schedules triggers a backup on the current primary HANA system. If the policy is not assigned to both the primary and secondary HANA system, after failover, the schedules will fail.

If different policies are assigned to the HSR systems, the scheduled backup triggers for both the primary and secondary HANA systems and the backup will fail for the secondary HANA system.

Steps

1. In the Applications page, if the database is not protected using any policy, click **Assign Policy**.

Though the database is protected using one or more policies, if needed, you can continue to assign more policies by clicking **...** > **Assign Policy**.

2. Select the policy and click **Assign**.

The backups are created as per the schedule defined in the policy.




The service account (*SnapCenter-account-`<account_id>`*) is used to run the scheduled backup operations.

Create on-demand backup of the SAP HANA database

After assigning the policy, you can create an on-demand backup of the application.

Steps

1. In the **Applications** page, click  corresponding to the application and click **On-Demand Backup**.
2. Select On-demand backup type.
3. For Policy Based backup, select the policy, retention tier and then click **Create Backup**.
4. For One time, select either Snapshot copy based, or File based perform the following steps:
 - a. Select the retention value and specify the backup name.
 - b. (Optional) Specify the scripts, and path for the scripts.

For more information, see [Prescripts and Postscripts](#)

- c. Click **Create Backup**.

Back up cloud-native SQL Server databases using REST APIs

Quick start

Get started quickly by following these steps.

1

Verify support for your configuration

- Operating System:
 - Windows 2016
 - Windows 2019
 - Windows 2022
- NetApp Cloud Storage: Amazon FSx for NetApp ONTAP
- Storage layouts: SAN (iSCSI)

NAS configuration is not supported.

- Database versions:
 - Microsoft SQL Server 2016
 - Microsoft SQL Server 2019
 - Microsoft SQL Server 2022
- Database configuration:
 - Standalone

2

Sign up to BlueXP

BlueXP is accessible from a web-based console. When you get started with BlueXP, your first step is to sign up

using your existing NetApp Support Site credentials or by creating a NetApp cloud login. For information, refer to [Sign up to BlueXP](#).

3

Log into BlueXP

After you sign up to BlueXP, you can log in from the web-based console. For information, refer to [Log into BlueXP](#).

4

Manage your BlueXP account

You can administer your account by managing users, service accounts, workspaces, and Connectors. For information, refer to [Manage your BlueXP account](#).

Configure FSx for ONTAP

Using BlueXP you should create an FSx for ONTAP working environment to add and manage volumes and additional data services. You should also create a Connector in AWS that enables BlueXP to manage resources and processes within your public cloud environment.

Create FSx for ONTAP working environment

You should create the FSx for ONTAP working environments where your databases are hosted. For information, refer to [Get started with Amazon FSx for ONTAP](#) and [Create and manage an Amazon FSx for ONTAP working environment](#).

You can create the FSx for ONTAP working environment either using BlueXP or AWS. If you have created using AWS, then you should discover the FSx for ONTAP systems in BlueXP.

Create a Connector

An Account Admin needs to create a Connector in AWS that enables BlueXP to manage resources and processes within your public cloud environment.

For information, refer to [Creating a Connector in AWS from BlueXP](#).

- You should use the same connector to manage both FSx for ONTAP working environment and databases.
- If you have the FSx for ONTAP working environment and databases in the same virtual private cloud (VPC), you can deploy the connector in the same VPC.
- If you have the FSx for ONTAP working environment and databases in different VPCs:
 - If you have NAS (NFS) workloads configured on FSx for ONTAP, then you can create the connector on either of the VPCs.
 - If you have only SAN workloads configured and not planning to use any NAS (NFS) workloads, then you should create the connector in the VPC where the FSx for ONTAP system is created.



For using NAS (NFS) workloads, you should have transit gateway between the database VPC and Amazon VPC. The NFS IP address which is a floating IP address can be accessed from another VPC only through transit gateway. We cannot access the floating IP addresses by peering the VPCs.

After creating the Connector, click **Storage > Canvas > My Working Environments > Add Working Environment** and follow the prompts to add the working environment.

Ensure that there is connectivity from the Connector to the Oracle database hosts and FSx working environment. The Connector should be able to connect to the cluster management IP address of the FSx working environment.

- Add the working environment by clicking **Storage > Canvas > My Working Environments > Add Working Environment**.

Ensure that there is connectivity from the connector to the database hosts and FSx for ONTAP working environment. The connector should connect to the cluster management IP address of the FSx for ONTAP working environment.

- Copy the Connector ID by clicking **Connector > Manage Connectors** and selecting the Connector name.

Install SnapCenter Plug-in for SQL Server and add database hosts

You should install the SnapCenter Plug-in for SQL Server on each of the SQL database hosts, add the database hosts, discover the database instances, and configure the credentials for the database instances.

Install the SnapCenter Plug-in for SQL Server

You should download plug-in **snapcenter_service_windows_host_plugin.exe** and then run the silent installer command to install the plug-in on the database host.

Before you begin

- You should ensure that the following prerequisites are met.
 - .Net 4.7.2 is installed
 - PowerShell 4.0 is installed
 - Minimum disk space of 5 GB is available
 - Minimum RAM size of 4 GB is available
- You should run the API to complete the customer on-boarding. For more information, refer to:
<https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/Tenant%20Registration/createTenant>

Steps

1. Download the plug-in by running the API from the Connector host.

```
docker exec -it cloudmanager_scs_cloud curl  
'http://127.0.0.1/api/v2/pluginpackage/windows'
```

The location of the file is `/var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/<agent_version>/sc-windows-host-plugin/snapcenter_service_windows_host_plugin.exe`.

2. Copy *snapcenter_service_windows_host_plugin.exe* from the connector to each of the MSSQL Server database hosts either using scp or other alternate methods.

3. Install the plug-in.

```
"C://<install_folder>/snapcenter_service_windows_host_plugin.exe"/silent/debuglog  
"C://<install_folder>/HA_Suite_Silent_Install_SCSQL_FRESH.log" /log"C://install_folder/"  
BI_SNAPCENTER_PORT=8145 ISFeatureInstall=SCSQL'
```

4. Copy the self signed certificate from `/var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/client/certificate/certificate.pem` to the MSSQL Server database hosts.

You can also generate a self signed certificate or a CA signed certificate if you do not use the default one.

5. Convert the certificate from .pem to .crt format in the Connector host.
'openssl x509 -outform der -in certificate.pem -out certificate.crt'
6. Double-click the certificate to add it to the **Personal** and **Trusted Root Certification Authorities** store.

Add the SQL Server database host

You should add the MSSQL database host using the host FQDN.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/v1/hosts'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/Host%20Management/AddHosts>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameters

Name	Type	Required
addr	string	True
connector_id	string	True
plugin_type	string	True
install_method	string	True
plugin_port	number	True
username	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

View the added SQL Server database hosts

You can run this API to view all the added SQL Server database hosts.

'GET snapcenter.cloudmanager.cloud.netapp.com/api/v1/hosts'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/Host%20Management/GetHosts>

Response

If the API is executed successfully, response code 200 is displayed.

Example:

```

{
  "num_records": 1,
  "total_records": 1,
  "records": [
    {
      "id": "85bd4603-08f7-45f4-ba8e-a0b1e2a0f4d0",
      "addr": "scspa2722211001.rtp.openenglab.netapp.com",
      "status": "Running",
      "connector_id": "fBf8Iwbp4BscBfD02qBwWm6I03gGAesRclients",
      "plugin_port": 8145,
      "plugins": [
        {
          "type": "mssql"
        }
      ],
      "os_type": "windows",
      "platform": "onprem",
      "username": "administrator",
      "operating_mode": "production"
    }
  ],
  "_links": {
    "next": {}
  }
}

```

Discover the database instances

You can run this API and enter the host ID to discover all the MSSQL instances.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/instances/discovery'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Instances/MSSQLInstancesDiscoveryRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameter

Name	Type	Required
host_id	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

View the discovered database instances

You can run this API to view all the discovered database instances.

'GET snapcenter.cloudmanager.cloud.netapp.com/api/mssql/instances'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/MSSQL%20Instances/GetMSSQLInstancesRequest>

Response

If the API is executed successfully, response code 200 is displayed.

Example:

```

{
  "num_records": 2,
  "total_records": 2,
  "records": [
    {
      "id": "953e66de-10d9-4fd9-bdf2-bf4b0eaabfd7",
      "name": "scspa2722211001\\NAMEDINSTANCE1",
      "host_id": "85bd4603-08f7-45f4-ba8e-a0b1e2a0f4d0",
      "status": "Running",
      "auth_mode": 0,
      "version": "",
      "is_clustered": false,
      "is_credentials_configured": false,
      "protection_mode": ""
    },
    {
      "id": "18e1b586-4c89-45bd-99c8-26268def787c",
      "name": "scspa2722211001",
      "host_id": "85bd4603-08f7-45f4-ba8e-a0b1e2a0f4d0",
      "status": "Stopped",
      "auth_mode": 0,
      "version": "",
      "is_clustered": false,
      "is_credentials_configured": false,
      "protection_mode": ""
    }
  ],
  "_links": {
    "next": {}
  }
}

```

Configure the database instance credentials

You can run this API to validate and set credentials for the database instances.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql//api/mssql/credentials-configuration'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/MSSQL%20Instances/ConfigureCredentialRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameter

Name	Type	Required
host_id	string	True
instance_ids	string	True
username	string	True
password	string	True
auth_mode	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

Back up cloud-native Microsoft SQL Server databases

You can create scheduled or on-demand backups by assigning the policies that you created.

Create backup policy

You can run this API to create the backup policy.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/backup/policies'

For more information, refer to: https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Backup%20Policies/MSSQLBackupPolicyService_CreateMSSQLBackupPolicy

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameters

Name	Type	Required
name	string	True
backup_type	string	True
copy_only_backup	string	False
is_system_defined	string	False
backup_name_format	string	True
schedule_type	string	True
start_time	number	True
hours_interval	number	True
minutes_interval	number	True
retention_type	string	True
retention_count	number	True
end_time	number	True

Response

If the API is executed successfully, response code 201 is displayed.

Example:

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}
```

Assign policy to SQL database instance

You can run this API to assign policy to SQL database instance.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/instances/{id}/policy-assignment'

Where, *id* is MSSQL instance ID obtained by running the discover database instance API. For more information, refer to [Discover the database instances](#).

Array of IDs is the input here. For example:

```
[
  "c9f3e68d-1f9c-44dc-b9af-72a9dfc54320"
]
```

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Policy%20Assignment/PostMSSQLInstanceAssignPolicyRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

Create an on-demand backup

You can run this API to create an on-demand backup.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/backups'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Backups/CreateMSSQLBackupRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameters

Name	Type	Required
id	string	True
 This is ID of the MSSQL database instance.		
resource_type	string	True
policy_id	string	True
schedule_type	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

View the backups

You can run these APIs to list view all the backups and also to view details of a particular backup.

'GET snapcenter.cloudmanager.cloud.netapp.com/api/mssql/backups'

'GET snapcenter.cloudmanager.cloud.netapp.com/api/mssql/backups/{id}'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Backups/MSSQLGetBackupsRequest>

Response

If the API is executed successfully, response code 200 is displayed.

Example:

```
{
  "total_records": 1,
  "num_records": 1,
  "records": [
    {
      "backup_id": "602d7796-8074-43fc-a178-eee8c78566ac",
      "resource_id": "a779578d-cf78-46f3-923d-b9223255938c",
      "backup_name":
"Hourly_policy2_scspa2722211001_NAMEDINSTANCE1_2023_08_08_07_02_01_81269_0",
      "policy_name": "policy2",
      "schedule_type": "Hourly",
      "start_time": "2023-08-08T07:02:10.203Z",
      "end_time": "0001-01-01T00:00:00Z",
      "backup_status": "success",
      "backup_type": "FullBackup"
    }
  ],
  "_links": {
    "next": {}
  }
}
```

Restore cloud-native Oracle databases

Restore cloud-native Oracle databases to original location


In the event of data loss, you can restore the data files, control files, or both to original location and then recover the database.

Before you begin

If the Oracle 21c database is in STARTED state, the restore operation fails. You should run the following command to restore the database successfully.

```
cp -f <ORACLE_HOME>/jdbc/lib/ojdbc8.jar
/opt/NetApp/snapcenter/spl/plugins/sco/lib/ojdbc8-8.jar
```

Steps

1. Click  corresponding to the database that you want to restore and click **Restore**.
2. Select the restore point to which the database should be restored and click **Restore to original location**.
3. In the Restore Scope section, perform the following actions:

If you...	Do this...
Want to restore only the data files	Select All Data Files .

If you...	Do this...
Want to restore only the control files	Select Control Files
Want to restore both data files and control files	Select All Data Files and Control Files .

You can also select **Force in-place restore** checkbox.

In Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP SAN layout, if SnapCenter Plug-in for Oracle finds any foreign files other than Oracle data files on the ASM diskgroup, connect and copy restore method is performed. The foreign files could be one or more of the following types:

- Parameter
- Password
- archive log
- online log
- ASM parameter file.

The **Force in-place restore** option overrides the foreign files of type parameter, password, and archive log. You should use the latest backup when **Force in-place restore** option is selected.

4. In the Recovery Scope section, perform the following actions:

If you...	Do this...
Want to recover to the last transaction	Select All Logs .
Want to recover to a specific System Change Number (SCN)	Select Until SCN and specify the SCN.
Want to recover to a specific date and time	Select Date and Time .
Do not want to recover	Select No recovery .

For the selected recovery scope, in the **Archive Log Files Locations** field you can optionally specify the location that contains the archive logs required for recovery.

Select the check box if you want to open the database in READ-WRITE mode after recovery.

5. Click **Next** and review the details.

6. Click **Restore**.

Restore cloud-native Oracle databases to alternate location

In the event of data loss, you can restore the Oracle database to alternate location only on Azure NetApp Files. The alternate location can be on a different host or on the same host.

Before you begin

- If the Oracle 21c database is in STARTED state, the restore operation fails. You should run the following command to restore the database successfully.

```
cp -f <ORACLE_HOME>/jdbc/lib/ojdbc8.jar  
/opt/NetApp/snapcenter/spl/plugins/sco/lib/ojdbc8-8.jar
```
- You should ensure that the Oracle version on the alternate host is same as that of the original host.


About this task

While initiating the restore operation, you are not allowed to modify the configurations except the Oracle home, maximum volume throughput, Oracle SID, and database credentials.

Full recovery is enabled by default with *Until cancel* set to true.

Archive log mode is turned off by default for the restored database. You can enable archive log mode and keep the archive logs on the NetApp volume if required.

Steps

1. Click  corresponding to the database that you want to restore and click **Restore**.
2. Select the restore point to which the database should be restored and click **Restore to alternate location** > **Next**.
3. In the Configuration page, specify the details of the alternate location, SID, Oracle_Home, database credentials, and storage throughput.

For the database credential, if the OS User authentication is disabled, you should provide a password for the sys user to connect to the restored database on the same or target host.

4. Click **Next**, review the details and click **Restore**.

The progress of the restore operation can be viewed in the Job Monitor page. After the job is completed, click **Refresh Discovery** to view the restored database. However, you cannot protect the database that is restored to alternate location.

Restore cloud-native SAP HANA databases

In the event of data loss, you can restore the data and non-data files and then recover the database.

Before you begin

- The SAP HANA system must be in a stopped state.
- If the SAP HANA system is up and running, you can provide a prescript to stop the system.

About this task

- If you enable the ANF backups on a volume, Single File SnapRestore operation is performed.
- For Non-Data Volumes and Global Non-Data Volumes, connect and copy restore operation is performed.
 - The Quality of Service (QoS) values for connect and copy restore operation are picked up from the source volumes of Non-Data Volumes or Global Non-Data Volumes.



QoS is applicable only for capacity pools of type "Manual".

Steps

1. Click [...](#) corresponding to the database that you want to restore and click **View Details**.
2. Click [...](#) corresponding to the data backup that you want to restore and click **Restore**.
3. In the **Restore System** page, enter the scripts. [Prescripts and postscripts](#).

For the restore workflow, the following environmental variables are available as part of prescript and postscript.

Environmental Variable	Description
SID	The System Identifier of the HANA Database chosen for restore
BackupName	Backup name chosen for restore operation
UserStoreKeyNames	Configured userstore key for the HANA database
OSDBUser	Configured OSDBUser for the HANA database

4. Click **Restore**.

What's next

After restoring, manually recover the SAP HANA system or provide a postscript, which performs the SAP HANA system recovery.

Restore Non-Data Volume

About this task

For connect and copy restore operation, go to Microsoft Azure portal, select the volume, click **Edit**, and enable **Hide snapshot path**.

Steps


1. In the **Applications** page, select Non-Data Volume from the drop-down box.
2. Click [...](#) corresponding to the backup that you want to restore and click **Restore**.

Restore Global Non-Data Volume

About this task

For connect and copy restore operation, go to Microsoft Azure portal, select the volume, click **Edit**, and enable **Hide snapshot path**.

Steps

1. In the **Applications** page, click on the Global Non-Data Volume that you want to restore.
2. Click  corresponding to the Global Non-Data Volume that you want to restore and click **Restore**.

Restore Microsoft SQL Server database

You can restore Microsoft SQL Server database to the same host. You should first get list of databases and then restore the database.

View the list of databases

You can run this API to view the list of databases.

'GET snapcenter.cloudmanager.cloud.netapp.com/api/mssql/databases'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Databases/GetMSSQLDatabasesRequest>

Response

If the API is executed successfully, response code 200 is displayed.

Example:

```

{
  "num_records": 3,
  "total_records": 3,
  "records": [
    {
      "id": "348901e5-aeaa-419f-88b1-80240de3b1fe",
      "name": "DB4",
      "hostname": "scspa2722211001.rtp.openenglab.netapp.com",
      "size": 0.078125,
      "instance_id": "454c8413-5351-41fc-88bf-f20fb050ec87",
      "instance": "scspa2722211001\\\\NAMEDINSTANCE1",
      "db_status": "Normal",
      "db_access": "eUndefined",
      "db_type": "User",
      "recovery_mode": "Full"
    },
    {
      "id": "c79d33ab-7322-4ed6-92f5-51ad7a6944e0",
      "name": "DB5",
      "hostname": "scspa2722211001.rtp.openenglab.netapp.com",
      "size": 0.078125,
      "instance_id": "454c8413-5351-41fc-88bf-f20fb050ec87",
      "instance": "scspa2722211001\\\\NAMEDINSTANCE1",
      "db_status": "Normal",
      "db_access": "eUndefined",
      "db_type": "User",
      "recovery_mode": "Full"
    },
    {
      "id": "40d6f35a-f4fb-48bc-8e0a-0ac93ddf0888",
      "name": "model",
      "hostname": "scspa2722211001.rtp.openenglab.netapp.com",
      "size": 0.015625,
      "instance_id": "454c8413-5351-41fc-88bf-f20fb050ec87",
      "instance": "scspa2722211001\\\\NAMEDINSTANCE1",
      "db_status": "Normal",
      "db_access": "eUndefined",
      "db_type": "System",
      "recovery_mode": "Full"
    }
  ],
  "_links": {
    "next": {}
  }
}

```


Restore and recover the MSSQL database

You can run this API to restore the MSSQL database.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/databases/{id}/restore'

Where, *id* is MSSQL database ID obtained by running the view database API. For more information, refer to [View the list of databases](#).

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/MSSQL%20Database%20Restore/RestoreMSSQLDatabaseRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameters

Name	Type	Required
backup_id	string	True
overwrite_database	bool	True
retain_replication_settings	bool	False
recovery_mode	string The 3 supported strings are <i>Operational</i> , <i>NonOperational</i> , and <i>ReadOnly</i> .	True
undo_file_directory	string	True
restore_type	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

Clone cloud-native Oracle databases

Clone concepts and requirements

You can clone an Oracle database residing on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP using the backup of the database either to the source database host or to an alternate host. You can clone the backup from primary storage systems.

Before cloning the database, you should understand the clone concepts and ensure that all the requirements are met.

Requirements for cloning an Oracle database

Before cloning an Oracle database, you should ensure that prerequisites are completed.

- You should have created a backup of the database.
You should have successfully created online data and log backup for the cloning operation to succeed.
- In the `asm_diskstring` parameter, you should configure:
 - `AFD:*` if you are using ASMFD
 - `ORCL:*` if you are using ASMLIB
 - `/dev/<exact_device_location>` if you are using ASMUDEV
- If you are creating the clone on an alternate host, the alternate host should meet the following requirements:
 - The plug-in should be installed on the alternate host.
 - Oracle software should be installed on the alternate host.
 - The clone host should be able to discover LUNs from storage if you are cloning a database residing on iSCSI SAN storage.
If you are cloning to an alternate host, then make sure that an iSCSI session is established between the storage and the alternate host.
 - If the source database is an ASM database:
 - The ASM instance should be up and running on the host where the clone will be performed.
 - The ASM diskgroup should be provisioned prior to the clone operation if you want to place archive log files of the cloned database in a dedicated ASM diskgroup.

- The name of the data diskgroup can be configured but ensure that the name is not used by any other ASM diskgroup on the host where the clone will be performed.
- Data files residing on the ASM diskgroup are provisioned as part of the clone workflow.

Limitations

- Cloning of databases residing on Azure NetApp Files is not supported.
- Cloning of databases residing on Qtree is not supported.
- Backing up a cloned database is not supported.
- If daily automatic backups are enabled on Amazon FSx for NetApp ONTAP, the cloned volumes on Amazon FSx for NetApp ONTAP cannot be deleted from BlueXP UI because FSx would have created backups on the cloned volumes.
You should delete the cloned volumes after deleting all the backups for the volume from FSx UI and then delete the clones from the BlueXP UI using force option.

Clone methods

You can create clone either using the basic method or using the clone specification file.

Clone using basic method

You can create the clone with the default configurations based on the source database and the selected backup.

- The database parameters, home, and OS user are defaulted to the source database.
- The data file paths are named based on the naming scheme selected.
- The pre-script, post-script, and SQL statements cannot be specified.
- The recovery option is by default **until cancel** and it uses the log backup associated with the data backup for recovery

Clone using specification file

You can define the configurations in the clone specification file and use it to clone the database. You can download the specification file, modify it to your requirement, and then upload the file. [Learn more](#).

The different parameters defined in the specification file and that can be modified are as follows:

Parameter	Description
control_files	<p>Location of control files for the clone database.</p> <p>The number of control files will be same as source database.</p> <p>If you want to override the control file path, you can provide a different control file path. The file system or the ASM diskgroup should exist on the host.</p>

Parameter	Description
redo_logs	<p>Location, size, redo group number of redo logs.</p> <p>A minimum of two redo log groups are required to clone the database. If you want to override the redo log file path, you can customize the redo log file path to a different file system than that of the source database. The file system or the ASM diskgroup should exist on the host.</p>
oracle_version	Version of Oracle on the target host.
oracle_home	Oracle home on the target host.
enable_archive_log_mode	Controls the archive log mode for the clone database
database_parameters	Database parameters for the cloned database
sql_statements	The SQL statements to be executed on the database after cloning
os_user_detail	Oracle OS user on the target clone database
database_port	Port used for communicating with the database if OS authentication is disabled on the host.
asm_port	Port used for communicating with ASM database if credentials are provided in the create clone input.
skip_recovery	Does not perform recovery operation.
until_scn	Recovers the database up to the specified system change number (scn).
until_time	<p>Recovers the database up to the specified date and time.</p> <p>The accepted format is <i>mm/dd/yyyy hh:mm:ss</i>.</p>
until_cancel	<p>Recovers by mounting the log backup associated with the data backup that was selected for cloning.</p> <p>The cloned database is recovered till the missing or corrupt log file.</p>
log_paths	Additional locations of archive log paths to be used for recovering the cloned database.

Parameter	Description
source_location	Location of the diskgroup or mount point on the source database host.
clone_location	Location of the diskgroup or mount point that needs to be created on the target host corresponding to the source location.
location_type	It can be either ASM_Diskgroup Or mountpoint. The values are auto-populated at the time of downloading the file. You should not edit this parameter.
pre_script	Script to be executed on the target host before creating the clone.
post_script	Script to be executed on the target host after creating the clone.
path	Absolute Path of the script on the clone host. You should store the script either in /var/opt/snapcenter/spl/scripts or in any folder inside this path.
timeout	The timeout time specified for the script running on the target host.
arguments	Arguments specified for the scripts.

Clone naming scheme

Clone naming scheme defines what will be the location of the mount points and name of the diskgroups of the cloned database. You can either select **Identical** or **Auto-generated**.

Identical naming scheme

If you select the clone naming scheme as **Identical**, the location of mount points and the name of the diskgroups of the cloned database will be same as the source database.

For example, if the mount point of the source database is `/netapp_sourcedb/data_1 , +DATA1_DG`, for the cloned database the mount point remains the same for both NFS and ASM on SAN.

- Configurations like number and path of control files and redo files will be same as source.



If the redo logs or control file paths are located on the non-data volumes, then the user should have provisioned the ASM diskgroup or mountpoint in the target host.

- Oracle OS user and Oracle version will be same as source database.
- Clone storage volume name will be in the following format `sourceVolNameSCS_Clone_CurrentTimeStampNumber`.

For example, if the volume name on the source database is *sourceVolName*, the cloned volume name will be *sourceVolNameSCS_Clone_1661420020304608825*.



The *CurrentTimeStampNumber* provides the uniqueness in volume name.

Auto-generated naming scheme

If you select the cloning scheme as **Auto-generated**, the location of mount points and the name of the diskgroups of the cloned database will be appended with a suffix.

- If you have selected the basic clone method, the suffixed will be the **Clone SID**.
- If you have selected the specification file method, the suffix will be the **Suffix** that was specified while downloading the clone specification file.

For example, if the mount point of the source database is */netapp_sourcedb/data_1* and the **Clone SID** or the **Suffix** is *HR*, then the mount point of the cloned database will be */netapp_sourcedb/data_1_HR*.

- Number of control files and redo log files will be same as the source.
- All redo log files and control files will be located on one of the cloned data mount points or data ASM diskgroups.
- Clone storage volume name will be in the following format `sourceVolNameSCS_Clone_CurrentTimeStampNumber`.

For example, if the volume name on the source database is *sourceVolName*, the cloned volume name will be *sourceVolNameSCS_Clone_1661420020304608825*.



The *CurrentTimeStampNumber* provides the uniqueness in volume name.

- The format of the NAS mount point will be *SourceNASMountPoint_suffix*.
- The format of the ASM diskgroup will be *SourceDiskgroup_suffix*.



If the number of characters in the clone diskgroup is greater than 25 then it will have *SC_HashCode_suffix*.

Database parameters

The value of the following database parameters will be same as that of the source database irrespective of the clone naming scheme.

- `log_archive_format`
- `audit_trail`
- `processes`
- `pga_aggregate_target`
- `remote_login_passwordfile`

- undo_tablespace
- open_cursors
- sga_target
- db_block_size

The value of the following database parameters will be appended with a suffix based on the clone SID.

- audit_file_dest = {sourcedatabase_parametervalue}_suffix
- log_archive_dest_1 = {sourcedatabase_oraclehome}_suffix

Supported predefined environment variables for clone specific prescript and postscript

You can use the supported predefined environment variables when you execute the prescript and postscript while cloning a database.

- SC_ORIGINAL_SID specifies the SID of the source database.
This parameter will be populated for application volumes. Example: NFSB32
- SC_ORIGINAL_HOST specifies the name of the source host.
This parameter will be populated for application volumes. Example: asmrac1.gdl.englab.netapp.com
- SC_ORACLE_HOME specifies the path of the target database's Oracle home directory.
Example: /ora01/app/oracle/product/18.1.0/db_1
- SC_BACKUP_NAME specifies the name of the backup.
This parameter will be populated for application volumes. Examples:
 - If the database is not running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
 - If the database is running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1, RG2_scspr2417819002_07-21-2021_12.16.48.9267_1, RG2_scspr2417819002_07-22-2021_12.16.48.9267_1
- SC_ORIGINAL_OS_USER specifies the operating system owner of the source database.
Example: oracle
- SC_ORIGINAL_OS_GROUP specifies the operating system group of the source database.
Example: oinstall
- SC_TARGET_SID specifies the SID of the cloned database.
For PDB clone workflow, the value of this parameter will not be predefined. This parameter will be populated for application volumes.
Example: clonedb
- SC_TARGET_HOST specifies the name of the host where the database will be cloned.
This parameter will be populated for application volumes. Example: asmrac1.gdl.englab.netapp.com
- SC_TARGET_OS_USER specifies the operating system owner of the cloned database.
For PDB clone workflow, the value of this parameter will not be predefined. Example: oracle
- SC_TARGET_OS_GROUP specifies the operating system group of the cloned database.
For PDB clone workflow, the value of this parameter will not be predefined. Example: oinstall
- SC_TARGET_DB_PORT specifies the database port of the cloned database.
For PDB clone workflow, the value of this parameter will not be predefined. Example: 1521

Supported delimiters

- @ is used to separate data from its database name and to separate the value from its key.
Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
- | is used to separate the data between two different entities for SC_BACKUP_NAME parameter.
Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
- , is used to separate set of variables for the same key.
Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1, RG2_scspr2417819002_07-21-2021_12.16.48.9267_1, RG2_scspr2417819002_07-22-2021_12.16.48.9267_1

Clone cloud-native Oracle databases

You can clone an Oracle database residing on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP using the backup of the database either to the source database host or to an alternate host.



You might clone databases for the following reasons:


- To test functionality that must be implemented using the current database structure and content during application development cycles.
- To populate data warehouses using data extraction and manipulation tools.
- To recover data that was mistakenly deleted or changed.


Before you begin

You should understand the clone concepts and ensure that all the requirements are met. [Learn more](#).

Steps

1. Click  corresponding to the database that you want to clone and click **View Details**.
2. Click  corresponding to the data backup and click **Clone**.
3. In the Clone Details page, select one of the clone options.
4. Depending on the option selected, perform the following actions:

If you have selected...	Do this...
Basic	<ol style="list-style-type: none"> 1. Select the clone host. If you want to create the clone on an alternate host, select the host having the same version of Oracle and OS as that of the source database host. 2. Specify the SID of the clone. 3. Select the clone naming scheme. If the database is cloned to the source host, the clone naming scheme will be auto-generated. If the database is cloned to an alternate host, clone naming scheme will be identical. 4. Specify the Oracle home path. 5. (Optional) Specify the database credentials. <ul style="list-style-type: none"> ◦ Database credential: If the OS User authentication is disabled, you should provide a password for the sys user to connect to the cloned database on the same or target host. ◦ ASM credential: If the OS user authentication is disabled on the target host, you should provide a sysasm privileged user's credentials to connect to the ASM instance on the target host. <div data-bbox="966 1199 1023 1255"></div> <div data-bbox="1073 1178 1446 1276">Ensure that the listener is up and running on the target host.</div> 6. Click Next. 7. Click Clone.

If you have selected...	Do this...
Specification file	<ol style="list-style-type: none"> 1. Click Download File to download the specification file. 2. Select the clone naming scheme. If you select, Auto-generated, you should specify the suffix. 3. Edit the specification file as per the requirement and upload it by clicking the Browse button. 4. Select the clone host. If you want to create the clone on an alternate host, select the host having the same version of Oracle and OS as that of the source database host. 5. Specify the SID of the clone. 6. (Optional) Specify the database credentials. <ul style="list-style-type: none"> ◦ Database credential: If the OS User authentication is disabled, you should provide a password for the sys user to connect to the cloned database on the same or target host. ◦ ASM credential: If the OS user authentication is disabled on the target host, you should provide a sysasm privileged user's credentials to connect to the ASM instance on the target host. <div>  <p>Ensure that the listener is up and running on the target host.</p> </div> 7. Click Next. 8. Click Clone.

5. Click  adjacent to **Filter By** and select **Clone options > Clones** to view the clones.

Refresh SAP HANA target system

You can perform a refresh of a SAP HANA target system with the data of a SAP HANA source system. This can be used to provide the current production data into a test system. BlueXP backup and recovery allows you to select a Snapshot copy from a source system and creates a new Azure NetApp Files volume based on the Snapshot copy. Example scripts are available, which executes the required operations on the

database host to recover the SAP HANA database.

Before you begin

- You should install the SAP HANA target system before you execute the first refresh operation.
- You should add the source and target HANA systems manually into BlueXP backup and recovery.
- Ensure that the SAP HANA database version is same on source and the target system.
- You should have decided on which refresh scripts to be used. The refresh scripts are available in the solution technical report.

Automation example scripts

You can customize the refresh scripts.

- The following environmental variables are available as part of the prescript and postscript:
 - CLONED_VOLUMES_MOUNT_PATH
 - <SOURCEVOLUME>_DESTINATION
 - HANA_DATABASE_TYPE
 - TENANT_DATABASE_NAMES
- You must upgrade the plug-in to 3.0 version.
- The mount paths should be the same for the data volume on both the source and the target SAP HANA systems.
- Before the first refresh operation, ensure that the '/etc/fstab' file does not have entries for the data volumes of the target SAP HANA system.

About this task

- System refresh is supported only for multi-tenant database container HANA system.
- The existing policies will be valid after the system refresh.
- The new volumes created will have the following naming convention: <sourcevolumename>-<timestamp>
 - Timestamp format: <year><month><day>-<hour><minute><second>

For example, if the source volume is vol1, the refreshed volume name will be vol1-20230109-184501




The new volume will be placed in the same capacity pool as that of the target volumes.

- The junction path will be the same as the volume name.
- The “max throughput number” for the new volume is picked from the volume of the target system with manual Quality of Service (QoS) capacity pools.
For auto QoS capacity pools the throughput is defined by the capacity of the source volume.
- During system refresh, the auto mount and unmount of the volumes are performed using workflows instead of scripts.

Steps

1. In the BlueXP UI, click **Protection > Backup and recovery > Applications**.

2. In the **Applications** page, click  icon to select the action corresponding to the system that you want to refresh and select **System Refresh**.
3. In the **System Refresh** page, perform the following actions:
 - a. Select source system and Snapshot copy.
 - b. (Optional) Enter Export addresses from which the new volumes can be accessed.
 - c. (Optional) Enter Maximum storage throughput (MIBs).
 - d. Enter prescript, postscript, and on failure script paths.
On failure script is executed only when the system refresh operation fails.
 - e. Click **Refresh**.

Manage protection of cloud-native application data

Monitor jobs

You can monitor the status of the jobs that have been initiated in your working environments. This allows you to see the jobs that have completed successfully, those currently in progress, and those that have failed so you can diagnose and fix any problems.



The scheduled jobs will be listed in the BlueXP Job monitor page after a delay of 5 minutes (maximum) from the job completion time.

For more information, refer to [Monitor job status](#).

Maintenance of Oracle database hosts

An admin can manually put the database hosts in maintenance mode to perform maintenance tasks on the hosts. During upgrade, the hosts are automatically put to maintenance mode and after upgrade, the hosts are automatically switched to production mode.


When the hosts are put in maintenance mode, the on-demand operations fail and the scheduled jobs are skipped.




You cannot verify if any jobs are running for the resources on the hosts before putting the hosts in maintenance mode.

Steps

1. In the BlueXP UI, click **Protection > Backup and recovery > Applications**
2. Select **Oracle** as the application type.
3. Click **Settings > Hosts**.
4. Perform one of the following actions:

If you...	Do this...
Want to put the host in maintenance mode	Click  corresponding to the host and select Enable maintenance mode .

If you...	Do this...
Want to bring the host out of maintenance mode	Click  corresponding to the host that is under maintenance and select Disable maintenance mode .

Audit data


When you either run an API directly or use the UI to make the API call to any of the externally exposed APIs of the BlueXP backup and recovery for applications, the request details such as headers, role, request body, and API information will be logged in the BlueXP timeline and the audit entries are retained in the timeline forever. The status and error response of the API call are also audited post operation completion. In the case of asynchronous API responses like jobs, the job id also gets logged as part of response.

BlueXP backup and recovery for applications log the entries such as host IP, request body, operation name, who triggered, some headers, and the operation state of the API.

View backup details

You can view total number of backups created, policies used for creating backups, database version, and agent ID.

Steps

1. Click **Backup and recovery > Applications**.
2. Click  corresponding to the application and click **View Details**.







The agent ID is associated to the Connector. If a Connector that was used during registering the SAP HANA host no longer exists, the subsequent backups of that application will fail because the agent ID of the new Connector is different. You should modify the Connector id in the host. For information, see [Update the Connector Details](#).

Delete clone

You can delete a clone if you no longer require.

Steps

1. Click  adjacent to **Filter By** and select **Clone options > Clone parents**.
2. Click  corresponding to the application and click **View Details**.
3. In the Database Details page, click  adjacent to **Filter By** and select **Clone**.
4. Click  corresponding to the clone that you want to delete and click **Delete**.
5. (Optional) Select the **force delete** checkbox.

Update the Connector Details

You should deploy a new Connector, if the Connector that was used during registering the application host no longer exists or is corrupted. After deploying the new connector, you should run the **connector-update** API to update the Connector details for all hosts registered using the old connector.

After updating the Connector details for Oracle or SAP HANA hosts, perform the following to ensure that the Connector details were updated successfully.

Steps

- 1. Login into BlueXP Connector VM and perform the following steps:
 - a. Validate that the plug-in is reachable from the Connector by running the below command from the Connector.
`docker exec -it cloudmanager_scs_cloud curl -ik https://<FQDN or IP of the plug-in host>:<plug-in port>/getVersion --cert/config/client/certificate/certificate.pem --key/config/client/certificate/key.pem`
 - b. Obtain the base mount path.
`sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs sudo docker volume inspect | grep Mountpoint`
 - c. Copy certificate.pem from <base_mount_path>/client/certificate/ path of the Connector VM to /var/opt/snapcenter/spl/etc/ on the plug-in host.
- 2. Log in to the plug-in host and perform the following steps:
 - a. Navigate to /var/opt/snapcenter/spl/etc and run the keytool command to import the certificate.pem file.
`keytool -import -alias agentcert -file certificate.pem -keystore keystore.jks -deststorepass snapcenter -noprompt`
 - b. Restart SPL: `systemctl restart spl`
 - c. Perform one of the following:

If you are on...	Do this...
Oracle database host	<ul style="list-style-type: none">1. Ensure that all the prerequisites are met.2. Click Backup and recovery > Applications3. Click ... corresponding to the application and click View Details.4. Modify Connector ID.

If you are on...	Do this...
SAP HANA database host	<ol style="list-style-type: none"> 1. Ensure that all the prerequisites are met. 2. Run the following command: <div data-bbox="885 304 1469 945"> <pre>curl --location --request PATCH 'https://snapcenter.cloudmanager .cloud.netapp.com/api/saphana/ho sts/connector/update' \ --header 'x-account-id: <CM account-id>' \ --header 'Authorization: Bearer token' \ --header 'Content-Type: application/json' \ --data-raw '{ "old_connector_id": "Old connector id that no longer exists", "new_connector_id": "New connector Id"}</pre> </div> <p>Connector details will get updated successfully if all the hosts have SnapCenter Plug-in for SAP HANA service installed and running and also if they are all reachable from the new Connector.</p>

Configure CA signed certificate

You can configure CA signed certificate if you want to include additional security to your environment.

Configure CA signed certificate for BlueXP Connector

The connector uses a self-signed certificate to communicate with plug-in. The self-signed certificate is imported to the keystore by the installation script. You can perform the following steps to replace the self-signed certificate with CA signed certificate.

Steps

1. Perform the following steps on the Connector to use the CA certificate as the client certificate when the Connector is connecting with the plug-in.
 - a. Login to Connector.
 - b. Run the following command to get the `<base_mount_path>`:

```
sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs
sudo docker volume inspect | grep Mountpoint
```
 - c. Delete all the existing files located at `<base_mount_path>/client/certificate` in the Connector.

- d. Copy the CA signed certificate and key file to the `<base_mount_path>/client/certificate` in the Connector.

The file name should be `certificate.pem` and `key.pem`. The `certificate.pem` should have the entire chain of the certificates like intermediate CA and root CA.

- e. Create the PKCS12 format of the certificate with the name `certificate.p12` and keep at `<base_mount_path>/client/certificate`.

Example: `openssl pkcs12 -inkey key.pem -in certificate.pem -export -out certificate.p12`

2. Perform the following steps on the plug-in host to validate the certificate sent by the Connector.

- a. Log in to the plug-in host.
- b. Copy the `certificate.pem` and certificates for all the intermediate CA and root CA from the Connector to the plug-in host at `/var/opt/snapcenter/spl/etc/`.



The format of the Intermediate CA and root CA certificate should be in `.crt` format.

- c. Navigate to `/var/opt/snapcenter/spl/etc` and run the `keytool` command to import the `certificate.pem` file.
`keytool -import -alias agentcert -file certificate.pem -keystore keystore.jks -deststorepass snapcenter -noprompt`
- d. Import the root CA and intermediate certificates.
`keytool -import -trustcacerts -keystore keystore.jks -storepass snapcenter -alias trustedca -file <certificate.crt>`



The `certificate.crt` refers to the certificates of root CA as well as intermediate CA.

- e. Restart SPL: `systemctl restart spl`

Configure CA signed certificate for the plug-in

The CA certificate should have the same name as registered in Cloud Backup for the plug-in host.

Steps

1. Perform the following steps on the plug-in host to host the plug-in using the CA certificate.
 - a. Navigate to the folder containing the SPL's keystore `/var/opt/snapcenter/spl/etc`.
 - b. Create the PKCS12 format of the certificate having both certificate and key with alias `splkeystore`.

The `certificate.pem` should have the entire chain of the certificates like intermediate CA and root CA.

Example: `openssl pkcs12 -inkey key.pem -in certificate.pem -export -out certificate.p12 -name splkeystore`

- c. Add the CA certificate created in the above step.
`keytool -importkeystore -srckeystore certificate.p12 -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS -srcalias splkeystore -destalias splkeystore -noprompt`
- d. Verify the certificates.
`keytool -list -v -keystore keystore.jks`

- e. Restart SPL: `systemctl restart spl`
2. Perform the following steps on the Connector so that the Connector can verify the plug-in's certificate.
 - a. Log in to the Connector as non-root user.
 - b. Run the following command to get the `<base_mount_path>`:

```
sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs
sudo docker volume inspect | grep Mountpoint
```
 - c. Copy the the root CA and intermediate CA files under the server directory.

```
cd <base_mount_path>
mkdir server
```

The CA files should be in pem format.
 - d. Connect to the `cloudmanager_scs_cloud` and modify the **enableCACert** in `config.yml` to **true**.

```
sudo docker exec -t cloudmanager_scs_cloud sed -i 's/enableCACert:
false/enableCACert: true/g' /opt/netapp/cloudmanager-scs-
cloud/config/config.yml
```
 - e. Restart `cloudmanager_scs_cloud` container.

```
sudo docker restart cloudmanager_scs_cloud
```

Access REST APIs

The REST APIs to protect the applications to cloud is available at:
<https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/>.

You should obtain the user token with federated authentication to access the REST APIs. For information to obtain the user token, refer to [Create a user token with federated authentication](#).

Back up and restore virtual machines data

Protect your virtual machines data

BlueXP backup and recovery for virtual machines provides data protection capabilities by backing up datastores and restoring virtual machines.

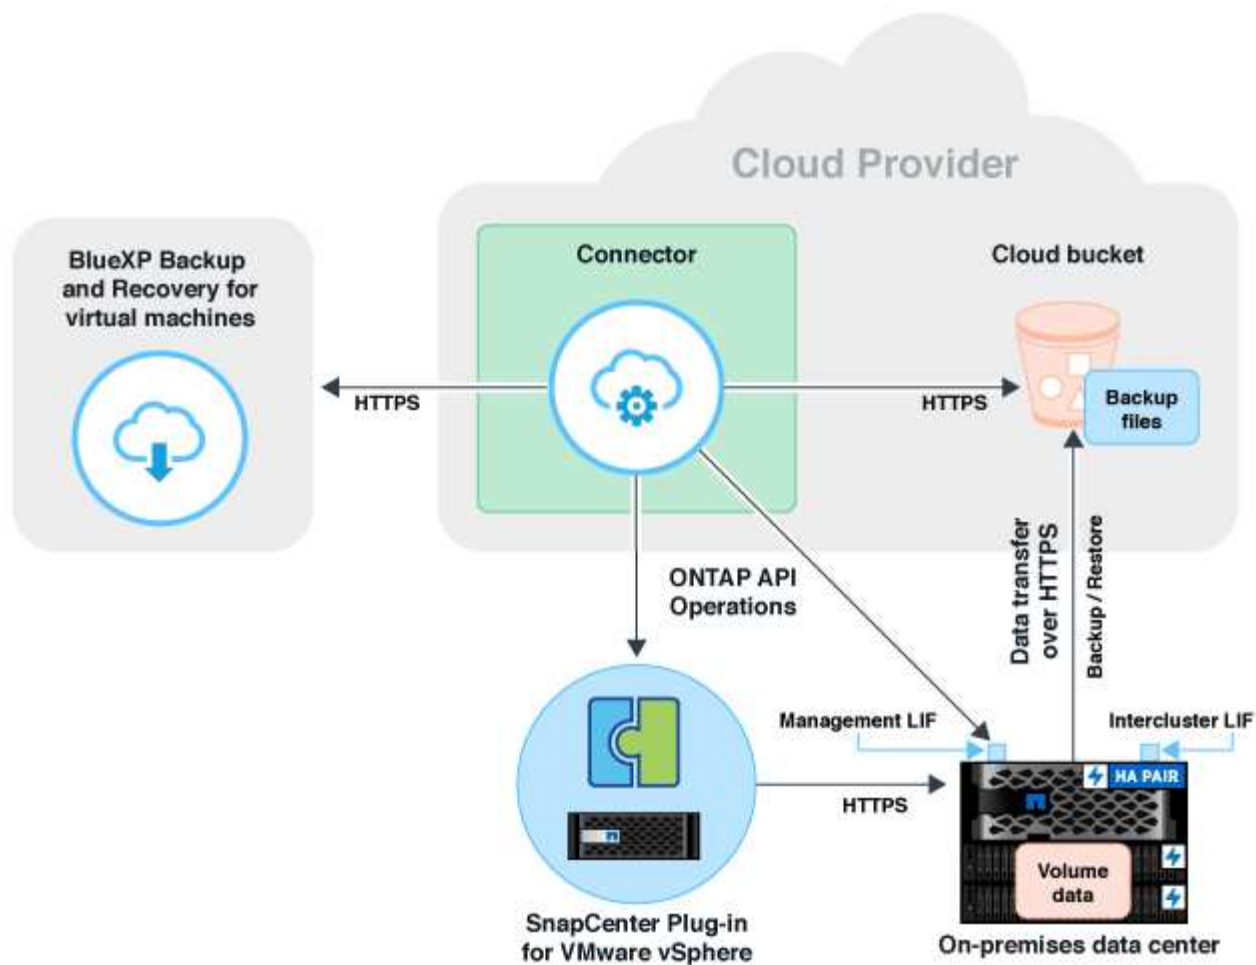
You can back up datastores to Amazon Web Services S3, Microsoft Azure Blob, Google Cloud Platform, and StorageGRID and restore virtual machines back to the on-premises SnapCenter Plug-in for VMware vSphere host. BlueXP backup and recovery for virtual machines also supports connector proxy deployment model.

Before you begin

Read the following requirements to make sure that you have a supported configuration before you start backing up datastores and virtual machines to a cloud provider.

- SnapCenter Plug-in for VMware vSphere 4.6P1 or later
 - You should be using SnapCenter Plug-in for VMware vSphere 4.7P1 or later to back up datastores from on-premises secondary storage.
- ONTAP 9.8 or later
- BlueXP
- NFS and VMFS datastores are supported. vVols are not supported.
- For VMFS support, SnapCenter Plug-in for VMware vSphere host should be running on 4.9 or later. Ensure to take a backup of the VMFS datastore if the SnapCenter Plug-in for VMware vSphere host was upgraded from an earlier version to the 4.9 release.
- At least one backup should have been taken in SnapCenter Plug-in for VMware vSphere 4.6P1.
- At least one daily, weekly, or monthly policy in SnapCenter Plug-in for VMware vSphere with no label or same label as that of the Virtual Machines policy in BlueXP.
- For pre-canned policy, the schedule tier should be the same for the datastore in SnapCenter Plug-in for VMware vSphere and in the cloud.
- Ensure that there are no FlexGroup volumes in the datastore because backing up and restoring FlexGroup volumes are not supported.
- Disable "**_recent**" on the required resource groups. If you have "**_recent**" enabled for the resource group, then the backups of those resource groups cannot be used for data protection to cloud and subsequently cannot be used for the restore operation.
- Ensure that the destination datastore where the virtual machine will be restored has enough space to accommodate a copy of all virtual machine files such as VMDK, VMX, VMSSD, and so on.
- Ensure that the destination datastore does not have stale virtual machine files in the format of restore_XXX_XXXXXX_filename from the previous restore operation failures. You should delete the stale files before triggering a restore operation.
- To deploy a connector with proxy configured, ensure that all outgoing connector calls are routed through the proxy server.

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



Register SnapCenter Plug-in for VMware vSphere host

You should register the SnapCenter Plug-in for VMware vSphere host in BlueXP for the datastores and virtual machines to be displayed. Only a user with administrative access can register the SnapCenter Plug-in for VMware vSphere host.



You can register multiple SnapCenter Plug-in for VMware vSphere hosts in BlueXP. However, once registered, you cannot remove the SnapCenter Plug-in for VMware vSphere host.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines**.
2. From the **Settings** drop-down, click **SnapCenter Plug-in for VMware vSphere**.
3. Click **Register SnapCenter Plug-in for VMware vSphere**.
4. Specify the following details:
 - a. In the SnapCenter Plug-in for VMware vSphere field, specify the FQDN or IP address of the SnapCenter Plug-in for VMware vSphere host.
 - b. In the Port field, specify the port number on which the SnapCenter Plug-in for VMware vSphere host is running.

You should ensure that communication is open between on-premises SnapCenter Plug-in for VMware vSphere host which is running on the default 8144 port and BlueXP Connector instance which could be either running in any cloud providers (Amazon Web Services, Microsoft Azure, Google Cloud Platform) or on-premises.

- c. In the Username and Password field, specify the credentials of the vCenter user with the administrator role.

5. Click **Register**.

After you finish

Click **Backup and recovery > Virtual Machines** to view all the datastores and virtual machines that are protected using the registered SnapCenter Plug-in for VMware vSphere host.

Create a policy to back up datastores

You can create a policy or use one of the following predefined policies that are available in BlueXP.

Before you begin

- You should create policies if you do not want to edit the predefined policies.
- To move backups from object store to archival storage, you should be running ONTAP 9.10.1 or later and Amazon Web Services or Microsoft Azure should be the cloud provider.
- You should configure the archive access tier for each cloud provider.

About this task

The following predefined policies are available in BlueXP:

Policy Name	Label	Retention Value
1 Year Daily LTR (Long Term Retention)	Daily	366
5 Years Daily LTR	Daily	1830
7 Year Weekly LTR	Weekly	370
10 Year Monthly LTR	Monthly	120

Steps

1. In the Virtual machines page, from the Settings drop-down list, select **Policies**.
2. Click **Create policy**.
3. In the Policy Details section, specify the policy name.
4. In the Retention section, select one of the retention type and specify the number of backups to retain.
5. Select Primary or Secondary as the backup storage source.
6. (Optional) If you want to move backups from object store to archival storage after a certain number of days for cost optimization, select the **Tier Backups to Archival** checkbox and enter the number of days after

which the backup should be archived.

7. Click **Create**.



You cannot edit or delete a policy, which is associated with a datastore.

Back up datastores to Amazon Web Services

You can back up and archive one or more datastores to Amazon Web Services to improve storage efficiency and cloud transition.

If the datastore is associated with an archival policy, you have an option to select the archival tier. The supported archival tiers are Glacier and Glacier Deep.

Before you begin

Ensure that you have met all the [requirements](#) before backing up datastores to the cloud.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines**.
2. Click **...** corresponding to the datastore that you want to back up and click **Activate Backup**.
3. In the Assign Policy page, select the policy and click **Next**.
4. Add the working environment.

Configure the cluster management LIF that you want BlueXP to discover. After adding the working environment for one of the datastores, it can be reused for all the other datastores residing on the same ONTAP cluster.

- a. Click **Add Working Environment** corresponding to the SVM.
 - b. In the Add Working Environment wizard:
 - i. Specify the IP address of the cluster management LIF.
 - ii. Specify the credentials of the ONTAP cluster user.
 - c. Click **Add Working Environment**.
5. Select **Amazon Web Services** to configure it as the cloud provider.
 - a. Specify the AWS account.
 - b. In the AWS Access Key field, specify the key for data encryption.
 - c. In the AWS Secret Key field, specify the password for data encryption.
 - d. Select the region where you want to create the backups.
 - e. Specify the IP addresses of the cluster management LIF that were added as the working environments.
 - f. Select the archival tier.

It is recommended to set the archival tier because this is an one time activity and you cannot set it up later.

6. Review the details and click **Activate Backup**.

Back up datastores to Microsoft Azure

You can back up one or more datastores to Microsoft Azure by integrating the SnapCenter Plug-in for VMware vSphere host with BlueXP. This will help the VM administrators to easily and quickly back up and archive data for storage efficiency and accelerate cloud transition.

If the datastore is associated with an archival policy, you will be provided with an option to select the archival tier. The supported archival tier is Azure Archive Blob Storage.

Before you begin

Ensure that you have met all the [requirements](#) before backing up datastores to the cloud.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines**.
2. Click **...** corresponding to the datastore that you want to back up and click **Activate Backup**.
3. In the Assign Policy page, select the policy and click **Next**.
4. Add the working environment.

Configure the cluster management LIF that you want BlueXP to discover. After adding the working environment for one of the datastores, it can be reused for all the other datastores residing on the same ONTAP cluster.

- a. Click **Add Working Environment** corresponding to the SVM.
 - b. In the Add Working Environment wizard:
 - i. Specify the IP address of the cluster management LIF.
 - ii. Specify the credentials of the ONTAP cluster user.
 - c. Click **Add Working Environment**.
5. Select **Microsoft Azure** to configure it as the cloud provider.
 - a. Specify the Azure subscription ID.
 - b. Select the region where you want to create the backups.
 - c. Create a new resource group or use an existing resource group.
 - d. Specify the IP addresses of the cluster management LIF that were added as the working environments.
 - e. Select the archival tier.

It is recommended to set the archival tier because this is an one time activity and you will not be allowed to set it up later.

6. Review the details and click **Activate Backup**.

Back up datastores to Google Cloud Platform

You can back up one or more datastores to Google Cloud Platform by integrating the SnapCenter Plug-in for VMware vSphere host with BlueXP. This will help the VM administrators to easily and quickly back up and archive data for storage efficiency and

accelerate cloud transition.

Before you begin

Ensure that you have met all the [requirements](#) before backing up datastores to the cloud.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines**.
2. Click **...** corresponding to the datastore that you want to back up and click **Activate Backup**.
3. In the Assign Policy page, select the policy and click **Next**.
4. Add the working environment.

Configure the cluster management LIF that you want BlueXP to discover. After adding the working environment for one of the datastores, it can be reused for all the other datastores residing on the same ONTAP cluster.

- a. Click **Add Working Environment** corresponding to the SVM.
 - b. In the Add Working Environment wizard:
 - i. Specify the IP address of the cluster management LIF.
 - ii. Specify the credentials of the ONTAP cluster user.
 - c. Click **Add Working Environment**.
5. Select **Google Cloud Platform** to configure it as the cloud provider.
 - a. Select the Google Cloud Project where you want the Google Cloud Storage bucket to be created for backups.
 - b. In the Google Cloud Access Key field, specify the key.
 - c. In the Google Cloud Secret Key field, specify the password.
 - d. Select the region where you want to create the backups.
 - e. Specify the IP space.
 6. Review the details and click **Activate Backup**.

Back up datastores to StorageGRID

You can back up one or more datastores to StorageGRID by integrating the SnapCenter Plug-in for VMware vSphere host with BlueXP. This will help the VM administrators to easily and quickly back up and archive data for storage efficiency and accelerate cloud transition.

Before you begin

Ensure that you have met all the [requirements](#) before backing up datastores to the cloud.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines**.
2. Click **...** corresponding to the datastore that you want to back up and click **Activate Backup**.
3. In the Assign Policy page, select the policy and click **Next**.
4. Add the working environment.

Configure the cluster management LIF that you want BlueXP to discover. After adding the working environment for one of the datastores, it can be reused for all the other datastores residing on the same ONTAP cluster.

- a. Click **Add Working Environment** corresponding to the SVM.
 - b. In the Add Working Environment wizard:
 - i. Specify the IP address of the cluster management LIF.
 - ii. Specify the credentials of the ONTAP cluster user.
 - c. Click **Add Working Environment**.
5. Select **StorageGRID**.
 - a. Specify the Storage Server IP.
 - b. Select the access key and secret key.
 6. Review the details and click **Activate Backup**.

Manage protection of datastores and virtual machines data

You can view policies, datastores, and virtual machines before you back up and restore data. Depending upon the change in database, policies, or resource groups, you can view the updates from the BlueXP UI.

View policies

You can view all the default pre-canned policies. For each of these policies, when you view the details, all the associated policies and virtual machines are listed.

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines**.
2. From the **Settings** drop-down, click **Policies**.
3. Click **View Details** corresponding to policy whose details you want to view.

The associated policies and virtual machines are listed.

View datastores and virtual machines

The datastores and virtual machines that are protected using the registered SnapCenter Plug-in for VMware vSphere host are displayed.


Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines > Settings > SnapCenter Plug-in for VMware vSphere**.
2. Click the SnapCenter Plug-in for VMware vSphere host for which you want to see the datastores and virtual machines.

Unprotect datastores

You can unprotect a datastore which was already protected earlier. You can unprotect a datastore when you want to delete the cloud backups or do not want to back it up to the cloud anymore. The datastore can be protected again after the unprotection is successful.


Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines**.
2. Click  corresponding to the datastore that you want to unprotect and click **Unprotect**.

Edit the SnapCenter Plug-in for VMware vSphere Instance


You can edit the details of the SnapCenter Plug-in for VMware vSphere host in BlueXP.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines > Settings > SnapCenter Plug-in for VMware vSphere**.
2. Click  and select **Edit**.
3. Modify the details as required.
4. Click **Save**.

Refresh resources and backups

If you want to view the latest datastores and backups that have been added to the application, you should refresh the resources and backups. This will initiate the discovery of the resources and backups and the latest details will be displayed.

1. Click **Backup and recovery > Virtual Machines**.
2. From the **Settings** drop-down, click **SnapCenter Plug-in for VMware vSphere**.
3. Click  corresponding to the SnapCenter Plug-in for VMware vSphere host and click **Refresh Resources and Backups**.


Refresh policy or resource group

If there is a change to the policy or resource group, you should refresh the protection relationship.

1. Click **Backup and recovery > Virtual Machines**.
2. Click  corresponding to the datastore and click **Refresh Protection**.

Unregister SnapCenter Plug-in for VMware vSphere host

All datastores and virtual machines associated with the SnapCenter Plug-in for VMware vSphere host will be unprotected.

1. Click **Backup and recovery > Virtual Machines**.
2. From the **Settings** drop-down, click **SnapCenter Plug-in for VMware vSphere**.
3. Click  corresponding to the SnapCenter Plug-in for VMware vSphere host and click **Unregister**.

Monitor Jobs

Jobs are created for all the BlueXP backup and recovery operations. You can monitor all the jobs and all the sub tasks that are performed as part of each task.

1. Click **Backup and recovery > Job Monitoring**.

When you initiate an operation, a window appears stating that the job is initiated. You can click the link to monitor the job.

2. Click the primary task to view the sub tasks and status of each of these sub tasks.

Restore virtual machines data from the cloud

You can restore virtual machines data from the cloud back to the on-premises vCenter. You can restore the virtual machine to the exact same location from where the backup was taken or to an alternate location. If the virtual machine was backed up using archival policy, then you can set the archival restore priority.



You cannot restore virtual machines that span across datastores.

Before you begin

- Ensure that you have met all the [requirements](#) before restoring virtual machines from the cloud.
- If you are restoring to an alternate location:
 - Ensure that the source and destination vCenters are in linked mode.
 - Ensure that the source and destination cluster details are added in BlueXP Canvas and in linked mode vCenters in both SnapCenter Plug-in for VMware vSphere host.
 - Ensure that the Working Environment (WE) is added corresponding to the alternate location in BlueXP Canvas.

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Virtual Machines > SnapCenter Plug-in for VMware vSphere** and select the SnapCenter Plug-in for VMware vSphere host.



If the source virtual machine is moved to another location (vMotion), and if the user triggers a restore of that virtual machine from BlueXP, then the virtual machine is restored to the source location from where the backup was taken.

2. You can restore the virtual machine to the original location or to an alternate location from the datastore or from virtual machines:

If you want to restore the virtual machine...	Do this...
to the original location from datastore	<ol style="list-style-type: none"> 1. Click ... corresponding to the datastore that you want to restore and click View Details. 2. Click Restore corresponding to the backup you want to restore. 3. Select the virtual machine that you want to restore from the backup and click Next. 4. Ensure that Original is selected and click Continue. 5. If the virtual machine is protected using a policy where archival settings are configured, select the Archival Restore Priority and click Next. The supported archival restore priority for Amazon Web Services are high, standard, and low and the supported archival restore priority for Microsoft Azure are high and standard. 6. Review the details and click Restore.
to an alternate location from datastore	<ol style="list-style-type: none"> 1. Click ... corresponding to the datastore that you want to restore and click View Details. 2. Click Restore corresponding to the backup you want to restore. 3. Select the virtual machine that you want to restore from the backup and click Next. 4. Select Alternate. 5. Select the alternate vCenter Server, ESXi host, datastore, and network. 6. Provide a name for the VM after restore and click Continue. 7. If the virtual machine is protected using a policy where archival settings are configured, select the Archival Restore Priority and click Next. The supported archival restore priority for Amazon Web Services are high, standard, and low and the supported archival restore priority for Microsoft Azure are high and standard. 8. Review the details and click Restore.

If you want to restore the virtual machine...	Do this...
to the original location from virtual machines	<ol style="list-style-type: none"> 1. Click ... corresponding to the virtual machine that you want to restore and click Restore. 2. Select the backup through which you want to restore the virtual machine. 3. Ensure that Original is selected and click Continue. 4. If the virtual machine is protected using a policy where archival settings are configured, select the Archival Restore Priority and click Next. The supported archival restore priority for Amazon Web Services are high, standard, and low and the supported archival restore priority for Microsoft Azure are high and standard. 5. Review the details and click Restore.
to an alternate location from virtual machines	<ol style="list-style-type: none"> 1. Click ... corresponding to the virtual machine that you want to restore and click Restore. 2. Select the backup through which you want to restore the virtual machine. 3. Select Alternate. 4. Select the alternate vCenter Server, ESXi host, datastore, and network. 5. Provide a name for the VM after restore and click Continue. 6. If the virtual machine is protected using a policy where archival settings are configured, select the Archival Restore Priority and click Next. The supported archival restore priority for Amazon Web Services are high, standard, and low and the supported archival restore priority for Microsoft Azure are high and standard. 7. Review the details and click Restore.

Back up and restore Kubernetes data

Protect your Kubernetes cluster data using BlueXP backup and recovery

BlueXP backup and recovery provides backup and restore capabilities for protection and long-term archive of your Kubernetes cluster data. Backups are automatically generated and stored in an object store in your public or private cloud account.

When necessary, you can restore an entire *volume* from a backup to the same or different working environment.

Features

Backup features:

- Back up independent copies of your persistent 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.
- Backup data is secured with AES-256 bit encryption at-rest and TLS 1.2 HTTPS connections in-flight.
- Support for up to 4,000 backups of a single volume.

Restore features:

- Restore data from a specific point in time.
- Restore a volume to the source system or to a different system.
- Restores data on a block level, placing the data directly in the location you specify, all while preserving the original ACLs.

Supported Kubernetes working environments and object storage providers

BlueXP backup and recovery enables you to back up Kubernetes volumes from the following working environments to object storage in the following public and private cloud providers:

Source Working Environment	Backup File Destination
Kubernetes cluster in AWS	Amazon S3
Kubernetes cluster in Azure	Azure Blob
Kubernetes cluster in Google	Google Cloud Storage

You can restore a volume from a Kubernetes backup file to the following working environments:

Backup File Location	Destination Working Environment
Amazon S3	Kubernetes cluster in AWS
Azure Blob	Kubernetes cluster in Azure
Google Cloud Storage	Kubernetes cluster in Google

Cost

There are two types of costs associated with using BlueXP backup and recovery: resource charges and service charges.

Resource charges

Resource charges are paid to the cloud provider for object storage capacity in the cloud. 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).

Service charges

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

There are two 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 purchase licenses directly from NetApp. Read the [Licensing](#) section for details.

Licensing

BlueXP backup and recovery is available in two licensing options: Pay As You Go (PAYGO), and Bring Your Own License (BYOL). A 30-day free trial is available if you don't have a license.

Free trial

When using the 30-day free trial, you are notified about the number of free trial days that remain. At the end of your free trial, backups stop being created. You must subscribe to the service or purchase a license to continue using the service.

Backup files are not deleted when the service is disabled. You'll continue to be charged by your cloud provider for object storage costs for the capacity that your backups use unless you delete the backups.

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 GB for data that's backed up—there's no up-front payment. You are billed by your cloud provider through your monthly bill.

You should subscribe even if you have a free trial or if you bring your own license (BYOL):

- Subscribing ensures that there's no disruption of service after your free trial ends.

When the trial ends, you'll be charged hourly according to the amount of data that you back up.

- If you back up more data than allowed by your BYOL license, then data backup continues through your pay-as-you-go subscription.

For example, if you have a 10 TB BYOL license, all capacity beyond the 10 TB is charged through the PAYGO subscription.

You won't be charged from your pay-as-you-go subscription during your free trial or if you haven't exceeded your BYOL license.

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

Bring your own license

BYOL is term-based (12, 24, or 36 months) *and* capacity-based in 1 TB increments. You pay NetApp to use the service for a period of time, say 1 year, and for a maximum amount capacity, say 10 TB.

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.](#)


How BlueXP backup and recovery works

When you enable BlueXP backup and recovery on a Kubernetes 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.



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

The following image shows the relationship between each component:



Supported storage classes or access tiers

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

- In Azure, backups are associated with the *Cool* access tier.
- In GCP, backups are associated with the *Standard* storage class by default.

Customizable backup schedule and retention settings per cluster

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 define. 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.

You can choose a combination of hourly, daily, weekly, and monthly backups of all volumes.

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.

Supported volumes

BlueXP backup and recovery supports Persistent volumes (PVs).

Limitations

- When creating or editing a backup policy when no volumes are assigned to the policy, the number of retained backups can be a maximum of 1018. As a workaround you can reduce the number of backups to create the policy. Then you can edit the policy to create up to 4000 backups after you assign volumes to the policy.
- Ad-hoc volume backups using the **Backup Now** button aren't supported on Kubernetes volumes.

Backing up Kubernetes persistent volume data to Amazon S3

Complete a few steps to get started backing up data from your persistent volumes on EKS Kubernetes clusters to Amazon S3 storage.

Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.

1

Review prerequisites

- You have discovered the Kubernetes cluster as a BlueXP working environment.
 - Trident must be installed on the cluster, and the Trident version must be 21.1 or greater.
 - All PVCs that will be used to create persistent volumes that you want to back up must have "snapshotPolicy" set to "default".
 - The cluster must be using Cloud Volumes ONTAP on AWS for its' backend storage.
 - The Cloud Volumes ONTAP system must be running ONTAP 9.7P5 or later.
- 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.

- The IAM role that provides the BlueXP Connector with permissions includes S3 permissions from the latest [BlueXP policy](#).

2

Enable BlueXP backup and recovery on your existing Kubernetes cluster

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



3

Define the backup policy

The default policy backs up volumes every day and retains the most recent 30 backup copies of each volume. Change to hourly, daily, weekly, or monthly backups, or select one of the system-defined policies that provide more options. You can also change the number of backup copies you want to retain.

 A screenshot of a 'Define Policy' wizard screen. The title 'Define Policy' is at the top in a light blue font. Below it, the section 'Policy - Retention & Schedule' is highlighted in blue. The screen displays four backup frequency options, each with a checkbox and a corresponding 'Number of backups to retain' value in a dropdown menu:

Frequency	Number of backups to retain
<input type="checkbox"/> Hourly	24
<input checked="" type="checkbox"/> Daily	30
<input type="checkbox"/> Weekly	52
<input type="checkbox"/> Monthly	12

 At the bottom left, there is a link 'S3 Bucket'. At the bottom right, a message states: 'Cloud Manager will create the S3 bucket after you complete the wizard'.

4


Select the volumes that you want to back up

Identify which volumes you want to back up in the Select Volumes page. An S3 bucket is created automatically in the same AWS account and Region as the Cloud Volumes ONTAP system, and the backup files are stored there.

Requirements

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

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



Note that the VPC Endpoint is optional.

Kubernetes cluster requirements

- You have discovered the Kubernetes cluster as a BlueXP working environment. [See how to discover the Kubernetes cluster](#).
- Trident must be installed on the cluster, and the Trident version must be a minimum of 21.1. See [how to install Trident](#) or [how to upgrade the Trident version](#).
- The cluster must be using Cloud Volumes ONTAP on AWS for its' backend storage.
- The Cloud Volumes ONTAP system must be in the same AWS region as the Kubernetes cluster, and it must be running ONTAP 9.7P5 or later (ONTAP 9.8P11 and later is recommended).

Note that Kubernetes clusters in on-premises locations are not supported. Only Kubernetes clusters in cloud deployments that are using Cloud Volumes ONTAP systems are supported.

- All Persistent Volume Claim objects that will be used to create the persistent volumes that you want to back up must have "snapshotPolicy" set to "default".

You can do this for individual PVCs by adding `snapshotPolicy` under annotations:

```

kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: full
  annotations:
    trident.netapp.io/snapshotPolicy: "default"
spec:
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 1000Mi
  storageClassName: silver

```

You can do this for all PVCs associated with a particular backend storage by adding the `snapshotPolicy` field under `defaults` in the `backend.json` file:

```

apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-ontap-nas-advanced
spec:
  version: 1
  storageDriverName: ontap-nas
  managementLIF: 10.0.0.1
  dataLIF: 10.0.0.2
  backendName: tbc-ontap-nas-advanced
  svm: trident_svm
  credentials:
    name: backend-tbc-ontap-nas-advanced-secret
  limitAggregateUsage: 80%
  limitVolumeSize: 50Gi
  nfsMountOptions: nfsvers=4
  defaults:
    spaceReserve: volume
    exportPolicy: myk8scluster
    snapshotPolicy: default
    snapshotReserve: '10'
    deletionPolicy: retain

```

License requirements

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

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

Supported AWS regions

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

AWS Backup permissions required

The IAM role that provides BlueXP with permissions must include S3 permissions from the latest [BlueXP policy](#).

Here are the specific S3 permissions from the policy:

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

Enabling BlueXP backup and recovery

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

Steps

1. Select the working environment and click **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 Kubernetes cluster onto the Amazon S3 working environment to initiate the setup wizard.



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

You can define the backup schedule and choose the number of backups to retain.

Define Policy

Policy - Retention & Schedule

☐ Hourly

☒ Daily

☐ Weekly

☐ Monthly

Number of backups to retain

S3 Bucket Cloud Manager will create the S3 bucket after you complete the wizard

3. Select the persistent volumes that you want to back up.

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

Select Volumes

57 Volumes

<input checked="" type="checkbox"/>	Persistent Volume Name	Namespace	Allocated Capacity	Backup Status
<input checked="" type="checkbox"/>	Persistent Volume 1 <small>On</small>	Namespace 1	10 TB	<input type="radio"/> Not Active
<input checked="" type="checkbox"/>	Persistent Volume 2 <small>On</small>	Namespace 1	10 TB	<input type="radio"/> Not Active
<input checked="" type="checkbox"/>	Persistent Volume 3 <small>On</small>	Namespace 1	10 TB	<input type="radio"/> Not Active
<input checked="" type="checkbox"/>	PV 1 <small>On</small>	Namespace 2	10 TB	<input type="radio"/> Not Active
<input checked="" type="checkbox"/>	PV 2 <small>On</small>	Namespace 2	10 TB	<input type="radio"/> Not Active

☒ Automatically back up all existing and future persistent volumes with the selected backup policy

4. If you want all current and future volumes to have backup enabled, just leave the checkbox for "Automatically back up future volumes..." checked. If you disable this setting, you'll need to manually enable backups for future volumes.
5. Click **Activate Backup** and BlueXP backup and recovery starts taking the initial backups of each selected volume.

Result

An S3 bucket is created automatically in the same AWS account and Region as the Cloud Volumes ONTAP system, and the backup files are stored there.

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

What's next?

You can [start and stop backups for volumes or change the backup schedule](#).

You can also [restore entire volumes from a backup file](#) as a new volume on the same or different Kubernetes cluster in AWS (in the same region).

Backing up Kubernetes persistent volume data to Azure Blob storage

Complete a few steps to get started backing up data from your persistent volumes on AKS Kubernetes clusters to Azure Blob storage.

Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.

1


Review prerequisites

- You have discovered the Kubernetes cluster as a BlueXP working environment.
 - Trident must be installed on the cluster, and the Trident version must be 21.1 or greater.
 - All PVCs that will be used to create persistent volumes that you want to back up must have "snapshotPolicy" set to "default".
 - The cluster must be using Cloud Volumes ONTAP on Azure for its' backend storage.
 - The Cloud Volumes ONTAP system must be running ONTAP 9.7P5 or later.
- 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

Enable BlueXP backup and recovery on your existing Kubernetes cluster

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



3

Define the backup policy

The default policy backs up volumes every day and retains the most recent 30 backup copies of each volume. Change to hourly, daily, weekly, or monthly backups, or select one of the system-defined policies that provide more options. You can also change the number of backup copies you want to retain.

Define Policy

Policy - Retention & Schedule

<input type="checkbox"/> Hourly	Number of backups to retain	24
<input checked="" type="checkbox"/> Daily	Number of backups to retain	30
<input type="checkbox"/> Weekly	Number of backups to retain	52
<input type="checkbox"/> Monthly	Number of backups to retain	12

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

4

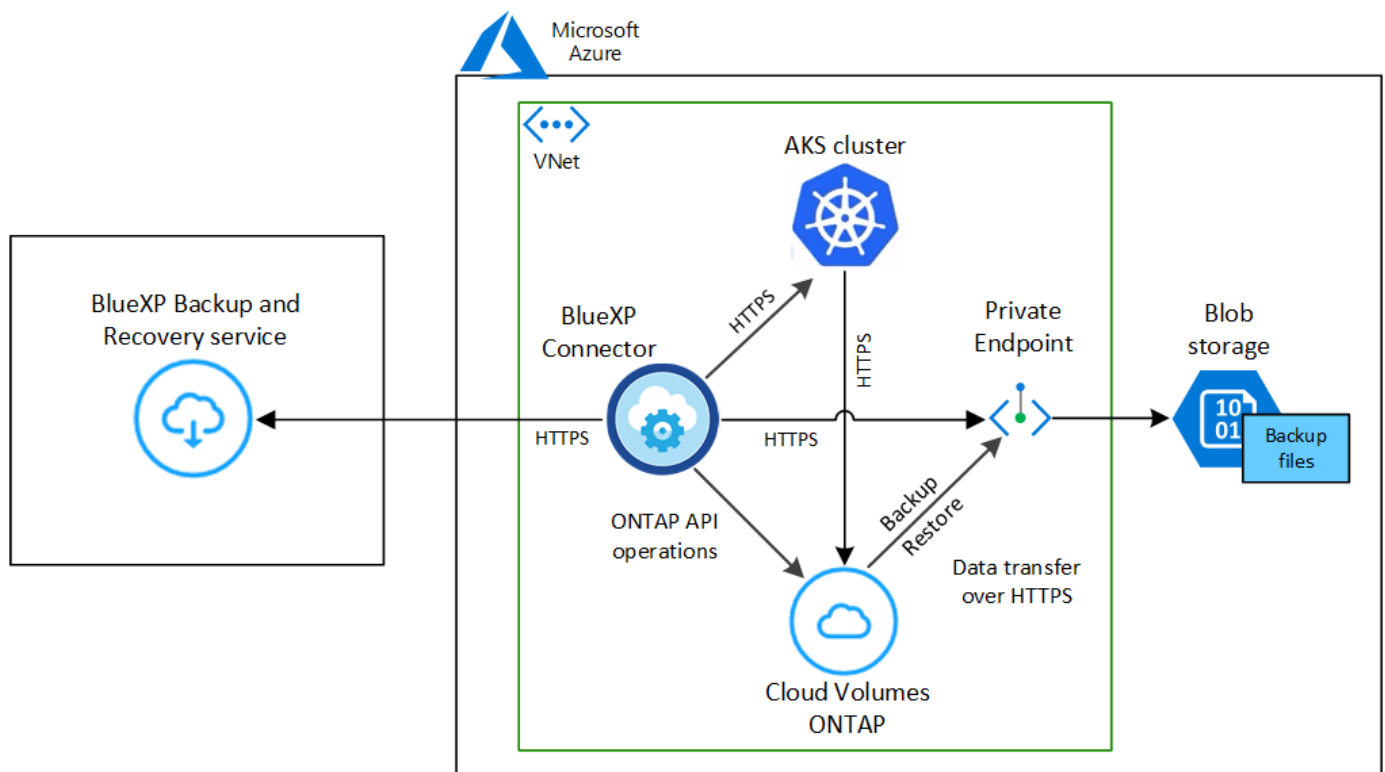
Select the volumes that you want to back up

Identify which volumes you want to back up in the Select Volumes page. The backup files are stored in a Blob container using the same Azure subscription and Region as the Cloud Volumes ONTAP system.

Requirements

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

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



Note that the Private Endpoint is optional.

Kubernetes cluster requirements

- You have discovered the Kubernetes cluster as a BlueXP working environment. [See how to discover the Kubernetes cluster](#).
- Trident must be installed on the cluster, and the Trident version must be a minimum of 21.1. See [how to install Trident](#) or [how to upgrade the Trident version](#).
- The cluster must be using Cloud Volumes ONTAP on Azure for its' backend storage.
- The Cloud Volumes ONTAP system must be in the same Azure region as the Kubernetes cluster, and it must be running ONTAP 9.7P5 or later (ONTAP 9.8P11 and later is recommended).

Note that Kubernetes clusters in on-premises locations are not supported. Only Kubernetes clusters in cloud deployments that are using Cloud Volumes ONTAP systems are supported.

- All Persistent Volume Claim objects that will be used to create the persistent volumes that you want to back up must have "snapshotPolicy" set to "default".

You can do this for individual PVCs by adding `snapshotPolicy` under annotations:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: full
  annotations:
    trident.netapp.io/snapshotPolicy: "default"
spec:
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 1000Mi
  storageClassName: silver
```

You can do this for all PVCs associated with a particular backend storage by adding the `snapshotPolicy` field under defaults in the `backend.json` file:

```

apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-ontap-nas-advanced
spec:
  version: 1
  storageDriverName: ontap-nas
  managementLIF: 10.0.0.1
  dataLIF: 10.0.0.2
  backendName: tbc-ontap-nas-advanced
  svm: trident_svm
  credentials:
    name: backend-tbc-ontap-nas-advanced-secret
  limitAggregateUsage: 80%
  limitVolumeSize: 50Gi
  nfsMountOptions: nfsvers=4
  defaults:
    spaceReserve: volume
    exportPolicy: myk8scluster
    snapshotPolicy: default
    snapshotReserve: '10'
    deletionPolicy: retain

```

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.](#)

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

Supported Azure regions

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

Enabling BlueXP backup and recovery

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

Steps

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



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

You can define the backup schedule and choose the number of backups to retain.

3. Select the persistent volumes that you want to back up.

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

<input checked="" type="checkbox"/>	Persistent Volume Name	Namespace	Allocated Capacity	Backup Status
<input checked="" type="checkbox"/>	Persistent Volume 1 On	Namespace 1	10 TB	Not Active
<input checked="" type="checkbox"/>	Persistent Volume 2 On	Namespace 1	10 TB	Not Active
<input checked="" type="checkbox"/>	Persistent Volume 3 On	Namespace 1	10 TB	Not Active
<input checked="" type="checkbox"/>	PV1 On	Namespace 2	10 TB	Not Active
<input checked="" type="checkbox"/>	PV2 On	Namespace 2	10 TB	Not Active

☒ Automatically back up all existing and future persistent volumes with the selected backup policy

4. If you want all current and future volumes to have backup enabled, just leave the checkbox for "Automatically back up future volumes..." checked. If you disable this setting, you'll need to manually enable backups for future volumes.
5. Click **Activate Backup** and BlueXP backup and recovery starts taking the initial backups of each selected volume.

Result

The backup files are stored in a Blob container using the same Azure subscription and Region as the Cloud Volumes ONTAP system.

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

What's next?

You can [start and stop backups for volumes or change the backup schedule](#).

You can also [restore entire volumes from a backup file](#) as a new volume on the same or different Kubernetes cluster in Azure (in the same region).

Backing up Kubernetes persistent volume data to Google Cloud storage

Complete a few steps to get started backing up data from your persistent volumes on GKE Kubernetes clusters to Google Cloud storage.

Quick start

Get started quickly by following these steps or scroll down to the remaining sections for full details.

1


Review prerequisites

- You have discovered the Kubernetes cluster as a BlueXP working environment.
 - Trident must be installed on the cluster, and the Trident version must be 21.1 or greater.
 - All PVCs that will be used to create persistent volumes that you want to back up must have "snapshotPolicy" set to "default".
 - The cluster must be using Cloud Volumes ONTAP on GCP for its' backend storage.
 - The Cloud Volumes ONTAP system must be running ONTAP 9.7P5 or later.
- 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 the predefined Storage Admin role.
- 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

Enable BlueXP backup and recovery on your existing Kubernetes cluster

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



3

Define the backup policy

The default policy backs up volumes every day and retains the most recent 30 backup copies of each volume. Change to hourly, daily, weekly, or monthly backups, or select one of the system-defined policies that provide more options. You can also change the number of backup copies you want to retain.

Define Policy

Policy - Retention & Schedule

☐ Hourly

Number of backups to retain

24

☒ Daily

Number of backups to retain

30

☐ Weekly

Number of backups to retain

52

☐ Monthly

Number of backups to retain

12

Storage Account

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

4


Select the volumes that you want to back up

Identify which volumes you want to back up in the Select Volumes page. The backup files are stored in a Google Cloud Storage bucket using the same GCP subscription and Region as the Cloud Volumes ONTAP system.

Requirements

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

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



Note that the Private Endpoint is optional.

Kubernetes cluster requirements

- You have discovered the Kubernetes cluster as a BlueXP working environment. [See how to discover the Kubernetes cluster](#).
- Trident must be installed on the cluster, and the Trident version must be a minimum of 21.1. See [how to install Trident](#) or [how to upgrade the Trident version](#).
- The cluster must be using Cloud Volumes ONTAP on GCP for its' backend storage.
- The Cloud Volumes ONTAP system must be in the same GCP region as the Kubernetes cluster, and it must be running ONTAP 9.7P5 or later (ONTAP 9.8P11 and later is recommended).

Note that Kubernetes clusters in on-premises locations are not supported. Only Kubernetes clusters in cloud deployments that are using Cloud Volumes ONTAP systems are supported.

- All Persistent Volume Claim objects that will be used to create the persistent volumes that you want to back up must have "snapshotPolicy" set to "default".

You can do this for individual PVCs by adding `snapshotPolicy` under annotations:

```
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: full
  annotations:
    trident.netapp.io/snapshotPolicy: "default"
spec:
  accessModes:
    - ReadWriteMany
  resources:
    requests:
      storage: 1000Mi
  storageClassName: silver
```

You can do this for all PVCs associated with a particular backend storage by adding the `snapshotPolicy` field under defaults in the `backend.json` file:

```

apiVersion: trident.netapp.io/v1
kind: TridentBackendConfig
metadata:
  name: backend-tbc-ontap-nas-advanced
spec:
  version: 1
  storageDriverName: ontap-nas
  managementLIF: 10.0.0.1
  dataLIF: 10.0.0.2
  backendName: tbc-ontap-nas-advanced
  svm: trident_svm
  credentials:
    name: backend-tbc-ontap-nas-advanced-secret
  limitAggregateUsage: 80%
  limitVolumeSize: 50Gi
  nfsMountOptions: nfsvers=4
  defaults:
    spaceReserve: volume
    exportPolicy: myk8scluster
    snapshotPolicy: default
    snapshotReserve: '10'
    deletionPolicy: retain

```

Supported GCP regions

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

License requirements

For BlueXP backup and recovery PAYGO licensing, a subscription through the [GCP 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](#).

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

GCP Service Account


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

Enabling BlueXP backup and recovery

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

Steps

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



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

You can define the backup schedule and choose the number of backups to retain.

3. Select the persistent volumes that you want to back up.

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

4. If you want all current and future volumes to have backup enabled, just leave the checkbox for "Automatically back up future volumes..." checked. If you disable this setting, you'll need to manually enable backups for future volumes.

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

Result

The backup files are stored in a Google Cloud Storage bucket using the same GCP subscription and Region as the Cloud Volumes ONTAP system.

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

What's next?

You can [start and stop backups for volumes or change the backup schedule](#).

You can also [restore entire volumes from a backup file](#) as a new volume on the same or different Kubernetes cluster in GCP (in the same region).

Managing backups for your Kubernetes systems

You can manage backups for your Kubernetes systems by changing the backup schedule, enabling/disabling volume backups, deleting backups, and more.



Do not manage or change backup files directly from your cloud provider environment. This may corrupt the files and will result in an unsupported configuration.

Viewing the volumes that are being backed up

You can view a list of all the volumes that are currently being backed up by BlueXP backup and recovery.

Steps

1. From the BlueXP menu, select **Protection > Backup and recovery**.
2. Click the **Kubernetes** tab to view the list of persistent volumes for Kubernetes systems.

Source K8s Cluster	Source Persistent Volume	Source Namespace	Last Backup	Backup Copies	Backup Status
aws eks1 Unknown	pvc-1704aa1f-af1d-49e9-87fd-6edd86125855 Unknown	default	Jun 09 2022, 10:00:24 am	20	Unknown
aws eks1 Unknown	pvc-1615f0a8-2d5d-44d0-b4e4-f365cc3fb4a6 Unknown	default	Jun 09 2022, 10:00:24 am	20	Unknown
aws eks1 Unknown	pvc-d1f839c1-d932-4f49-b620-33321dbe939e Unknown	trident	Jun 09 2022, 10:00:24 am	20	Unknown


If you are looking for specific volumes in certain clusters, you can refine the list by cluster and volume, or you can use the search filter.

Enabling and disabling backups of volumes


You can stop backing up a volume if you do not need backup copies of that volume and you do not want to pay for the cost to store the backups. You can also add a new volume to the backup list if it is not currently being backed up.

Steps


1. From the **Kubernetes** tab, select **Backup Settings**.



2. From the *Backup Settings* page, click ... for the Kubernetes cluster and select **Manage Persistent Volumes**.



3. Select the checkbox for a volume, or volumes, that you want to change, and then click **Activate** or **Deactivate** depending on whether you want to start or stop backups for the volume.



4. Click **Save** to commit your changes.


Note: When stopping a volume from being backed up you'll continue to be charged by your cloud provider for object storage costs for the capacity that the backups use unless you [delete the backups](#).

Editing an existing backup 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.

Steps

1. From the **Kubernetes** tab, select **Backup Settings**.




2. From the *Backup Settings* page, click ... for the working environment where you want to change the settings, and select **Manage Policies**.



3. From the *Manage Policies* page, click **Edit Policy** for the backup policy you want to change in that working environment.



- From the *Edit Policy* page, change the schedule and backup retention and click **Save**.



Setting a backup policy to be assigned to new volumes


If you did not select the option to automatically assign a backup policy to newly created volumes when you first activated BlueXP backup and recovery on your Kubernetes cluster, you can choose this option in the *Backup Settings* page later. Having a backup policy assigned to newly created volumes ensures that all your data is protected.

Note that the policy that you want to apply to the volumes must already exist.

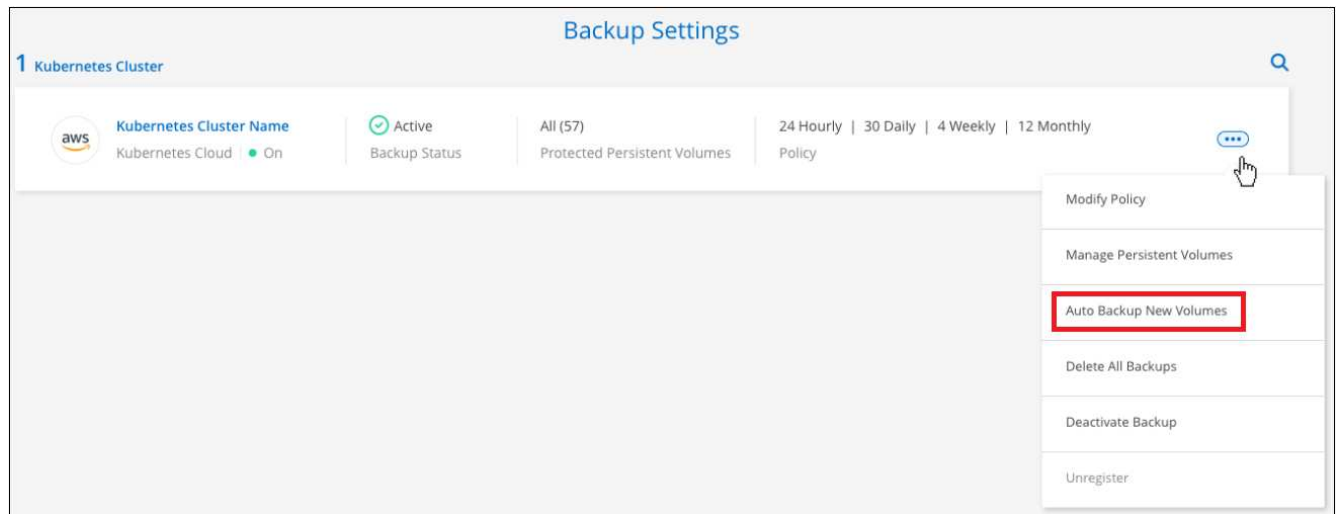
You can also disable this setting so that newly created volumes do not get backed up automatically. In that case you'll need to manually enable backups for any specific volumes that you do want to back up in the future.

Steps

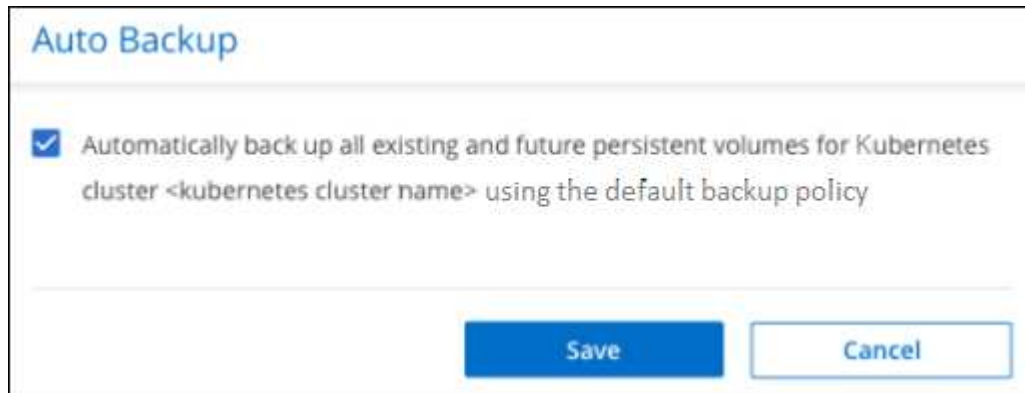
- From the **Kubernetes** tab, select **Backup Settings**.



- From the *Backup Settings* page, click ... for the Kubernetes cluster where the volumes exist, and select **Auto Backup New Volumes**.



3. Select the checkbox "Automatically back up future persistent volumes...", choose the backup policy that you want to apply to new volumes, and click **Save**.



Result

Now this backup policy will be applied to any new volumes created in this Kubernetes cluster.

Viewing 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.

This page also enables you perform the following tasks:

- Delete all backup files for the volume
- Delete individual backup files for the volume
- Download a backup report for the volume

Steps

1. From the **Kubernetes** tab, click ... for the source volume and select **Details & Backup List**.

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

All Kubernetes Clusters

Backup Settings

1 Kubernetes Clusters | 57 Protected PVs | 15.1 TB Total Backups Size

Protected Persistent Volumes Status: 57 Healthy Backup | 0 Failed Backup

57 Backups

Source Kubernetes Cluster	Source Persistent Volume	Source Namespace	Last Backup	Backups	Backup Status
Kubernetes_Cloud_AWS	Source Persistent Volume	Source Namespace	May 22 2019, 00:00:00	2,050 Backups	Active
Kubernetes_Cloud_AWS	Source Persistent Volume	Source Namespace	May 22 2019, 00:00:00	2,050 Snapshot	
Kubernetes_Cloud_AWS	Source Persistent Volume	Source Namespace	May 22 2019, 00:00:00	2,050 Snapshot	

Details & Backup List | Backup Now | Pause Backups

The list of all backup files is displayed along with details about the source volume, destination location, and backup details.

Source

Kubernetes Cluster: eks1

Type: EKS

Provider: AWS

Persistent Volume: pvc-05881c70-cf5f-4edc-8537...

Namespace: default

Destination

Cloud Provider: AWS

Bucket: netapp-backup-vsa5twmc9ae

Region: us-west-1

Account ID: 123456789012

Backup Information

Relationship Status: enabled

Last Backup: Dec 07 2021, 2:20:30 pm

Lag Duration: 1 hour

Backups: 2

Backup Policy: 24 hourly | 30 daily | 52 weekly

2 Backups

Backup Name	Date	Size
daily-dem-163887957011628bef197-34b5-11ec-8916-5b2669f1987a	Dec 07 2021, 2:19:30 pm	9.77 KB
daily-dem-163887963015128bef197-34b5-11ec-8916-5b2669f1987a	Dec 07 2021, 2:20:30 pm	9.77 KB

Restore

Deleting 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 Kubernetes cluster. 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.




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.

Deleting all backup files for a working environment


Deleting all backups 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](#).

Steps

1. From the **Kubernetes** tab, select **Backup Settings**.



2. Click **...** for the Kubernetes cluster 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**.

Deleting all backup files for a volume

Deleting all backups for a volume also disables future backups for that volume.

You can [restart making backups for the volume](#) at any time from the Manage Backups page.

Steps

1. From the **Kubernetes** tab, click **...** for the source volume and select **Details & Backup List**.

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

All Kubernetes Clusters Backup Settings

1
Kubernetes Clusters

57
Protected PVs

15.1 TB
Total Backups Size

Protected Persistent Volumes Status

57
Healthy Backup

0
Failed Backup

57 Backups

Source Kubernetes Cluster	Source Persistent Volume	Source Namespace	Last Backup	Backups	Backup Status
Kubernetes_Cloud_AWS On	Source Persistent Volume On	Source Namespace	May 22 2019, 00:00:00	2,050 Backups	Active
Kubernetes_Cloud_AWS On	Source Persistent Volume On	Source Namespace	May 22 2019, 00:00:00	2,050 Snapshot	<div> <div>Details & Backup List</div> <div>Backup Now</div> <div>Pause Backups</div> </div>
Kubernetes_Cloud_AWS On	Source Persistent Volume On	Source Namespace	May 22 2019, 00:00:00	2,050 Snapshot	

The list of all backup files is displayed.

Source

Working Environment Working Environment N...

Type Cloud Volumes ONTAP (HA)

Provider AWS

Volume Volume Name

SVM SVM Name

Destination

Cloud Provider AWS

Region us-east-1

Bucket netapp-backup

Account ID 012345678901234567890

Backup Information

Relationship Status Active

Last Backup Oct 05 2021, 2:41:33 pm

Lag Duration 14 days 3 hours, 38 mi...

Backups 2,050

Backup Policy Netapp7YearsRetention

2,050 Backups

Select Timeframe Actions

Backup Name	Date	Size
Backup_2020_Jan	May 22 2019, 00:00:00	19,001
Backup_2020_Mar	May 22 2019, 00:00:00	19,002
Backup_2020_Apr	May 22 2019, 00:00:00	19,009

2. Click **Actions** > **Delete all Backups**.

2,050 Backups

Select Timeframe Actions

Backup Name	Date
Backup_2020_Jan	May 22 2019, 00:00:00
Backup_2020_Mar	May 22 2019, 00:00:00

Delete All Backups

Download Backup Report

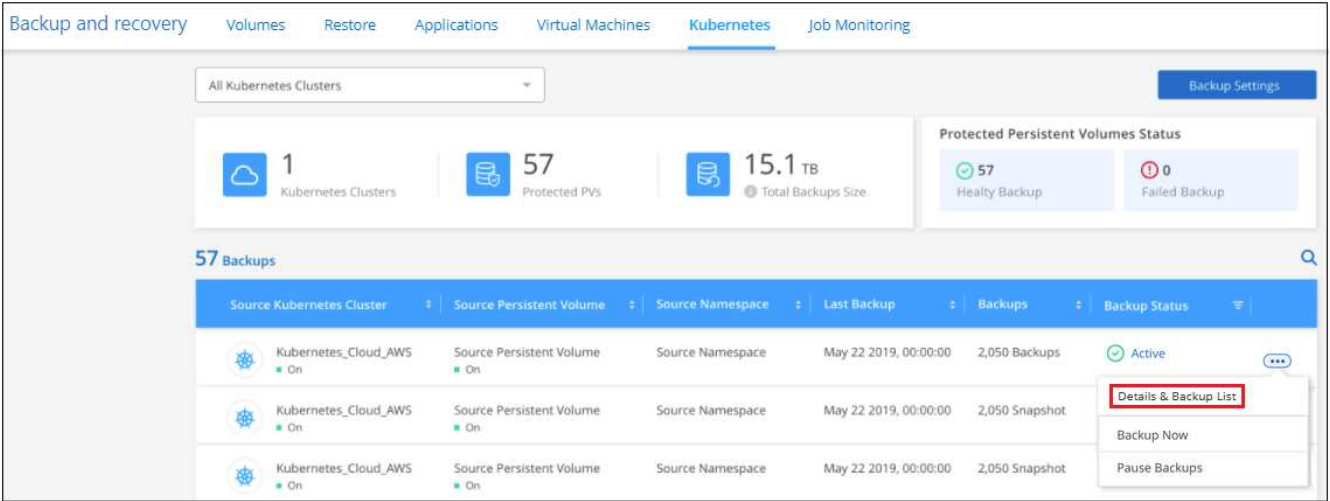
3. In the confirmation dialog box, enter the volume name and click **Delete**.

Deleting a single backup file for a volume

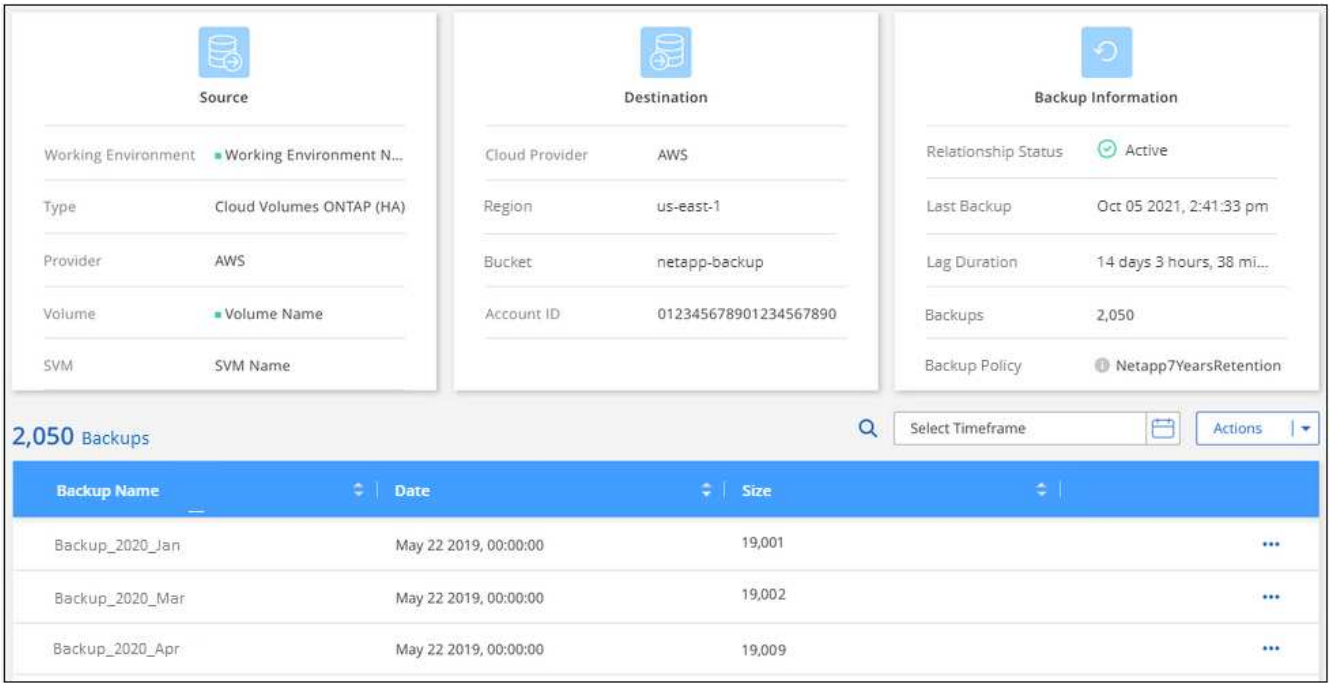
You can delete a single backup file. This feature is available only if the volume backup was created from a system with ONTAP 9.8 or greater.

Steps


- 1. From the **Kubernetes** tab, click **...** for the source volume and select **Details & Backup List**.



The list of all backup files is displayed.



- 2. Click **...** for the volume backup file you want to delete and click **Delete**.



3. In the confirmation dialog box, click **Delete**.


Disabling BlueXP backup and recovery for a working environment

Disabling 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 **Kubernetes** tab, select **Backup Settings**.



2. From the *Backup Settings* page, click **...** for the working environment, or the Kubernetes cluster, 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.

Unregistering 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 Kubernetes cluster, 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 **Kubernetes** tab, select **Backup Settings**.



2. From the *Backup Settings* page, click ... for the Kubernetes cluster where you want to unregister the backup service and select **Unregister**.



3. In the confirmation dialog box, click **Unregister**.

Restoring Kubernetes data from backup files

Backups are stored in an object store in your cloud account so that you can restore data from a specific point in time. You can restore an entire Kubernetes persistent volume from a saved backup file.

You can restore a persistent volume (as a new volume) to the same working environment or to a different working environment that's using the same cloud account.

Supported working environments and object storage providers

You can restore a volume from a Kubernetes backup file to the following working environments:

Backup File Location	Destination Working Environment
Amazon S3	Kubernetes cluster in AWS
Azure Blob	Kubernetes cluster in Azure
Google Cloud Storage	Kubernetes cluster in Google

Restoring volumes from a Kubernetes backup file

When you restore a persistent volume from a backup file, BlueXP creates a *new* volume using the data from the backup. You can restore the data to a volume in the same Kubernetes cluster or to a different Kubernetes cluster that's located in the same cloud account as the source Kubernetes cluster.

Before you start, you should know the name of the volume you want to restore and the date of the backup file you want to use to create the newly restored volume.

Steps

1. From the BlueXP menu, select **Protection > Backup and recovery**.

- Click the **Kubernetes** tab and the Kubernetes Dashboard is displayed.

The screenshot shows the 'Kubernetes' tab selected in the top navigation bar. Below the navigation bar, there's a summary section with three cards: '1 Kubernetes Clusters', '57 Protected PVS', and '15.1 TB Total Backups Size'. To the right, a 'Protected Persistent Volumes Status' section shows '57 Healthy Backup' and '0 Failed Backup'. Below this, a table lists '57 Backups'. The table has columns: Source Kubernetes Cluster, Source Persistent Volume, Source Namespace, Last Backup, Backups, and Backup Status. The first row shows 'Kubernetes_Cloud_AWS' as the source cluster, 'Source Persistent Volume' as the source PVS, 'Source Namespace' as the source namespace, 'May 22 2019, 00:00:00' as the last backup, '2,050 Backups' as the number of backups, and 'Active' as the backup status. A red box highlights the 'Details & Backup List' link in the 'Backup Status' column of the first row.

- Locate the volume you want to restore, click **...**, and then click **Details & Backup List**.

The list of all backup files for that volume is displayed along with details about the source volume, destination location, and backup details.

The screenshot shows the 'Details & Backup List' page. It has three main sections: 'Source', 'Destination', and 'Backup Information'. The 'Source' section shows 'Kubernetes Cluster: eks1', 'Type: EKS', 'Provider: AWS', 'Persistent Volume: pvc-05881c70-cf5f-4edc-8537...', and 'Namespace: default'. The 'Destination' section shows 'Cloud Provider: AWS', 'Bucket: netapp-backup-vsa5twmc9ae', 'Region: us-west-1', and 'Account ID: 123456789012'. The 'Backup Information' section shows 'Relationship Status: enabled', 'Last Backup: Dec 07 2021, 2:20:30 pm', 'Lag Duration: 1 hour', 'Backups: 2', and 'Backup Policy: 24 hourly | 30 daily | 52 weekly'. Below these sections, a table lists '2 Backups'. The table has columns: Backup Name, Date, and Size. The first row shows 'daily.dem-163887957011628bef197-34b5-11ec-8916-5b2669f1987a' as the backup name, 'Dec 07 2021, 2:19:30 pm' as the date, and '9.77 KB' as the size. A red box highlights the 'Restore' button in the 'Backup Name' column of the first row.

- Locate the specific backup file that you want to restore based on the date/time stamp, click **...**, and then **Restore**.
- In the *Select Destination* page, select the *Kubernetes cluster* where you want to restore the volume, the *Namespace*, the *Storage Class*, and the new *Persistent volume name*.

A dialog box titled "Select Destination" with a light gray background. It contains four dropdown menus: "Select Kubernetes Cluster" with "eks1" selected, "Namespace" with "default" selected, "Storage Class" with "basic" selected, and "PVC Name" with "pvc-05881c70-cf5f-4edc-8537-a0a5ce36f9a1-restore" selected. At the bottom are two buttons: "Cancel" (light blue) and "Restore" (dark blue).

Select Destination

Select Kubernetes Cluster

eks1

Namespace

default

Storage Class

basic

PVC Name

pvc-05881c70-cf5f-4edc-8537-a0a5ce36f9a1-restore

Cancel Restore

6. Click **Restore** and you are returned to the Kubernetes Dashboard so you can review the progress of the restore operation.

Result

BlueXP creates a new volume in the Kubernetes cluster based on the backup you selected. You can [manage the backup settings for this new volume](#) as required.

BlueXP backup and recovery APIs

The BlueXP backup and recovery capabilities that are available through the web UI are also available through the RESTful API.

There are ten categories of endpoints defined within BlueXP backup and recovery:

- backup - manages backup operations of cloud and on-premises resources, and retrieves details of the backup data
- catalog - manages the indexed catalog search for files based on a query (Search & Restore)
- cloud - retrieves information about various cloud provider resources from the BlueXP
- job - manages job detail entries on the BlueXP database
- license - retrieves the license validity of the working environments from BlueXP
- ransomware scan - initiates a ransomware scan on a specific backup file
- restore - enables you to perform volume, file, and folder-level restore operations
- sfr - retrieves files from a backup file for single file-level restore operations (Browse & Restore)
- storagegrid - retrieves details about a StorageGRID server, and enables you to discover a StorageGRID server
- working environment - manages the backup policies, and configures the destination object store associated with a working environment

Getting started

To get started with the BlueXP backup and recovery APIs, you'll need to obtain a user token, your BlueXP account ID, and the BlueXP Connector ID.

When making API calls, you'll add the user token in the Authorization header, and the BlueXP Connector ID in the x-agent-id header. You should use the BlueXP account ID in the APIs.

Steps

1. Obtain a user token from the NetApp BlueXP web site.

Make sure you generate the refresh token from the following xref:./ <https://services.cloud.netapp.com/refresh-token/>. The refresh token is an alpha-numeric string that you'll use to generate a user token.

```
curl --location --request POST 'https://netapp-cloud-  
account.auth0.com/oauth/token?=' \\  
--header 'Content-Type: application/json' \\  
-d '{  
  "grant_type": "refresh_token",  
  "refresh_token": "JxaVHn9cGkX92aPVCkhat3zxxxxxwsC9qM1_pLHkZtsVA",  
  "client_id": "Mu0V1ywgYteI6w1MbD15fKfVIUrNXGWC"  
}'
```



The user token from the BlueXP web site has an expiration date. The API response includes an "expires_in" field that states when the token expires. To refresh the token, you'll need to call this API again.

2. Obtain your BlueXP account ID.

```
GET 'https://api.bluexp.netapp.com/tenancy/account' -H 'authority:
api.bluexp.netapp.com'
Header:
-H 'accept: application/json'
-H 'accept-language: en-GB,en;q=0.9'
-H 'authorization: Bearer eyJhbGciOiJSUzI1NiIsInR.....
```

This API will return a response like the following. You can retrieve the account ID by parsing the output from **[0].[accountPublicId]**.

```
[{"accountPublicId":"account-
i6vJXvZW","accountName":"rashidn","isSaas":true,"isGov":false,"isPrivate
PreviewEnabled":false,"is3rdPartyServicesEnabled":false,"accountSerial":
"96064469711530003565","userRole":"Role-1"}].....
```

3. Obtain the x-agent-id which contains the BlueXP Connector ID.

```
GET curl 'https://api.services.cloud.netapp.com/occm/list-occms/account-
OOnAR4ZS?excludeStandalone=true&source=saas' \
Header:
-H 'authority: api.services.cloud.netapp.com' \
-H 'accept: application/json' \
-H 'accept-language: en-GB,en;q=0.9' \
-H 'authorization: Bearer eyJhbGciOiJSUzI1NiIsInR5.....
```

This API will return a response like the following. You can retrieve the agent id by parsing the output from **occm.[0].[agent].[agentId]**.


```
{
  "occms": [
    {
      "account": "account-OOOnAR4ZS",
      "accountName": "cbs",
      "occm": "imEdsEW4HyYTFbt8ZcNKTKDF05jMIe6Z",
      "agentId": "imEdsEW4HyYTFbt8ZcNKTKDF05jMIe6Z",
      "status": "ready",
      "occmName": "cbsgcpdevcntsg-asia",
      "primaryCallbackUri": "http://34.93.197.21",
      "manualOverrideUris": [],
      "automaticCallbackUris": [
        "http://34.93.197.21",
        "http://34.93.197.21/occmui",
        "https://34.93.197.21",
        "https://34.93.197.21/occmui",
        "http://10.138.0.16",
        "http://10.138.0.16/occmui",
        "https://10.138.0.16",
        "https://10.138.0.16/occmui",
        "http://localhost",
        "http://localhost/occmui",
        "http://localhost:1337",
        "http://localhost:1337/occmui",
        "https://localhost",
        "https://localhost/occmui",
        "https://localhost:1337",
        "https://localhost:1337/occmui"
      ],
      "createDate": "1652120369286",
      "agent": {
        "useDockerInfra": true,
        "network": "default",
        "name": "cbsgcpdevcntsg-asia",
        "agentId": "imEdsEW4HyYTFbt8ZcNKTKDF05jMIe6Zclients",
        "provider": "gcp",
        "systemId": "a3aa3578-bfee-4d16-9e10-"
      }
    }
  ]
}
```

Example using the APIs

The following example shows an API call to activate BlueXP backup and recovery on a working environment with a new policy that has daily, hourly, and weekly labels set, archive after days set to 180 days, in East-US-2 region in Azure cloud. Note that this only enables backup on the working environment, but no volumes are backed up.

API Request

You'll see that we use the BlueXP account ID `account-DpTFcxN3`, BlueXP Connector ID `iZwFFeVCZjWnzGlw8RgD0QQNANZvpP7Iclients`, and user token `Bearer eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IksrSXlPVFUzUWpZek1E...y6nyhBjwkeMwHc4ValobjUmju2x0xUH48g` in this command.

```

curl --location --request POST
'https://api.blueexp.netapp.com/account/account-
DpTFcxN3/providers/cloudmanager_cbs/api/v3/backup/working-
environment/VsaWorkingEnvironment-99hPYEgk' \
--header 'x-agent-id: iZwFFeVCZjWnzGlw8RgD0QQNANZvpP7Iclients' \
--header 'Accept: application/json' \
--header 'Content-Type: application/json' \
--header 'Authorization: Bearer
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6Iks5rSXlPVFUzUWpZek1E...y6nyhBjwk
eMwHc4ValobjUmju2x0xUH48g' \
--data-raw '{
  "provider": "AZURE",
  "backup-policy": {
    "archive-after-days": 180,
    "rule": [
      {
        "label": "hourly",
        "retention": "2"
      },
      {
        "label": "daily",
        "retention": "30"
      },
      {
        "label": "weekly",
        "retention": "52"
      }
    ]
  },
  "ip-space": "Default",
  "region": "eastus2",
  "azure": {
    "resource-group": "rn-test-backup-rg",
    "subscription": "3beb4dd0-25d4-464f-9bb0-303d7cf5c0c2"
  }
}'

```

Response is a job ID that you can then monitor.

```

{
  "job-id": "1b34b6f6-8f43-40fb-9a52-485b0dfe893a"
}

```

Monitor the response.

```
curl --location --request GET
'https://api.bluexp.netapp.com/account/account-
DpTFcxN3/providers/cloudmanager_cbs/api/v1/job/1b34b6f6-8f43-40fb-9a52-
485b0dfe893a' \
--header 'x-agent-id: iZwFFeVCZjWnzGlw8RgD0QQNANZvpP7Iclients' \
--header 'Accept: application/json' \
--header 'Content-Type: application/json' \
--header 'Authorization: Bearer
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6Ikp5rSXlPVFUzUWpZek1E...hE9ss2Nub
K6wZRHUdSaORI7JvcOorUhJ8srqdiUiW6MvuGIFAQIh668of2M3dLbhVDBe8BBMtsa939UGnJx
7Qz6Eg'
```

Response.

```
{
  "job": [
    {
      "id": "1b34b6f6-8f43-40fb-9a52-485b0dfe893a",
      "type": "backup-working-environment",
      "status": "PENDING",
      "error": "",
      "time": 1651852160000
    }
  ]
}
```

Monitor until "status" is "COMPLETED".

```
{
  "job": [
    {
      "id": "1b34b6f6-8f43-40fb-9a52-485b0dfe893a",
      "type": "backup-working-environment",
      "status": "COMPLETED",
      "error": "",
      "time": 1651852160000
    }
  ]
}
```

API reference

Documentation for each BlueXP backup and recovery API is available from <https://docs.netapp.com/us-en/>

Reference

BlueXP backup and recovery policy configuration settings

BlueXP backup and recovery enables you to create backup policies with a variety of configuration 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 schedules

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.

Name	Default_Policy_Name
Labels & Retention	
12 Labels	Selected Labels (2) (Select up to 5 Labels)
<input checked="" type="checkbox"/> Hourly	Hourly Number of Backups to Retain 12
<input checked="" type="checkbox"/> Daily	Daily Number of Backups to Retain 30
<input type="checkbox"/> Weekly	
<input type="checkbox"/> Monthly	
<input type="checkbox"/> Yearly	
DataLock & Ransomware Protection	None
Archival Policy	Disabled

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

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 DataLock and Ransomware Protection 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 must enable it when creating your first backup policy and activating BlueXP backup and recovery for that cluster.
- DataLock and Ransomware protection can't be disabled for a cluster once it has been configured. Do not enable this feature on a cluster to try it out.
- 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 14 day buffer. Any DataLock retention policy that is less than 30 days is rounded up to 30 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 14 days) 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 14 days buffer), the "7 weekly" backups would be retained for 9 weeks (7 weeks plus 14 days), and the "12 monthly" backups would be retained for 12 months (plus 14 days).
- 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 14 days buffer).

You can see in this last case that if each backup file is locked for 44 days, 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 3 points in the backup and restore process:

- When a backup file is created

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 settings

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

AWS

DataLock & Ransomware Protection

Backup copies are protected from being modified or deleted, and they are scanned for ransomware threats.

- ☒ **None**
- ☐ **Governance**
Users with specific permissions can overwrite or delete protected backup files during the retention period
- ☐ **Compliance**
No users can overwrite or delete protected backup files during the retention period

Azure

DataLock & Ransomware Protection

Backup copies are protected from being modified or deleted, and they are scanned for ransomware threats.

- ☒ **None**
- ☐ **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.
- ☐ **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.

StorageGRID

DataLock & Ransomware Protection

Backup copies are protected from being modified or deleted, and they are scanned for ransomware threats.

- ☒ **None**
- ☐ **Compliance**
No users can overwrite or delete protected backup files during the retention period

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-*":
 - 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:
 - 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

- DataLock and Ransomware protection 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't enable this feature after BlueXP backup and recovery has been activated.
- FlexGroup volumes can use DataLock and Ransomware protection only when using ONTAP 9.13.1 or greater.

Archival storage settings

When using AWS, Azure, or Google cloud storage you can move older backup files to a less expensive archival 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. 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.



- 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	<div>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.</div> <div><input checked="" type="checkbox"/> Tier Backups to Archive</div> <div>Archive After (Days) <input type="text" value="30"/></div> <div>Storage Class <input type="text" value="S3 Glacier"/></div>	

- 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.](#)

AWS S3 archival storage classes and restore retrieval times

BlueXP backup and recovery supports two S3 archival storage classes and most regions.

Supported S3 archival storage classes for BlueXP backup and recovery

When backup files are initially created they're stored in *S3 Standard* storage. This tier is optimized for storing data that's infrequently accessed; but that also allows you to access it immediately. After 30 days the backups transition to the *S3 Standard-Infrequent Access* storage class to save on costs.

If your source clusters are running ONTAP 9.10.1 or greater, you can choose to tier backups to either *S3 Glacier* or *S3 Glacier Deep Archive* storage after a certain number of days (typically more than 30 days) for further cost optimization. Data in these tiers can't be accessed immediately when needed, and will require a higher retrieval cost, so you need to consider how often you may need to restore data from these archived backup files. See the section about [restoring data from 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.

Note that when you configure BlueXP backup and recovery with this type of lifecycle rule, you must not configure any lifecycle rules when setting up the bucket in your AWS account.

[Learn about S3 storage classes.](#)

Restoring data from archival storage

While storing older backup files in archival storage is much less expensive than Standard or Standard-IA storage, accessing data from a backup file in archive storage for restore operations will take a longer amount of time and will cost more money.

How much does it cost to restore data from Amazon S3 Glacier and Amazon S3 Glacier Deep Archive?

There are 3 restore priorities you can choose when retrieving data from Amazon S3 Glacier, and 2 restore priorities when retrieving data from Amazon S3 Glacier Deep Archive. S3 Glacier Deep Archive costs less than S3 Glacier:

Archive Tier	Restore Priority & Cost		
	High	Standard	Low
S3 Glacier	Fastest retrieval, highest cost	Slower retrieval, lower cost	Slowest retrieval, lowest cost
S3 Glacier Deep Archive		Faster retrieval, higher cost	Slower retrieval, lowest cost

Each method has a different per-GB retrieval fee and per-request fee. For detailed S3 Glacier pricing by AWS Region, visit the [Amazon S3 pricing page](#).

How long will it take to restore my objects archived in Amazon S3 Glacier?

There are 2 parts that make up the total restore time:

- **Retrieval time:** The time to retrieve the backup file from archive and place it in Standard storage. This is sometimes called the "rehydration" time. The retrieval time is different depending on the restore priority you choose.

Archive Tier	Restore Priority & Retrieval Time		
	High	Standard	Low
S3 Glacier	3-5 minutes	3-5 hours	5-12 hours
S3 Glacier Deep Archive		12 hours	48 hours

- **Restore time:** The time to restore the data from the backup file in Standard storage. This time is no different than the typical restore operation directly from Standard storage - when not using an archival tier.

For more information about Amazon S3 Glacier and S3 Glacier Deep Archive retrieval options, refer to [the Amazon FAQ about these storage classes](#).

Azure archival tiers and restore retrieval times

BlueXP backup and recovery supports one Azure archival access tier and most regions.

Supported Azure Blob access tiers for BlueXP backup and recovery

When backup files are initially created they're stored in the *Cool* access tier. This tier is optimized for storing data that's infrequently accessed; but when needed, can be accessed immediately.

If your source clusters are running ONTAP 9.10.1 or greater, you can choose to tier backups from *Cool* to *Azure Archive* storage after a certain number of days (typically more than 30 days) for further cost optimization. Data in this tier can't be accessed immediately when needed, and will require a higher retrieval cost, so you need to consider how often you may need to restore data from these archived backup files. See the next section about [restoring data from archival storage](#).

Note that when you configure BlueXP backup and recovery with this type of lifecycle rule, you must not configure any lifecycle rules when setting up the container in your Azure account.

[Learn about Azure Blob access tiers.](#)

Restoring data from archival storage

While storing older backup files in archival storage is much less expensive than Cool storage, accessing data from a backup file in Azure Archive for restore operations will take a longer amount of time and will cost more money.

How much does it cost to restore data from Azure Archive?

There are two restore priorities you can choose when retrieving data from Azure Archive:

- **High:** Fastest retrieval, higher cost
- **Standard:** Slower retrieval, lower cost

Each method has a different per-GB retrieval fee and per-request fee. For detailed Azure Archive pricing by Azure Region, visit the [Azure pricing page](#).



The High priority is not supported when restoring data from Azure to StorageGRID systems.

How long will it take to restore my data archived in Azure Archive?

There are 2 parts that make up the restore time:

- **Retrieval time:** The time to retrieve the archived backup file from Azure Archive and place it in Cool storage. This is sometimes called the "rehydration" time. The retrieval time is different depending on the restore priority you choose:
 - **High:** < 1 hour
 - **Standard:** < 15 hours
- **Restore time:** The time to restore the data from the backup file in Cool storage. This time is no different than the typical restore operation directly from Cool storage - when not using an archival tier.

For more information about Azure Archive retrieval options, refer to [this Azure FAQ](#).

Google archival storage classes and restore retrieval times

BlueXP backup and recovery supports one Google archival storage class and most regions.

Supported Google archival storage classes for BlueXP backup and recovery

When backup files are initially created they're stored in *Standard* storage. This tier is optimized for storing data that's infrequently accessed; but that also allows you to access it immediately.

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 (typically more than 30 days) for further cost optimization. Data in this tier will require a higher retrieval cost, so you need to consider how often you may need to restore data from these archived backup files. See the section about [restoring data from archival storage](#).

Note that when you configure BlueXP backup and recovery with this type of lifecycle rule, you must not configure any lifecycle rules when setting up the bucket in your Google account.

[Learn about Google storage classes.](#)

Restoring data from archival storage

While storing older backup files in Archive storage is much less expensive than Standard storage, accessing data from a backup file in Archive storage for restore operations will take a slightly longer amount of time and will cost more money.

How much does it cost to restore data from Google Archive?

For detailed Google Cloud Storage pricing by region, visit the [Google Cloud Storage pricing page](#).

How long will it take to restore my objects archived in Google Archive?

There are 2 parts that make up the total restore time:

- **Retrieval time:** The time to retrieve the backup file from Archive and place it in Standard storage. This is sometimes called the "rehydration" time. Unlike the "coldest" storage solutions provided by other cloud providers, your data is accessible within milliseconds.
- **Restore time:** The time to restore the data from the backup file in Standard storage. This time is no different than the typical restore operation directly from Standard storage - when not using an archival tier.

Configure backup for multi-account access in Azure

BlueXP backup and recovery enables you to create backup files in an Azure account that is different than where your source Cloud Volumes ONTAP volumes reside. Both of those accounts can be different than the account where the BlueXP Connector resides.

These steps are required only when you are [backing up Cloud Volumes ONTAP data to Azure Blob storage](#).

Just follow the steps below to set up your configuration in this manner.

Set up VNet peering between accounts

Note that if you want BlueXP to manage your Cloud Volumes ONTAP system in a different account/region, then you need to setup VNet peering. VNet peering is not required for storage account connectivity.

1. Log in to the Azure portal and from home, select Virtual Networks.
2. Select the subscription you are using as subscription 1 and click on the VNet where you want to set up peering.

Home > Virtual networks

NetApp HCL (netapphcl.onmicrosoft.com)

+ New Manage view Refresh Export to CSV Open query Assign tags Feedback

Filter for any field... Subscription == OCCM Dev Resource group == all Location == all Add filter

Showing 1 to 60 of 60 records.

<input type="checkbox"/> Name ↑↓	Resource group ↑↓	Location ↑↓
<input type="checkbox"/> cbsnetwork	occm_group_eastasia	East Asia
<input type="checkbox"/> Vnet1	occm_group_australiaeast	Australia East
<input type="checkbox"/> Vnet1	occm_group_australiasoutheast	Australia Southeast

3. Select **cbsnetwork** and from the left panel, click on **Peerings**, and then click **Add**.

Subscription * ⓘ

OCCM Automation

Virtual network *

cbse2evnet

Traffic to remote virtual network ⓘ

☒ Allow (default)

☐ Block all traffic to the remote virtual network

Traffic forwarded from remote virtual network ⓘ

☒ Allow (default)

☐ Block traffic that originates from outside this virtual network

Virtual network gateway or Route Server ⓘ

☐ Use this virtual network's gateway or Route Server

☐ Use the remote virtual network's gateway or Route Server

☒ None (default)

Add

4. Enter the following information on the Peering page and then click **Add**.

- Peering link name for this network: you can give any name to identify the peering connection.
- Remote virtual network peering link name: enter a name to identify the remote VNet.
- Keep all the selections as default values.
- Under subscription, select the subscription 2.
- Virtual network, select the virtual network in subscription 2 to which you want to set up the peering.

The screenshot shows the 'cbsnetwork | Peerings' page in the Azure portal. The left sidebar contains navigation links: Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, and Settings. The main content area has a search bar and a table of peerings. The table has three columns: Name, Peering status, and Peer. One peering is listed with the name 'cbsnetwork', status 'Connected', and peer 'cbse2evnet'.

Name	Peering status	Peer
cbsnetwork	Connected	cbse2evnet

5. Perform the same steps in subscription 2 VNet and specify the subscription and remote VNet details of subscription 1.

Subscription * ⓘ

OCCM Dev

Virtual network *

cbsnetwork

Traffic to remote virtual network ⓘ

☒ Allow (default)

☐ Block all traffic to the remote virtual network

Traffic forwarded from remote virtual network ⓘ

☒ Allow (default)

☐ Block traffic that originates from outside this virtual network

Virtual network gateway or Route Server ⓘ

☐ Use this virtual network's gateway or Route Server

☐ Use the remote virtual network's gateway or Route Server

☒ None (default)

Add

The peering settings are added.

cbse2evnet | Peerings ...

Virtual network

Search (Cmd+ /) << + Add ↻ Refresh

Filter by name...

Name	Peering status	Peer
cbsnetworkpeer	Connected	cbsnetwork

Overview

Activity log

Access control (IAM)

Tags

Diagnose and solve problems

Settings

Address space

Connected devices

Subnets

DDoS protection

Firewall

Security

DNS servers

Peerings

Create a private endpoint for the storage account

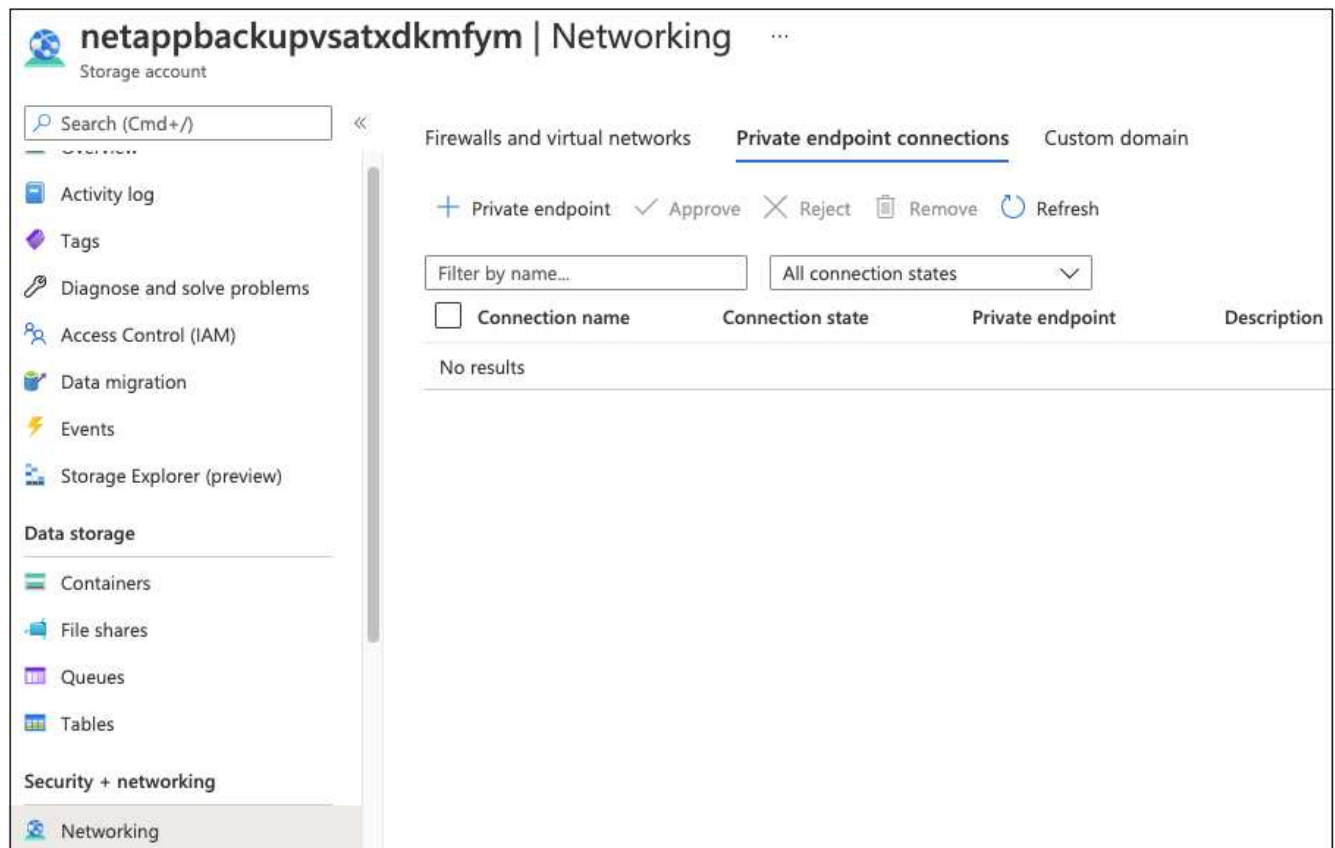
Now you need to create a private endpoint for the storage account. In this example, the storage account is created in subscription 1 and the Cloud Volumes ONTAP system is running in subscription 2.



You need network contributor permission to perform the following action.

```
{
  "id": "/subscriptions/d333af45-0d07-4154-943dc25fbbce1b18/providers/Microsoft.Authorization/roleDefinitions/4d97b98b-1d4f-4787-a291-c67834d212e7",
  "properties": {
    "roleName": "Network Contributor",
    "description": "Lets you manage networks, but not access to them.",
    "assignableScopes": [
      "/"
    ],
    "permissions": [
      {
        "actions": [
          "Microsoft.Authorization/*/read",
          "Microsoft.Insights/alertRules/*",
          "Microsoft.Network/*",
          "Microsoft.ResourceHealth/availabilityStatuses/read",
          "Microsoft.Resources/deployments/*",
          "Microsoft.Resources/subscriptions/resourceGroups/read",
          "Microsoft.Support/*"
        ],
        "notActions": [],
        "dataActions": [],
        "notDataActions": []
      }
    ]
  }
}
```

1. Go to the Storage account > Networking > Private endpoint connections and click **+ Private endpoint**.



2. In the Private Endpoint *Basics* page:

- Select subscription 2 (where the BlueXP Connector and Cloud Volumes ONTAP system are deployed) and the resource group.
- Enter an endpoint name.
- Select the region.

Create a private endpoint

1 Basics 2 Resource 3 Configuration 4 Tags 5 Review + create

Use private endpoints to privately connect to a service or resource. Your private endpoint must be in the same region as your virtual network, but can be in a different region from the private link resource that you are connecting to. [Learn more](#)

Project details

Subscription * ⓘ OCCM Dev

Resource group * ⓘ cbsoccmdevcvo-rg [Create new](#)

Instance details

Name * cbse2e ✓

Region * (Asia Pacific) East Asia

3. In the *Resource* page, select Target sub-resource as **blob**.

Create a private endpoint ...

✓ Basics **2 Resource** 3 Configuration 4 Tags 5 Review + create

Private Link offers options to create private endpoints for different Azure resources, like your private link service, a SQL server, or an Azure storage account. Select which resource you would like to connect to using this private endpoint. [Learn more](#)

Subscription	OCCM Dev (d333af45-0d07-4154-943d-c25fbbce1b18)
Resource type	Microsoft.Storage/storageAccounts
Resource	test150521
Target sub-resource * ⓘ	blob

4. In the Configuration page:

- Select the virtual network and subnet.
- Click the **Yes** radio button to "Integrate with private DNS zone".

Create a private endpoint ...

✓ Basics ✓ Resource **3 Configuration** 4 Tags 5 Review + create

Networking

To deploy the private endpoint, select a virtual network subnet. [Learn more](#)

Virtual network * ⓘ	cbsnetwork
Subnet * ⓘ	default (10.2.0.0/24)

i If you have a network security group (NSG) enabled for the subnet above, it will be disabled for private endpoints on this subnet only. Other resources on the subnet will still have NSG enforcement.

Private DNS integration

To connect privately with your private endpoint, you need a DNS record. We recommend that you integrate your private endpoint with a private DNS zone. You can also utilize your own DNS servers or create DNS records using the host files on your virtual machines. [Learn more](#)

Integrate with private DNS zone ☒ Yes ☐ No

Configuration name	Subscription	Private DNS zone
privatelink-blob-core-...	OCCM Dev	privatelink.blob.core.windows.net

Review + create < Previous Next : Tags >

5. In the Private DNS zone list, ensure that the Private Zone is selected from the correct Region, and click **Review + Create**.

Configuration name	Subscription	Private DNS zone
privatelink-blob-core-...	OCCM Dev	privatelink.blob.core.windows.net
		<input type="text" value="Filter private DNS zones"/> <ul style="list-style-type: none"> occm_group_centralus privatelink.blob.core.windows.net occm_group_eastus privatelink.blob.core.windows.net occm_group_eastus2 privatelink.blob.core.windows.net

Now the storage account (in subscription 1) has access to the Cloud Volumes ONTAP system which is running in subscription 2.

6. Retry enabling BlueXP backup and recovery on the Cloud Volumes ONTAP system and this time it should be successful.

Restore BlueXP backup and recovery data in a dark site

When using BlueXP backup and recovery in a site with no internet access, known as "private mode", the BlueXP backup and recovery configuration data is backed up to the StorageGRID or ONTAP S3 bucket where your backups are being stored. If you have an issue with the BlueXP Connector host system in the future, you can deploy a new Connector and restore the critical BlueXP backup and recovery data.

Note that when you use BlueXP backup and recovery in a SaaS environment where the BlueXP Connector is deployed at your cloud provider, or on your own host system that has internet access, all the important BlueXP backup and recovery configuration data is backed up and protected in the cloud. If you have an issue with the Connector, just create a new Connector and add your working environments and the backup details are automatically restored.

There are 2 types of data that are backed up:

- BlueXP backup and recovery database - contains a listing of all the volumes, backup files, backup policies, and configuration information.
- Indexed Catalog files - contains detailed indexes that are used for Search & Restore functionality that make your searches very quick and efficient when looking for volume data that you want to restore.

This data is backed up once per day at midnight, and a maximum of 7 copies of each file are retained. If the Connector is managing multiple on-premises ONTAP working environments, the BlueXP backup and recovery files will be located in the bucket of the working environment that was activated first.



No volume data is ever included in the BlueXP backup and recovery database or Indexed Catalog files.

Restore BlueXP backup and recovery data to a new Connector

If your on-premises Connector has a catastrophic failure, you'll need to install a new Connector, and then restore the BlueXP backup and recovery data to the new Connector.

There are 4 tasks you'll need to perform to return your BlueXP backup and recovery system to a working state:

- Install a new BlueXP Connector
- Restore the BlueXP backup and recovery database
- Restore the Indexed Catalog files
- Rediscover all of your on-prem ONTAP systems and StorageGRID systems to the BlueXP UI

Once you verify that your system is back in a working order, we recommend that you create new backup files.

What you'll need

You'll need to access the most recent database and index backups from the StorageGRID or ONTAP S3 bucket where your backup files are being stored:

- BlueXP backup and recovery MySQL database file

This file is located in the following location in the bucket `netapp-backup-<GUID>/mysql_backup/`, and it is named `CBS_DB_Backup_<day>_<month>_<year>.sql`.

- Indexed Catalog backup zip file

This file is located in the following location in the bucket `netapp-backup-<GUID>/catalog_backup/`, and it is named `Indexed_Catalog_DB_Backup_<db_name>_<day>_<month>_<year>.zip`.

Install a new Connector on a new on-premises Linux host

When installing a new BlueXP Connector, make sure you download the same release of software as you had installed on the original Connector. Periodic changes to the BlueXP backup and recovery database structure may make newer software releases incompatible with the original database backups. You can [upgrade the Connector software to the most current version after restoring the Backup database](#).

1. [Install the BlueXP Connector on a new on-premises Linux host](#)
2. Log into BlueXP using the admin user credentials that you just created.

Restore the BlueXP backup and recovery database

1. Copy the MySQL backup from the backup location to the new Connector host. We'll use the example file name "CBS_DB_Backup_23_05_2023.sql" below.
2. Copy the backup into the MySQL docker container using the following command:

```
docker cp CBS_DB_Backup_23_05_2023.sql ds_mysql_1:/. 
```

3. Enter the MySQL container shell using the following command:


```
docker exec -it ds_mysql_1 sh
```

4. In the container shell, deploy the "env".
5. You'll need the MySQL DB password, so copy the value of the key "MYSQL_ROOT_PASSWORD".
6. Restore the BlueXP backup and recovery MySQL DB using the following command:

```
mysql -u root -p cloud_backup < CBS_DB_Backup_23_05_2023.sql
```

7. Verify that the BlueXP backup and recovery MySQL DB has been restored correctly using the following SQL commands:

```
mysql -u root -p cloud_backup
```

Enter the password.

```
mysql> show tables;  
mysql> select * from volume;
```

Check if the volumes that are shown are the same as those that existed in your original environment.

Restore the Indexed Catalog files

1. Copy the Indexed Catalog backup zip file (we'll use the example file name "Indexed_Catalog_DB_Backup_catalogdb1_23_05_2023.zip") from the backup location to the new Connector host in the "/opt/application/netapp/cbs" folder.
2. Unzip the "Indexed_Catalog_DB_Backup_catalogdb1_23_05_2023.zip" file using the following command:

```
unzip Indexed_Catalog_DB_Backup_catalogdb1_23_05_2023.zip -d catalogdb1
```

3. Run the **ls** command to make sure that the folder "catalogdb1" has been created with the subfolders "changes" and "snapshots" underneath.

Discover your ONTAP clusters and StorageGRID systems

1. [Discover all the on-prem ONTAP working environments](#) that were available in your previous environment. This includes the ONTAP system you have used as an S3 server.
2. [Discover your StorageGRID systems](#).

Set up the StorageGRID environment details

Add the details of the StorageGRID system associated with your ONTAP working environments as they were set up on the original Connector setup using the [BlueXP APIs](#).

You'll need to perform these steps for each ONTAP system that is backing up data to StorageGRID.

1. Extract the authorization token using the following oauth/token API.

```
curl 'http://10.193.192.202/oauth/token' -X POST -H 'User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10.15; rv:100101 Firefox/108.0' -H 'Accept: application/json' -H 'Accept-Language: en-US,en;q=0.5' -H 'Accept-Encoding: gzip, deflate' -H 'Content-Type: application/json' -d '{"username":admin@netapp.com,"password":"Netapp@123","grant_type":"password"}'
> '
```

This API will return a response like the following. You can retrieve the authorization token as shown below.

```
{
  "expires_in": 21600,
  "access_token": "eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjRiIn0eyJzdWIiOiJvY2NtYXV0aHwxIiwiaXVkiJjpbImh0dHBzOi8vYXBpLmNsb3VkLm5ldGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm5ldGFwcC5jb20vZnVsbnB9YyW1lIjoiiYWRtaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS9lbWFPbCI6ImFkbWluQG5ldGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwcm9maWxlIiwiaWF0IjoxNjc5NzM2MDIzLCJleHAiOjE2NzI3NTc2MjMsIm1zcyI6Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9CjRrRDY23PokyLg1Ii67bmgnMcyXdcvBOY-ZUYWzhrWbbY_hqUH4T-114v_pNDsPyNDyWqHaKizThdjJHYHxm56vTz_Vdn4NqjaBDPwN9KANc6Z88WA1cJ4WRQqj5ykODNDmrv5At_f9HHp0-xVMYHqywZ4nNFaIMvAh4xESc5jfoKOZc-IOQdWm4F4LHpMzs4qFzCYthTuSKLYtqSTUrZB81-o-ipvrOqSolIwIeHXZJJV-Uswun9daNgiYd_wX-4WWJViGEnDzzwOKfUoUoe1Fg3ch--7JfKf1-rrXDOjklSUmumN3WHV9usplPgBE5HAcJPrEBm0ValSZcUbiA"}
}
```

2. Extract the Working Environment ID and the X-Agent-Id using the `tenancy/external/resource` API.

```
curl -X GET
http://10.193.192.202/tenancy/external/resource?account=account-
DARKSITE1 -H 'accept: application/json' -H 'authorization: Bearer
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjZjRiInOeyJzdWIiOiJvY
2NtYXV0aHwxiwiYXVkiJpbImh0dHBzOi8vYXBpLmNsb3VkLm5ldGFwCj5jb20iXSwiaHR0c
DovL2Nsb3VkLm5ldGFwCj5jb20vZnVsbnF9uYW11IjoieYWRtaW4iLCJodHRwOi8vY2xvdWQub
mV0YXBwLmNvbS9lbWFPbCI6ImFkbWluQG5ldGFwCj5jb20iLCJzY29wZSI6Im9wZW5pZCBw
c29maWxlIiwiaWF0IjoxNjc5NzIyNzEzNDQzMjM5ImlzcyI6Imh0dHA6L
y9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdu_kN-
fLWpdJjX98HODwPpVUitLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-
yE0fH9gr4XgkdsWjWcNvw2rRkFzjHpWrETgfgAmkZcAukV4DHuxogHWh6-
DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zC-
sp81GaqMahPf0KcFVyjbBL4krOewgKHGFo_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-
KbZqmmBX9vDnYp7SSxC1hHJRDStcFgJLdJHtowweNH2829KsjEGBTtcBdO8SvIdtctNH_GAX
wSqMT3zUfwaOimPw'
```

This API will return a response like the following. The value under the "resourceIdentifier" denotes the *WorkingEnvironment Id* and the value under "agentId" denotes *x-agent-id*.

3. Update the BlueXP backup and recovery database with the details of the StorageGRID system associated with the Working Environments. Make sure to enter the Fully Qualified Domain Name of the StorageGRID, as well as the Access-Key and Storage-Key as shown below:

```
curl -X POST 'http://10.193.192.202/account/account-DARKSITE1/providers/cloudmanager_cbs/api/v1/sg/credentials/working-environment/OnPremWorkingEnvironment-pMtZND0M' \
> --header 'authorization: Bearer eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjRiIn0eyJzdWIiOiJvY2NtYXV0aHwxIiwiaXVkaW50dHBzOi8vYXBpLmNsb3VkLm5ldGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm5ldGFwcC5jb20vZnVsbF9uYW11IjoieWRtaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS9lbWVpY29wZSI6Im9wZW5pZCBwcM9maWx1IiwiaWF0IjoxNjcyNzIyNzEzNDQzMjM5Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdu_kN-fLWpdJJX98HODwPpVUitLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-yE0fH9gr4XgkdsWjWcNvw2rRkFzjHpWrETgfqAMkZcAukV4DHuxogHWh6-DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zC-sp81GaqMahPf0KcFVybBL4krOewgKHGfo_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-KbZqmmBX9vDnYp7SSxC1hHJRdstcFgJLdJHtowweNH2829KsjEGBTtcBd08SvIDtctNH_GAxwSgMT3zUfwaOimPw' \
> --header 'x-agent-id: vB_1xShPpBtUosjD7wfBlLIhqDgIPA0wclients' \
> -d '{ "storage-server" : "sr630ip15.rtp.eng.netapp.com:10443", "access-key": "2ZMYOAVAS5E70MCNH9", "secret-password": "uk/6ikd4LjlXQOFnzSzP/T0zR4ZQlG0w1xgWsB" }'
```

Verify BlueXP backup and recovery settings

1. Select each ONTAP working environment and click **View Backups** next to the Backup and recovery service in the right-panel.

You should be able to see all the backups that have been created for your volumes.

2. From the Restore Dashboard, under the Search & Restore section, click **Indexing Settings**.

Make sure that the working environments which had Indexed Cataloging enabled previously remain enabled.

3. From the Search & Restore page, run a few catalog searches to confirm that the Indexed Catalog restore has been completed successfully.

Restart the BlueXP backup and recovery service

There may be situations where you'll need to restart the BlueXP backup and recovery service.

BlueXP backup and recovery functionality is built into the BlueXP Connector. You'll need to follow different initial steps to restart the service depending on whether you deployed the Connector in the cloud or whether you installed the Connector manually on a Linux system.

Steps

- 1. Connect to the Linux system that the Connector is running on.

Connector location	Procedure
Cloud deployment	Follow the instructions for connecting to the Connector Linux virtual machine depending on the cloud provider you're using.
Manual installation	Log in to the Linux system.

- 2. Enter the command to restart the service.

Connector location	Command
Cloud deployment	<code>docker restart cloudmanager_cbs</code>
Manual installation with internet access	<code>docker restart cloudmanager_cbs</code>
Manual installation without internet access	<code>docker restart ds_cloudmanager_cbs_1</code>

Knowledge and support

Register for support

Support registration is required to receive technical support specific to BlueXP and its storage solutions and services. Support registration is also required to enable key workflows for Cloud Volumes ONTAP systems.

Registering for support does not enable NetApp support for a cloud provider file service. For technical support related to a cloud provider file service, its infrastructure, or any solution using the service, refer to "Getting help" in the BlueXP documentation for that product.

- [Amazon FSx for ONTAP](#)
- [Azure NetApp Files](#)
- [Cloud Volumes Service for Google Cloud](#)

Support registration overview

There are two forms of registration to activate support entitlement:

- Registering your BlueXP account ID support subscription (your 20 digit 960xxxxxxx serial number located on the Support Resources page in BlueXP).

This serves as your single support subscription ID for any service within BlueXP. Each BlueXP account-level support subscription must be registered.

- Registering the Cloud Volumes ONTAP serial numbers associated with a subscription in your cloud provider's marketplace (these are 20 digit 909201xxxxxxx serial numbers).

These serial numbers are commonly referred to as *PAYGO serial numbers* and get generated by BlueXP at the time of Cloud Volumes ONTAP deployment.

Registering both types of serial numbers enables capabilities like opening support tickets and automatic case generation. Registration is completed by adding NetApp Support Site (NSS) accounts to BlueXP as described below.

Register your BlueXP account for NetApp support

To register for support and activate support entitlement, one user in your BlueXP account must associate a NetApp Support Site account with their BlueXP login. How you register for NetApp support depends on whether you already have a NetApp Support Site (NSS) account.

Existing customer with an NSS account


If you're a NetApp customer with an NSS account, you simply need to register for support through BlueXP.

Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.
2. Select **User Credentials**.

3. Select **Add NSS credentials** and follow the NetApp Support Site (NSS) Authentication prompt.
4. To confirm that the registration process was successful, select the Help icon, and select **Support**.

The **Resources** page should show that your account is registered for support.



Note that other BlueXP users will not see this same support registration status if they have not associated a NetApp Support Site account with their BlueXP login. However, that doesn't mean that your BlueXP account is not registered for support. As long as one user in the account has followed these steps, then your account has been registered.

Existing customer but no NSS account

If you're an existing NetApp customer with existing licenses and serial numbers but *no* NSS account, you need to create an NSS account and associate it with your BlueXP login.

Steps


1. Create a NetApp Support Site account by completing the [NetApp Support Site User Registration form](#)
 - a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
 - b. Be sure to copy the BlueXP account serial number (960xxxx) used above for the serial number field. This will speed up the account processing.
2. Associate your new NSS account with your BlueXP login by completing the steps under [Existing customer with an NSS account](#).

Brand new to NetApp



If you are brand new to NetApp and you don't have an NSS account, follow each step below.

Steps

1. In the upper right of the BlueXP console, select the Help icon, and select **Support**.



2. Locate your account ID serial number from the Support Registration page.

 96015585434285107893 Account serial number	 Not Registered Add your NetApp Support Site (NSS) credentials to BlueXP Follow these instructions to register for support in case you don't have an NSS account yet.
--	--

3. Navigate to [NetApp's support registration site](#) and select **I am not a registered NetApp Customer**.
4. Fill out the mandatory fields (those with red asterisks).
5. In the **Product Line** field, select **Cloud Manager** and then select your applicable billing provider.
6. Copy your account serial number from step 2 above, complete the security check, and then confirm that you read NetApp's Global Data Privacy Policy.

An email is immediately sent to the mailbox provided to finalize this secure transaction. Be sure to check your spam folders if the validation email doesn't arrive in few minutes.

7. Confirm the action from within the email.

Confirming submits your request to NetApp and recommends that you create a NetApp Support Site account.

8. Create a NetApp Support Site account by completing the [NetApp Support Site User Registration form](#)
 - a. Be sure to select the appropriate User Level, which is typically **NetApp Customer/End User**.
 - b. Be sure to copy the account serial number (960xxxx) used above for the serial number field. This will speed up the account processing.

After you finish

NetApp should reach out to you during this process. This is a one-time onboarding exercise for new users.

Once you have your NetApp Support Site account, associate the account with your BlueXP login by completing the steps under [Existing customer with an NSS account](#).

Associate NSS credentials for Cloud Volumes ONTAP support

Associating NetApp Support Site credentials with your BlueXP account is required to enable the following key workflows for Cloud Volumes ONTAP:

- Registering pay-as-you-go Cloud Volumes ONTAP systems for support

Providing your NSS account is required to activate support for your system and to gain access to NetApp technical support resources.

- Deploying Cloud Volumes ONTAP when you bring your own license (BYOL)

Providing your NSS account is required so that BlueXP can upload your license key and to enable the subscription for the term that you purchased. This includes automatic updates for term renewals.

- Upgrading Cloud Volumes ONTAP software to the latest release


Associating NSS credentials with your BlueXP account is different than the NSS account that is associated with a BlueXP user login.

These NSS credentials are associated with your specific BlueXP account ID. Users who belong to the BlueXP account can access these credentials from **Support > NSS Management**.

- If you have a customer-level account, you can add one or more NSS accounts.
- If you have a partner or reseller account, you can add one or more NSS accounts, but they can't be added alongside customer-level accounts.

Steps

1. In the upper right of the BlueXP console, select the Help icon, and select **Support**.



2. Select **NSS Management > Add NSS Account**.
3. When you're prompted, select **Continue** to be redirected to a Microsoft login page.

NetApp uses Microsoft Entra ID as the identity provider for authentication services specific to support and licensing.

4. At the login page, provide your NetApp Support Site registered email address and password to perform the authentication process.

These actions enable BlueXP to use your NSS account for things like license downloads, software upgrade verification, and future support registrations.


Note the following:


- The NSS account must be a customer-level account (not a guest or temp account). You can have multiple customer-level NSS accounts.
- There can be only one NSS account if that account is a partner-level account. If you try to add customer-level NSS accounts and a partner-level account exists, you'll get the following error message:

"The NSS customer type is not allowed for this account as there are already NSS Users of different type."

The same is true if you have pre-existing customer-level NSS accounts and try to add a partner-level account.

- Upon successful login, NetApp will store the NSS user name.

This is a system-generated ID that maps to your email. On the **NSS Management** page, you can display your email from the  menu.

- If you ever need to refresh your login credential tokens, there is also an **Update Credentials** option in the  menu.

Using this option prompts you to log in again. Note that the token for these accounts expire after 90 days. A notification will be posted to alert you of this.

Get help

NetApp provides support for BlueXP and its cloud services in a variety of ways. Extensive free self-support options are available 24x7, such as knowledgebase (KB) articles and a community forum. Your support registration includes remote technical support via web ticketing.

Get support for a cloud provider file service

For technical support related to a cloud provider file service, its infrastructure, or any solution using the service, refer to "Getting help" in the BlueXP documentation for that product.

- [Amazon FSx for ONTAP](#)
- [Azure NetApp Files](#)
- [Cloud Volumes Service for Google Cloud](#)

To receive technical support specific to BlueXP and its storage solutions and services, use the support options described below.

Use self-support options

These options are available for free, 24 hours a day, 7 days a week:

- Documentation

The BlueXP documentation that you're currently viewing.

- [Knowledge base](#)

Search through the BlueXP knowledge base to find helpful articles to troubleshoot issues.

- [Communities](#)

Join the BlueXP community to follow ongoing discussions or create new ones.

Create a case with NetApp support

In addition to the self-support options above, you can work with a NetApp Support specialist to resolve any issues after you activate support.

Before you get started

- To use the **Create a Case** capability, you must first associate your NetApp Support Site credentials with your BlueXP login. [Learn how to manage credentials associated with your BlueXP login.](#)
- If you're opening a case for an ONTAP system that has a serial number, then your NSS account must be associated with the serial number for that system.

Steps

1. In BlueXP, select **Help > Support**.
2. On the **Resources** page, choose one of the available options under Technical Support:
 - a. Select **Call Us** if you'd like to speak with someone on the phone. You'll be directed to a page on netapp.com that lists the phone numbers that you can call.
 - b. Select **Create a Case** to open a ticket with a NetApp Support specialist:
 - **Service:** Select the service that the issue is associated with. For example, BlueXP when specific to a technical support issue with workflows or functionality within the service.
 - **Working Environment:** If applicable to storage, select **Cloud Volumes ONTAP** or **On-Prem** and then the associated working environment.


The list of working environments are within scope of the BlueXP account, workspace, and Connector you have selected in the top banner of the service.

- **Case Priority:** Choose the priority for the case, which can be Low, Medium, High, or Critical.

To learn more details about these priorities, hover your mouse over the information icon next to the field name.

- **Issue Description:** Provide a detailed description of your problem, including any applicable error messages or troubleshooting steps that you performed.
- **Additional Email Addresses:** Enter additional email addresses if you'd like to make someone else aware of this issue.
- **Attachment (Optional):** Upload up to five attachments, one at a time.

Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.

ntapitdemo 


NetApp Support Site Account

Service

Select ▼

Working Enviroment


Select ▼

Case Priority 

Low - General guidance ▼

Issue Description



Provide detailed description of problem, applicable error messages and troubleshooting steps taken.

Additional Email Addresses (Optional) 

Type here

Attachment (Optional)

No files selected

 Upload 

After you finish

A pop-up will appear with your support case number. A NetApp Support specialist will review your case and get back to you soon.

For a history of your support cases, you can select **Settings > Timeline** and look for actions named "create support case." A button to the far right lets you expand the action to see details.

It's possible that you might encounter the following error message when trying to create a case:

"You are not authorized to Create a Case against the selected service"

This error could mean that the NSS account and the company of record it's associated with is not the same company of record for the BlueXP account serial number (ie. 960xxxx) or the working environment serial number. You can seek assistance using one of the following options:

- Use the in-product chat
- Submit a non-technical case at <https://mysupport.netapp.com/site/help>

Manage your support cases (Preview)

You can view and manage active and resolved support cases directly from BlueXP. You can manage the cases associated with your NSS account and with your company.

Case management is available as a Preview. We plan to refine this experience and add enhancements in upcoming releases. Please send us feedback by using the in-product chat.

Note the following:

- The case management dashboard at the top of the page offers two views:
 - The view on the left shows the total cases opened in the past 3 months by the user NSS account you provided.
 - The view on the right shows the total cases opened in the past 3 months at your company level based on your user NSS account.

The results in the table reflect the cases related to the view that you selected.

- You can add or remove columns of interest and you can filter the contents of columns like Priority and Status. Other columns provide just sorting capabilities.

View the steps below for more details.

- At a per-case level, we offer the ability to update case notes or close a case that is not already in Closed or Pending Closed status.

Steps

1. In BlueXP, select **Help > Support**.
2. Select **Case Management** and if you're prompted, add your NSS account to BlueXP.

The **Case management** page shows open cases related to the NSS account that is associated with your BlueXP user account. This is the same NSS account that appears at the top of the **NSS management** page.

3. Optionally modify the information that displays in the table:
 - Under **Organization's cases**, select **View** to view all cases associated with your company.
 - Modify the date range by choosing an exact date range or by choosing a different time frame.

Search icon | Cases opened on the last 3 months | Create a case

Date created	Last updated		Status (5)	
December 22, 2022	December 29, 2022	Last 7 days	Assigned	...
December 21, 2022	December 28, 2022	Last 30 days	Active	...
December 15, 2022	December 27, 2022	Last 3 months	Pending customer	...
December 14, 2022	December 26, 2022	Medium (P3)	Solution proposed	...
		Low (P4)		

Apply | Reset

- Filter the contents of the columns.

Search icon | Cases opened on the last 3 months | Create a case

Last updated	Priority	Status (5)	
December 29, 2022	Critical (P1)	Active	...
December 28, 2022	High (P2)	Pending customer	...
December 27, 2022	Medium (P3)	Solution proposed	...
December 26, 2022	Low (P4)	Pending closed	...
		Closed	...

Apply | Reset

- Change the columns that appear in the table by selecting  and then choosing the columns that you'd like to display.

Search icon | Cases opened on the last 3 months | Create a case

Last updated	Priority	Status (5)	
December 29, 2022	Critical (P1)	Last updated	...
December 28, 2022	High (P2)	Priority	...
December 27, 2022	Medium (P3)	Cluster name	...
December 26, 2022	Low (P4)	Case owner	...
		Opened by	...

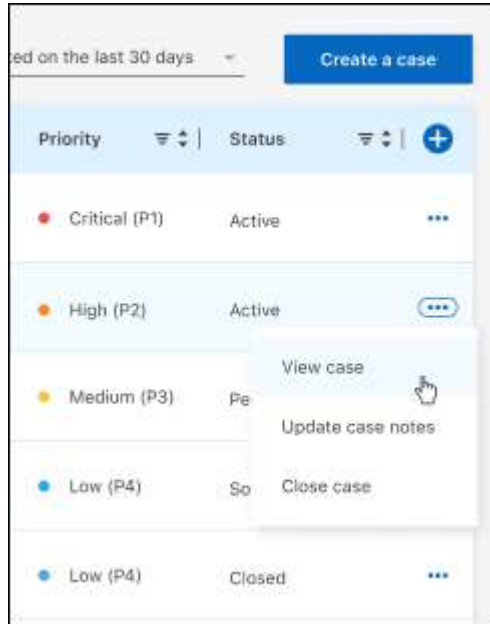
Apply | Reset

4. Manage an existing case by selecting ... and selecting one of the available options:

- **View case:** View full details about a specific case.
- **Update case notes:** Provide additional details about your problem or select **Upload files** to attach up to a maximum of five files.

Attachments are limited to 25 MB per file. The following file extensions are supported: txt, log, pdf, jpg/jpeg, rtf, doc/docx, xls/xlsx, and csv.

- **Close case:** Provide details about why you're closing the case and select **Close case**.



Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<https://www.netapp.com/company/legal/copyright/>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

- [Notice for BlueXP](#)
- [Notice for the BlueXP backup and recovery](#)
- [Notice for Single File Restore](#)

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.