



Documentation

## **BlueXP Private Mode (legacy interface)**

Documentation for BlueXP Private mode and supported services

### **Abstract**

Documentation for BlueXP private mode (legacy interface) and supported services: BlueXP backup and recovery, BlueXP classification, Cloud Volumes ONTAP in BlueXP, On-premises ONTAP cluster management using BlueXP, and BlueXP replication.



# **BlueXP setup and administration documentation**

## **BlueXP setup and administration**

NetApp  
August 29, 2025

# Table of Contents

- BlueXP setup and administration documentation . . . . . 1
- Release notes . . . . . 2
  - What’s new . . . . . 2
    - 11 August 2025 . . . . . 2
    - 31 July 2025 . . . . . 2
    - 21 July 2025 . . . . . 3
    - 14 July 2025 . . . . . 3
    - 9 June 2025 . . . . . 4
    - 29 May 2025 . . . . . 5
    - 12 May 2025 . . . . . 6
    - 14 April 2025 . . . . . 7
    - 28 March 2025 . . . . . 7
    - 10 March 2025 . . . . . 8
    - 6 March 2025 . . . . . 8
    - 18 February 2025 . . . . . 9
    - 10 February 2025 . . . . . 9
    - 13 January 2025 . . . . . 12
    - 16 December 2024 . . . . . 12
    - 9 December 2024 . . . . . 13
    - 26 November 2024 . . . . . 13
    - 11 November 2024 . . . . . 14
    - 10 October 2024 . . . . . 14
    - 7 October 2024 . . . . . 14
    - 30 September 2024 . . . . . 17
    - 9 September 2024 . . . . . 18
    - 22 August 2024 . . . . . 19
    - 8 August 2024 . . . . . 19
    - 31 July 2024 . . . . . 20
    - 15 July 2024 . . . . . 21
    - 8 July 2024 . . . . . 21
    - 12 June 2024 . . . . . 22
    - 4 June 2024 . . . . . 22
    - 17 May 2024 . . . . . 22
- Known limitations . . . . . 23
  - Connector limitations . . . . . 23
- Changes to supported Linux operating systems . . . . . 24
  - Supported operating systems . . . . . 24
  - Support for RHEL 8 and 9 . . . . . 25
  - End of support for RHEL 7 and CentOS 7 . . . . . 26
  - Related information . . . . . 26
- Get started . . . . . 28
  - Learn the basics . . . . . 28
  - Learn about BlueXP . . . . . 28

Learn about BlueXP Connectors . . . . .	31
Learn about BlueXP deployment modes . . . . .	35
Get started with standard mode . . . . .	45
Getting started workflow (standard mode) . . . . .	45
Prepare networking for the BlueXP console . . . . .	46
Sign up or log in to BlueXP . . . . .	48
Create a Connector . . . . .	50
Subscribe to NetApp Intelligent Services (standard mode) . . . . .	166
What you can do next (standard mode) . . . . .	171
Get started with restricted mode . . . . .	171
Getting started workflow (restricted mode) . . . . .	171
Prepare for deployment in restricted mode . . . . .	172
Deploy the Connector in restricted mode . . . . .	191
Subscribe to NetApp Intelligent Services (restricted mode) . . . . .	203
What you can do next (restricted mode) . . . . .	208
Get started with private mode . . . . .	209
Getting started workflow (private mode) . . . . .	209
Prepare for deployment in private mode . . . . .	209
Deploy the Connector in private mode . . . . .	225
What you can do next (private mode) . . . . .	230
Use BlueXP . . . . .	231
Log in to BlueXP . . . . .	231
Manage your BlueXP user settings . . . . .	233
Change your display name . . . . .	233
Configure multi-factor authentication . . . . .	233
Regenerate your MFA recovery code . . . . .	234
Delete your MFA configuration . . . . .	234
Contact your Organization administrator . . . . .	235
Configure dark mode (dark theme) . . . . .	235
Administer BlueXP . . . . .	236
Identity and access management . . . . .	236
Learn about BlueXP identity and access management . . . . .	236
Get started with BlueXP identity and access management . . . . .	243
Organize your resources in BlueXP IAM with folders and projects . . . . .	244
Add BlueXP members and service accounts . . . . .	249
Use roles to manage user access to resources . . . . .	254
Manage the resource hierarchy in your BlueXP organization . . . . .	255
Associate a BlueXP Connector with other folders and projects . . . . .	258
Switch between BlueXP organizations, projects, and Connectors . . . . .	259
Organization and project IDs . . . . .	261
Monitor or audit IAM activity from the BlueXP timeline . . . . .	262
BlueXP access roles . . . . .	263
Identity federation . . . . .	275
Enable single sign-on by using identity federation with BlueXP . . . . .	275
Domain verification . . . . .	277



Configure federations . . . . .	278
Manage federations in BLueXP . . . . .	285
Import your federation to BlueXP . . . . .	287
Connectors . . . . .	287
Maintain the Connector VM and operating system . . . . .	287
Install a CA-signed certificate for web-based console access . . . . .	290
Configure a Connector to use a proxy server . . . . .	292
Require the use of IMDSv2 on Amazon EC2 instances . . . . .	299
Manage connector upgrades . . . . .	301
Work with multiple Connectors . . . . .	303
Troubleshoot the Connector . . . . .	304
Uninstall and remove the Connector . . . . .	307
Default configuration for the Connector . . . . .	309
Enforce ONTAP permissions for ONTAP Advanced View (ONTAP System Manager) . . . . .	311
Credentials and subscriptions . . . . .	312
AWS . . . . .	312
Azure . . . . .	326
Google Cloud . . . . .	340
Manage NSS credentials associated with BlueXP . . . . .	346
Manage credentials associated with your BlueXP login . . . . .	351
Monitor BlueXP operations . . . . .	353
Audit user activity from the BlueXP timeline . . . . .	353
Monitor activities using the Notification Center . . . . .	354
Reference . . . . .	359
Connector maintenance console . . . . .	359
Connector maintenance console . . . . .	359
Permissions . . . . .	360
Permissions summary for BlueXP . . . . .	360
AWS permissions for the Connector . . . . .	364
Azure permissions for the Connector . . . . .	395
Google Cloud permissions for the Connector . . . . .	414
Ports . . . . .	420
Connector security group rules in AWS . . . . .	420
Connector security group rules in Azure . . . . .	421
Connector firewall rules in Google Cloud . . . . .	423
Ports for the on-premisesConnector . . . . .	424
Knowledge and support . . . . .	425
Register for support . . . . .	425
Support registration overview . . . . .	425
Register BlueXP for NetApp support . . . . .	425
Associate NSS credentials for Cloud Volumes ONTAP support . . . . .	427
Get help . . . . .	429
Get support for a cloud provider file service . . . . .	429
Use self-support options . . . . .	429
Create a case with NetApp support . . . . .	429

Manage your support cases (Preview) .....	432
Legal notices .....	435
Copyright .....	435
Trademarks .....	435
Patents .....	435
Privacy policy .....	435
Open source .....	435

# BlueXP setup and administration documentation

# Release notes

## What's new

Learn what's new with BlueXP administration features: identity and access management (IAM), Connectors, cloud provider credentials, and more.

### 11 August 2025

#### Connector 3.9.55

This release of the BlueXP Connector includes security improvements, and bug fixes.

The 3.9.55 release is available for standard mode and restricted mode.

#### Japanese language support

The BlueXP UI is now available in the Japanese language. If your browser language is Japanese, BlueXP displays in Japanese. To access documentation in Japanese, use the language menu on the documentation website.

#### Operational resiliency feature

The Operational resiliency feature has been removed from BlueXP. Contact NetApp support if you encounter issues.

#### BlueXP Identity and Access Management (IAM)

Identity and Access Management in BlueXP now provides the following feature.

#### New access role for operational support

BlueXP now supports an Operational support analyst role. This role grants a user permissions to monitor storage alerts, view the BlueXP audit timeline, and enter and track NetApp Support cases.

[Learn more about using access roles.](#)

### 31 July 2025

#### Private mode release (3.9.54)

A new private mode release is now available to download from the [NetApp Support Site](#)

The 3.9.54 release includes updates to the following BlueXP components and services.

Component or service	Version included in this release	Changes since the previous private mode release
Connector	3.9.54, 3.9.53	Go to the <a href="#">what's new in BlueXP page</a> and refer to the changes included for versions 3.9.54 and 3.9.53.

Component or service	Version included in this release	Changes since the previous private mode release
Backup and recovery	28 July 2025	Go to the <a href="#">what's new in BlueXP backup and recovery page</a> and refer to the changes included in the July 2025 release.
Classification	14 July 2025 (version 1.45)	Go to the <a href="#">what's new in BlueXP classification page</a> .

For more details about private mode, including how to upgrade, refer to the following:

- [Learn about private mode](#)
- [Learn how to get started with BlueXP in private mode](#)
- [Learn how to upgrade the Connector when using private mode](#)

## 21 July 2025

### Support for Google Cloud NetApp Volumes

You can now view Google Cloud NetApp Volumes in BlueXP. [Learn more about Google Cloud NetApp Volumes.](#)

### BlueXP Identity and Access Management (IAM)

#### New access role for Google Cloud NetApp Volumes

BlueXP now supports using an access role for the following storage system:

- Google Cloud NetApp Volumes

[Learn more about using access roles.](#)

## 14 July 2025

### Connector 3.9.54

This release of the BlueXP Connector includes security improvements, bug fixes, and the following new features:

- Support for transparent proxies for Connectors dedicated to supporting Cloud Volumes ONTAP services. [Learn more about configuring a transparent proxy.](#)
- Ability to use network tags to help route Connector traffic when the Connector is deployed in a Google Cloud environment.
- Additional in-product notifications for Connector health monitoring, including CPU and RAM usage.

At this time, the 3.9.54 release is available for standard mode and restricted mode.

### BlueXP Identity and Access Management (IAM)

Identity and Access Management in BlueXP now provides the following features:

- Support for IAM in private mode, allowing you to manage user access and permissions for BlueXP services and applications.

- Streamlined management of identity federations, including easier navigation, clearer options for configuring federated connections, and improved visibility into existing federations.
- Access roles for BlueXP backup and recovery, BlueXP disaster recovery, and federation management.

### **Support for IAM in private mode**

BlueXP now supports IAM in private mode, allowing you to manage user access and permissions for BlueXP services and applications. This enhancement enables private mode customers to leverage role-based access control (RBAC) for better security and compliance.

[Learn more about IAM in BlueXP.](#)

### **Streamlined management of identity federations**

BlueXP now offers a more intuitive interface for managing identity federation. This includes easier navigation, clearer options for configuring federated connections, and improved visibility into existing federations.

Enabling single sign-on (SSO) through identity federation lets users log in to BlueXP with their corporate credentials. This improves security, reduces password use, and simplifies onboarding.

You'll be prompted to import any existing federated connections to the new interface to gain access to the new management features. This allows you to take advantage of the latest enhancements without having to recreate your federated connections. [Learn more about importing your existing federated connection to BlueXP.](#)

Improved federation management allows you to:

- Add more than one verified domain to a federated connection, allowing you to use multiple domains with the same identity provider (IdP).
- Disable or delete federated connections when needed, giving you control over user access and security.
- Control access to federation management with IAM roles.

[Learn more about identity federation in BlueXP.](#)

### **New access roles for BlueXP backup and recovery, BlueXP disaster recovery, and federation management**

BlueXP now supports using IAM roles for the following features and data services:

- BlueXP backup and recovery
- BlueXP disaster recovery
- Federation

[Learn more about using access roles.](#)

## **9 June 2025**

### **Connector 3.9.53**

This release of the BlueXP Connector includes security improvements and bug fixes.

The 3.9.53 release is available for standard mode and restricted mode.

## Disk space usage alerts

The Notifications Center now includes alerts for disk space usage on the Connector. [Learn more.](#)

## Audit improvements

The Timeline now includes login and logout events for users. You can see when login activity, which can help with auditing and security monitoring. API users who have the Organization administrator role can view the email address of the user who logged in by including the `includeUserData=true`` parameter as in the following: `/audit/<account_id>?includeUserData=true.`

## Keystone subscription management available in BlueXP

You can manage your NetApp Keystone subscription from BlueXP.

[Learn about Keystone subscription management in BlueXP.](#)

## BlueXP Identity and Access Management (IAM)

### Multi-factor authentication (MFA)

Unfederated users can enable MFA for their BlueXP accounts to improve security. Administrators can manage MFA settings, including resetting or disabling MFA for users as needed. This is supported in standard mode only.

[Learn about setting up multi-factor authentication for yourself.](#)  
[Learn about administering multi-factor authentication for users.](#)

## Workloads

You can now view and delete Amazon FSx for NetApp ONTAP credentials from the Credentials page in BlueXP.

## 29 May 2025

### Private mode release (3.9.52)

A new private mode release is now available to download from the [NetApp Support Site](#)

The 3.9.52 release includes updates to the following BlueXP components and services.

Component or service	Version included in this release	Changes since the previous private mode release
Connector	3.9.52, 3.9.51	Go to the <a href="#">what's new in BlueXP connector page</a> and refer to the changes included for versions 3.9.52 and 3.9.50.
Backup and recovery	12 May 2025	Go to the <a href="#">what's new in BlueXP backup and recovery page</a> and refer to the changes included in the May 2025 release.
Classification	12 May 2025 (version 1.43)	Go to the <a href="#">what's new in BlueXP classification page</a> and refer to the changes included in the 1.38 to 1.371.41 releases.

For more details about private mode, including how to upgrade, refer to the following:

- [Learn about private mode](#)
- [Learn how to get started with BlueXP in private mode](#)
- [Learn how to upgrade the Connector when using private mode](#)

## 12 May 2025

### Connector 3.9.52

This release of the BlueXP Connector includes minor security improvements and bug fixes, as well as some additional updates.

At this time, the 3.9.52 release is available for standard mode and restricted mode.

#### Support for Docker 27 and Docker 28

Docker 27 and Docker 28 are now supported with the Connector.

#### Cloud Volumes ONTAP

Cloud Volumes ONTAP nodes no longer shutdown when the Connector is out of compliance or down for more than 14 days. Cloud Volumes ONTAP still sends Event Management messages when it loses access to the Connector. This change is to ensure that Cloud Volumes ONTAP can continue to operate even if the Connector is down for an extended period of time. It does not change compliance requirements for the Connector.

#### Keystone administration available in BlueXP

The beta for NetApp Keystone in BlueXP has added access to Keystone administration. You can access the sign-up page for the NetApp Keystone beta from the left navigation bar of BlueXP.

#### BlueXP Identity and Access Management (IAM)

##### New storage management roles

The Storage admin, System health specialist, and Storage viewer roles are available and can be assigned to users.

These roles enable you to manage who in your organization can discover and manage storage resources, as well as view storage health information and perform software updates.

These roles are supported for controlling access to the following storage resources:

- E-Series systems
- StorageGRID systems
- On-premises ONTAP systems

You can also use these roles to control access to the following BlueXP services:

- Software updates
- Digital advisor
- Operational resiliency
- Economic efficiency



- Sustainability

The following roles have been added:

- **Storage admin**

Administer storage health, governance, and discovery for the storage resources in the organization. This role can also perform software updates on storage resources.

- **System health specialist**

Administer storage health and governance for the storage resources in the organization. This role can also perform software updates on storage resources. This role cannot modify or delete working environments.

- **Storage viewer**

View storage health information and governance data.

[Learn about access roles.](#)

## 14 April 2025

### Connector 3.9.51

This release of the BlueXP Connector includes minor security improvements and bug fixes.

At this time, the 3.9.51 release is available for standard mode and restricted mode.

#### **Secure endpoints for Connector downloads now supported for Backup and recovery and Ransomware protection**

If you are using Backup and recovery or Ransomware protection, you can now use secure endpoints for Connector downloads. [Learn about secure endpoints for Connector downloads.](#)

### **BlueXP Identity and Access Management (IAM)**

- Users without the Org admin or Folder or project admin must be assigned a Ransomware protection role to have access to Ransomware protection. You can assign a user one of two roles: Ransomware protection admin or Ransomware protection viewer.
- Users without the Org admin or Folder or project admin must be assigned a Keystone role to have access to Keystone. You can assign a user one of two roles: Keystone admin or Keystone viewer.

[Learn about access roles.](#)

- If you have the Org admin or Folder or project admin role, you can now associate a Keystone subscription with an IAM project. Associating a Keystone subscription with an IAM project allows you to control access to Keystone within BlueXP.

## 28 March 2025

### **Private mode release (3.9.50)**

A new private mode release is now available to download from the [NetApp Support Site](#)

The 3.9.50 release includes updates to the following BlueXP components and services.

Component or service	Version included in this release	Changes since the previous private mode release
Connector	3.9.50, 3.9.49	Go to the <a href="#">what's new in BlueXP connector page</a> and refer to the changes included for versions 3.9.50 and 3.9.49.
Backup and recovery	17 March 2025	Go to the <a href="#">what's new in BlueXP backup and recovery page</a> and refer to the changes included in the March 2024 release.
Classification	10 March 2025 (version 1.41)	Go to the <a href="#">what's new in BlueXP classification page</a> and refer to the changes included in the 1.38 to 1.371.41 releases.

For more details about private mode, including how to upgrade, refer to the following:

- [Learn about private mode](#)
- [Learn how to get started with BlueXP in private mode](#)
- [Learn how to upgrade the Connector when using private mode](#)

## 10 March 2025

### Connector 3.9.50

This release of the BlueXP Connector includes minor security improvements and bug fixes.

- Management of Cloud Volumes ONTAP systems is now supported by Connectors that have SELinux enabled on the operating system.

[Learn more about SELinux](#)

At this time, the 3.9.50 release is available for standard mode and restricted mode.

### NetApp Keystone beta available in BlueXP

NetApp Keystone will soon be available from BlueXP and is now in beta. You can access the sign-up page for the NetApp Keystone beta from the left navigation bar of BlueXP.

## 6 March 2025

### Connector 3.9.49 update

#### ONTAP System Manager access when BlueXP uses a Connector

A BlueXP administrator (users with the Organization admin role) can configure BlueXP to prompt users to enter their ONTAP credentials in order to access ONTAP system manager. When this setting is enabled, users need enter their ONTAP credentials each time as they are not stored in BlueXP.

This feature is available in Connector version 3.9.49 and higher. [Learn how to configure credentials settings..](#)

### Connector 3.9.48 update

#### Ability to disable the auto-upgrade setting for the Connector

You can disable the auto-upgrade feature of the Connector.

When you use BlueXP in standard mode or restricted mode, BlueXP automatically upgrades your Connector to the latest release, as long as the Connector has outbound internet access to obtain the software update. If you need to manually manage when the connector is upgraded, you can now disable automatic upgrades for standard mode or restricted mode.



This change does not impact BlueXP private mode where you must always upgrade the connector yourself.

This feature is available in Connector version 3.9.48 and higher.

[Learn how to disable auto-upgrade for the Connector.](#)

## 18 February 2025

### Private mode release (3.9.48)

A new private mode release is now available to download from the [NetApp Support Site](#)

The 3.9.48 release includes updates to the following BlueXP components and services.

Component or service	Version included in this release	Changes since the previous private mode release
Connector	3.9.48	Go to the <a href="#">what's new in BlueXP connector page</a> and refer to the changes included for versions 3.9.48.
Backup and recovery	21 February 2025	Go to the <a href="#">what's new in BlueXP backup and recovery page</a> and refer to the changes included in the February 2025 release.
Classification	22 January 2025 (version 1.39)	Go to the <a href="#">what's new in BlueXP classification page</a> and refer to the changes included in the 1.39 release.

## 10 February 2025

### Connector 3.9.49

This release of the BlueXP Connector includes minor security improvements and bug fixes.

At this time, the 3.9.49 release is available for standard mode and restricted mode.

### BlueXP identity and access management (IAM)

- Support for assigning multiple roles to a BlueXP user.
- Support for assigning a role on multiple resources of the BlueXP organization (Org/folder/project)
- Roles are now associated with one of two categories: platform and data service.

### Restricted mode now uses BlueXP IAM

BlueXP identity and access management (IAM) is now used in restricted mode.

BlueXP identity and access management (IAM) is a resource and access management model that replaces and enhances the previous functionality provided by BlueXP accounts when using BlueXP in standard and restricted mode.

## Related information

- [Learn about BlueXP IAM](#)
- [Get started with BlueXP IAM](#)

BlueXP IAM provides more granular management of resources and permissions:

- A top-level *organization* enables you to manage access across your various *projects*.
- *Folders* enable you to group related projects together.
- Enhanced resource management enables you to associate a resource with one or more folders or projects.

For example, you can associate a Cloud Volumes ONTAP system with multiple projects.

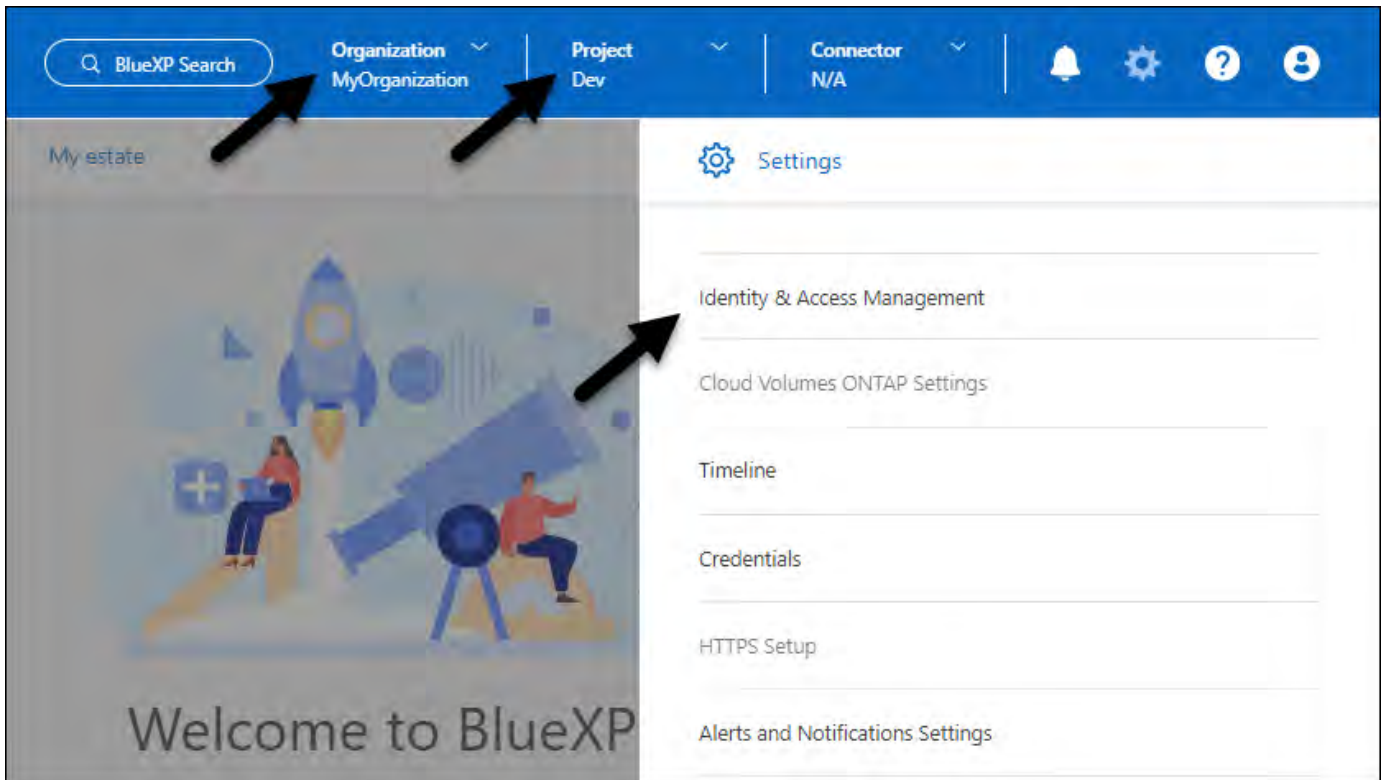
- Enhanced access management enables you to assign a role to members at different levels of the organization hierarchy.

These enhancements provide better control over the actions that users can perform and the resources that they can access.

## How BlueXP IAM affects your existing account in restricted mode

When you log in to BlueXP, you'll notice these changes:

- Your *account* is now called an *organization*
- Your *workspaces* are now called *projects*
- The names of user roles have changed:
  - *Account admin* is now *Organization admin*
  - *Workspace admin* is now *Folder or project admin*
  - *Compliance viewer* is now *Classification viewer*
- Under Settings, you can access BlueXP identity and access management to take advantage of these enhancements



Note the following:

- There are no changes to your existing users or working environments.
- While the names of the roles have changed, there are no differences from a permissions perspective. Users will continue to have access to the same working environments as before.
- There are no changes to how you log in to BlueXP. BlueXP IAM works with NetApp cloud logins, NetApp Support Site credentials, and federated connections just like BlueXP accounts did.
- If you had multiple BlueXP accounts, you now have multiple BlueXP organizations.

### API for BlueXP IAM

This change introduces a new API for BlueXP IAM, but it is backwards compatible with the previous tenancy API. [Learn about the API for BlueXP IAM](#)

### Supported deployment modes

BlueXP IAM is supported when using BlueXP in standard and restricted mode. If you're using BlueXP in private mode, then you'll continue using a BlueXP *account* to manage workspaces, users, and resources.

### Private mode release (3.9.48)

A new private mode release is now available to download from the [NetApp Support Site](#)

The 3.9.48 release includes updates to the following BlueXP components and services.

Component or service	Version included in this release	Changes since the previous private mode release
Connector	3.9.48	Go to the <a href="#">what's new in BlueXP connector page</a> and refer to the changes included for versions 3.9.48.

Component or service	Version included in this release	Changes since the previous private mode release
Backup and recovery	21 February 2025	Go to the <a href="#">what's new in BlueXP backup and recovery page</a> and refer to the changes included in the February 2025 release.
Classification	22 January 2025 (version 1.39)	Go to the <a href="#">what's new in BlueXP classification page</a> and refer to the changes included in the 1.39 release.

## 13 January 2025

### Connector 3.9.48

This release of the BlueXP Connector includes minor security improvements and bug fixes.

At this time, the 3.9.48 release is available for standard mode and restricted mode.

### BlueXP identity and access management

- The Resources page now displays undiscovered resources. Undiscovered resources are storage resources that BlueXP knows about but you have not created working environments for. For example, resources that display in digital advisor that do not yet have working environments display on the Resources page as undiscovered resources.
- Amazon FSx for NetApp ONTAP resources aren't displayed on the IAM resources page as you cannot associate them with an IAM role. You can view these resources on their respective canvas or from workloads.

### Create a support case for additional BlueXP services

After you register BlueXP for support, you can create a support case directly from the BlueXP web-based console. When you create the case, you need to select the service that the issue is associated with.

Starting with this release, you can now create a support case and associate it with additional BlueXP services:

- BlueXP disaster recovery
- BlueXP ransomware protection

[Learn more about creating a support case.](#)

## 16 December 2024

### New secure endpoints to obtain Connector images

When you install the Connector, or when an automatic upgrade occurs, the Connector contacts repositories to download images for the installation or upgrade. By default, the Connector has always contacted the following endpoints:

- [https://\\*.blob.core.windows.net](https://*.blob.core.windows.net)
- <https://cloudmanagerinfraproduct.azurecr.io>

The first endpoint includes a wild card because we can't provide a definitive location. The load balancing of the repository is managed by the service provider, which means the downloads can happen from different endpoints.

For increased security, the Connector can now download installation and upgrades images from dedicated endpoints:

- <https://bluexpinfraprod.eastus2.data.azurecr.io>
- <https://bluexpinfraprod.azurecr.io>

We recommend that you start using these new endpoints by removing the existing endpoints from your firewall rules and allowing the new endpoints.

These new endpoints are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.

Note the following:

- The existing endpoints are still supported. If you don't want to use the new endpoints, no changes are required.
- The Connector contacts the existing endpoints first. If those endpoints aren't accessible, the Connector automatically contacts the new endpoints.
- The new endpoints are not supported in the following scenarios:
  - If the Connector is installed in a Government region.
  - If you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection.

For both of these scenarios, you can continue to use the existing endpoints.

## 9 December 2024

### Connector 3.9.47

This release of the BlueXP Connector includes bug fixes and a change to the endpoints contacted during Connector installation.

At this time, the 3.9.47 release is available for standard mode and restricted mode.

#### Endpoint to contact NetApp support during installation

When you manually install the Connector, the installer no longer contacts <https://support.netapp.com>.

The installer still contacts <https://mysupport.netapp.com>.

#### BlueXP identity and access management

The Connectors page lists only currently available Connectors. It no longer displays Connectors that you have removed.

## 26 November 2024

### Private mode release (3.9.46)

A new private mode release is now available to download from the [NetApp Support Site](#)

The 3.9.46 release includes updates to the following BlueXP components and services.

Component or service	Version included in this release	Changes since the previous private mode release
Connector	3.9.46	Minor security improvements and bug fixes
Backup and recovery	22 November 2024	Go to the <a href="#">what's new in BlueXP backup and recovery page</a> and refer to the changes included in the November 2024 release
Classification	4 November 2024 (version 1.37)	Go to the <a href="#">what's new in BlueXP classification page</a> and refer to the changes included in the 1.32 to 1.37 releases
Cloud Volumes ONTAP management	11 November 2024	Go to the <a href="#">what's new with Cloud Volumes ONTAP management page</a> and refer to the changes included in the October 2024 and November 2024 releases
On-premises ONTAP cluster management	26 November 2024	Go to the <a href="#">what's new with on-premises ONTAP cluster management page</a> and refer to the changes included in the November 2024 release

While the BlueXP digital wallet and BlueXP replication are also included with private mode, there are no changes from the previous private mode release.

For more details about private mode, including how to upgrade, refer to the following:

- [Learn about private mode](#)
- [Learn how to get started with BlueXP in private mode](#)
- [Learn how to upgrade the Connector when using private mode](#)

## 11 November 2024

### Connector 3.9.46

This release of the BlueXP Connector includes minor security improvements and bug fixes.

At this time, the 3.9.46 release is available for standard mode and restricted mode.

### ID for IAM projects

You can now view the ID for a project from BlueXP identity and access management. You might need to use the ID when making an API call.

[Learn how to obtain the ID for a project.](#)

## 10 October 2024

### Connector 3.9.45 patch

This patch includes bug fixes.

## 7 October 2024

### BlueXP identity and access management

BlueXP identity and access management (IAM) is a new resource and access management model that



replaces and enhances the previous functionality provided by BlueXP accounts when using BlueXP in standard mode.

BlueXP IAM provides more granular management of resources and permissions:

- A top-level *organization* enables you to manage access across your various *projects*.
- *Folders* enable you to group related projects together.
- Enhanced resource management enables you to associate a resource with one or more folders or projects.

For example, you can associate a Cloud Volumes ONTAP system with multiple projects.


- Enhanced access management enables you to assign a role to members at different levels of the organization hierarchy.

These enhancements provide better control over the actions that users can perform and the resources that they can access.

### **How BlueXP IAM affects your existing account**

When you log in to BlueXP, you'll notice these changes:

- Your *account* is now called an *organization*
- Your *workspaces* are now called *projects*
- The names of user roles have changed:
  - *Account admin* is now *Organization admin*
  - *Workspace admin* is now *Folder or project admin*
  - *Compliance viewer* is now *Classification viewer*
- Under Settings, you can access BlueXP identity and access management to take advantage of these enhancements



Note the following:

- There are no changes to your existing users or working environments.
- While the names of the roles have changed, there are no differences from a permissions perspective. Users will continue to have access to the same working environments as before.
- There are no changes to how you log in to BlueXP. BlueXP IAM works with NetApp cloud logins, NetApp Support Site credentials, and federated connections just like BlueXP accounts did.
- If you had multiple BlueXP accounts, you now have multiple BlueXP organizations.

### API for BlueXP IAM

This change introduces a new API for BlueXP IAM, but it is backwards compatible with the previous tenancy API. [Learn about the API for BlueXP IAM](#)

### Supported deployment modes

BlueXP IAM is supported when using BlueXP in standard mode. If you're using BlueXP in restricted mode or private mode, then you'll continue using a BlueXP *account* to manage workspaces, users, and resources.

### Where to go next

- [Learn about BlueXP IAM](#)
- [Get started with BlueXP IAM](#)

### Connector 3.9.45

This release includes expanded operating system support and bug fixes.

The 3.9.45 release is available for standard mode and restricted mode.

### Support for Ubuntu 24.04 LTS

Starting with the 3.9.45 release, BlueXP now supports new installations of the Connector on Ubuntu 24.04 LTS hosts when using BlueXP in standard mode or restricted mode.

[View Connector host requirements.](#)

## Support for SELinux with RHEL hosts

BlueXP now supports the Connector with Red Hat Enterprise Linux hosts that have SELinux enabled in either enforcing mode or permissive mode.

Support for SELinux starts with the 3.9.40 release for standard mode and restricted mode and with the 3.9.42 release for private mode.

Note the following limitations:

- BlueXP does not support SELinux with Ubuntu hosts.
- Management of Cloud Volumes ONTAP systems is not supported by Connectors that have SELinux enabled on the operating system.

[Learn more about SELinux](#)

## 30 September 2024

### Private mode release (3.9.44)

A new private mode release is now available to download from the NetApp Support Site.

This release includes the following versions of the BlueXP components and services that are supported with private mode.

Service	Version included
Connector	3.9.44
Backup and recovery	27 September 2024
Classification	15 May 2024 (version 1.31)
Cloud Volumes ONTAP management	9 September 2024
Digital wallet	30 July 2023
On-premises ONTAP cluster management	22 April 2024
Replication	18 Sept 2022

For the Connector, the 3.9.44 private mode release includes the updates introduced in the August 2024 and September 2024 releases. Most notably, support for Red Hat Enterprise Linux 9.4.

To learn more about what's included in the versions of these BlueXP components and services, refer to the release notes for each BlueXP service:

- [What's new in the September 2024 release of the Connector](#)
- [What's new in the August 2024 release of the Connector](#)
- [What's new with BlueXP backup and recovery](#)

- [What's new with BlueXP classification](#)
- [What's new with Cloud Volumes ONTAP management in BlueXP](#)

For more details about private mode, including how to upgrade, refer to the following:

- [Learn about private mode](#)
- [Learn how to get started with BlueXP in private mode](#)
- [Learn how to upgrade the Connector when using private mode](#)

## 9 September 2024

### Connector 3.9.44

This release includes support for Docker Engine 26, an enhancement to SSL certificates, and bug fixes.

The 3.9.44 release is available for standard mode and restricted mode.

#### Support for Docker Engine 26 with new installations

Starting with the 3.9.44 release of the Connector, Docker Engine 26 is now supported with *new* Connector installations on Ubuntu hosts.

If you have an existing Connector created prior to the 3.9.44 release, then Docker Engine 25.0.5 is still the maximum supported version on Ubuntu hosts.

[Learn more about Docker Engine requirements.](#)

#### Updated SSL certificate for local UI access

When you use BlueXP in restricted mode or private mode, the user interface is accessible from the Connector virtual machine that's deployed in your cloud region or on-premises. By default, BlueXP uses a self-signed SSL certificate to provide secure HTTPS access to the web-based console running on the Connector.

In this release, we made changes to the SSL certificate for new and existing Connectors:

- The Common Name for the certificate now matches the short host name
- The Certificate Subject Alternative Name is the Fully Qualified Domain Name (FQDN) of the host machine

#### Support for RHEL 9.4

BlueXP now supports installing the Connector on a Red Hat Enterprise Linux 9.4 host when using BlueXP in standard mode or restricted mode.

Support for RHEL 9.4 starts with the 3.9.40 release of the Connector.

The updated list of supported RHEL versions for standard mode and restricted mode now includes the following:

- 8.6 to 8.10
- 9.1 to 9.4

[Learn about support for RHEL 8 and 9 with the Connector.](#)

## Support for Podman 4.9.4 with all RHEL versions

Podman 4.9.4 is now supported with all supported versions of Red Hat Enterprise Linux. Version 4.9.4 was previously supported with only RHEL 8.10.

The updated list of supported Podman versions includes 4.6.1 and 4.9.4 with Red Hat Enterprise Linux hosts.

Podman is required for RHEL hosts starting with the 3.9.40 release of the Connector.

[Learn about support for RHEL 8 and 9 with the Connector.](#)

## Updated AWS and Azure permissions

We updated the AWS and Azure policies for the Connector to remove permissions that are no longer required. The permissions were related to BlueXP edge caching and discovery and management of Kubernetes clusters, which are no longer supported as of August, 2024.

- [Learn what changed in the AWS policy.](#)
- [Learn what changed in the Azure policy.](#)

## 22 August 2024

### Connector 3.9.43 patch

We updated the Connector to support the Cloud Volumes ONTAP 9.15.1 release.

Support for this release includes an update to the Connector policy for Azure. The policy now includes the following permissions:

```
"Microsoft.Compute/virtualMachineScaleSets/write",  
"Microsoft.Compute/virtualMachineScaleSets/read",  
"Microsoft.Compute/virtualMachineScaleSets/delete"
```

These permissions are required for Cloud Volumes ONTAP support of Virtual Machine Scale Sets. If you have existing Connectors and you want to use this new feature, you'll need to add these permissions to the custom roles that are associated with your Azure credentials.

- [Learn about the Cloud Volumes ONTAP 9.15.1 release](#)
- [View Azure permissions for the Connector.](#)

## 8 August 2024

### Connector 3.9.43

This release includes minor improvements and bug fixes.

The 3.9.43 release is available for standard mode and restricted mode.

## Updated CPU and RAM requirements

To provide higher reliability and to improve the performance of BlueXP and the Connector, we now require

additional CPU and RAM for the Connector virtual machine:

- CPU: 8 cores or 8 vCPUs (the previous requirement was 4)
- RAM: 32 GB (the previous requirement was 14 GB)

As a result of this change, the default VM instance type when deploying the Connector from BlueXP or from the cloud provider's marketplace is as follows:

- AWS: t3.2xlarge
- Azure: Standard\_D8s\_v3
- Google Cloud: n2-standard-8

The updated CPU and RAM requirements apply to all new Connectors. For existing Connectors, increasing the CPU and RAM is recommended to provide improved performance and reliability.

### **Support for Podman 4.9.4 with RHEL 8.10**

Podman version 4.9.4 is now supported when installing the Connector on a Red Hat Enterprise Linux 8.10 host.

### **User validation for identity federation**

If you use identity federation with BlueXP, each user who logs in to BlueXP for the first time will need to complete a quick form to validate their identity.

## **31 July 2024**

### **Private mode release (3.9.42)**

A new private mode release is now available to download from the NetApp Support Site.

### **Support for RHEL 8 and 9**

This release includes support for installing the Connector on a Red Hat Enterprise Linux 8 or 9 host when using BlueXP in private mode. The following versions of RHEL are supported:

- 8.6 to 8.10
- 9.1 to 9.3

Podman is required as the container orchestration tool for these operating systems.

You should be aware of Podman requirements, known limitations, a summary of operating system support, what to do if you have a RHEL 7 host, how to get started, and more.

[Learn about support for RHEL 8 and 9 with the Connector.](#)

### **Versions included in this release**

This release includes the following versions of the BlueXP services that are supported with private mode.

<b>Service</b>	<b>Version included</b>
Connector	3.9.42
Backup and recovery	18 July 2024

<b>Service</b>	<b>Version included</b>
Classification	1 July 2024 (version 1.33)
Cloud Volumes ONTAP management	10 June 2024
Digital wallet	30 July 2023
On-premises ONTAP cluster management	30 July 2023
Replication	18 Sept 2022

To learn more about what's included in the versions of these BlueXP services, refer to the release notes for each BlueXP service.

- [Learn about private mode](#)
- [Learn how to get started with BlueXP in private mode](#)
- [Learn how to upgrade the Connector when using private mode](#)
- [Learn what's new with BlueXP backup and recovery](#)
- [Learn what's new with BlueXP classification](#)
- [Learn what's new with Cloud Volumes ONTAP management in BlueXP](#)

## 15 July 2024

### Support for RHEL 8.10

BlueXP now supports installing the Connector on a Red Hat Enterprise Linux 8.10 host when using standard mode or restricted mode.

Support for RHEL 8.10 starts with the 3.9.40 release of the Connector.

[Learn about support for RHEL 8 and 9 with the Connector.](#)

## 8 July 2024

### Connector 3.9.42

This release includes minor improvements, bug fixes, and support for the Connector in the AWS Canada West (Calgary) region.

The 3.9.42 release is available for standard mode and restricted mode.

### Updated Docker Engine requirements

When the Connector is installed on an Ubuntu host, the minimum supported version of Docker Engine is now 23.0.6. It was previously 19.3.1.

The maximum supported version is still 25.0.5.

[View Connector host requirements.](#)

## Email verification now required

New users who sign up to BlueXP are now required to verify their email address before they can log in.

## 12 June 2024

### Connector 3.9.41

This release of the BlueXP Connector includes minor security improvements and bug fixes.

The 3.9.41 release is available for standard mode and restricted mode.

## 4 June 2024

### Private mode release (3.9.40)

A new private mode release is now available to download from the NetApp Support Site. This release includes the following versions of the BlueXP services that are supported with private mode.

Note that this private mode release does *not* include support for the Connector with Red Hat Enterprise Linux 8 and 9.

Service	Version included
Connector	3.9.40
Backup and recovery	17 May 2024
Classification	15 May 2024 (version 1.31)
Cloud Volumes ONTAP management	17 May 2024
Digital wallet	30 July 2023
On-premises ONTAP cluster management	30 July 2023
Replication	18 Sept 2022

To learn more about what's included in the versions of these BlueXP services, refer to the release notes for each BlueXP service.

- [Learn about private mode](#)
- [Learn how to get started with BlueXP in private mode](#)
- [Learn how to upgrade the Connector when using private mode](#)
- [Learn what's new with BlueXP backup and recovery](#)
- [Learn what's new with BlueXP classification](#)
- [Learn what's new with Cloud Volumes ONTAP management in BlueXP](#)

## 17 May 2024

### Connector 3.9.40

This release of the BlueXP Connector includes support for additional operating systems, minor security improvements, and bug fixes.



At this time, the 3.9.40 release is available for standard mode and restricted mode.

### Support for RHEL 8 and 9

The Connector is now supported on hosts running the following versions of Red Hat Enterprise Linux with *new* Connector installations when using BlueXP in standard mode or restricted mode:

- 8.6 to 8.9
- 9.1 to 9.3

Podman is required as the container orchestration tool for these operating systems.

You should be aware of Podman requirements, known limitations, a summary of operating system support, what to do if you have a RHEL 7 host, how to get started, and more.

[Learn about support for RHEL 8 and 9 with the Connector.](#)

### End of support for RHEL 7 and CentOS 7

On June 30, 2024, RHEL 7 will reach end of maintenance (EOM), while CentOS 7 will reach end of life (EOL). NetApp will continue to support the Connector on these Linux distributions until June 30, 2024.

[Learn what to do if you have an existing Connector running on RHEL 7 or CentOS 7.](#)

### AWS permissions update

In the 3.9.38 release, we updated the Connector policy for AWS to include the "ec2:DescribeAvailabilityZones" permission. This permission is now required to support AWS Local Zones with Cloud Volumes ONTAP.

- [View AWS permissions for the Connector.](#)
- [Learn more about support for AWS Local Zones](#)

## Known limitations

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

These limitations are specific to BlueXP set up and administration: the Connector, the software as a service (SaaS) platform, and more.

### Connector limitations

#### Possible conflict with IP addresses in the 172 range

BlueXP deploys the Connector with two interfaces that have IP addresses in the 172.17.0.0/16 and 172.18.0.0/16 ranges.

If your network has a subnet configured with either of these ranges, then you might experience connectivity failures from BlueXP. For example, discovering on-premises ONTAP clusters in BlueXP might fail.

See Knowledge Base article [BlueXP Connector IP conflict with existing network](#) for instructions on how to change the IP address of the Connector's interfaces.

## SSL decryption isn't supported

BlueXP doesn't support firewall configurations that have SSL decryption enabled. If SSL decryption is enabled, error messages appear in BlueXP and the Connector instance displays as inactive.

For enhanced security, you have the option to [install an HTTPS certificate signed by a certificate authority \(CA\)](#).

## Blank page when loading the local UI

If you load the web-based console that's running on a Connector, the interface might fail to display sometimes, and you just get a blank page.

This issue is related to a caching problem. The workaround is to use an incognito or private web browser session.

## Shared Linux hosts are not supported

The Connector isn't supported on a VM that is shared with other applications. The VM must be dedicated to the Connector software.

## 3rd-party agents and extensions

3rd-party agents or VM extensions are not supported on the Connector VM.

# Changes to supported Linux operating systems

As we add and remove support for the Connector on specific Linux operating systems, you might have questions about how this support affects your existing Connector deployments.

## Supported operating systems

NetApp supports the BlueXP Connector with the following Linux operating systems.

## **Standard mode**

### **Manual installation**

- Ubuntu 24.04 LTS
- Ubuntu 22.04 LTS
- Red Hat Enterprise Linux
  - 8.6 to 8.10
  - 9.1 to 9.4

### **Deployment from BlueXP**

Ubuntu 22.04 LTS

### **Deployment from the AWS Marketplace**

Ubuntu 22.04 LTS

### **Deployment from the Azure Marketplace**

Ubuntu 22.04 LTS

## **Restricted mode**

### **Manual installation**

- Ubuntu 24.04 LTS
- Ubuntu 22.04 LTS
- Red Hat Enterprise Linux
  - 8.6 to 8.10
  - 9.1 to 9.4

### **Deployment from the AWS Marketplace**

Ubuntu 22.04 LTS

### **Deployment from the Azure Marketplace**

Ubuntu 22.04 LTS

## **Private mode**

### **Manual installation**

- Ubuntu 22.04 LTS
- Red Hat Enterprise Linux
  - 8.6 to 8.10
  - 9.1 to 9.4

## **Support for RHEL 8 and 9**

Note the following about support for RHEL 8 and 9:

### **Limitations**

BlueXP classification is supported if you install the Connector on a RHEL 8 or 9 host that resides on-premises. It's not supported if the RHEL 8 or 9 host resides in AWS, Azure, or Google Cloud.

## Container orchestration tool

You must use the Podman tool as the container orchestration tool when installing the Connector on a RHEL 8 or 9 host. Docker Engine is not supported with RHEL 8 and 9.

## Deployment mode

RHEL 8 and 9 are supported when using BlueXP in standard mode, restricted mode, and private mode.

## Supported Connector versions

NetApp supports RHEL 8 and 9 beginning with the following versions of the Connector:

- 3.9.40 when using BlueXP in standard mode or restricted mode
- 3.9.42 when using BlueXP in private mode

## New manual installations only

RHEL 8 and 9 are supported with *new* Connector installations when manually installing the Connector on hosts running on your premises or in the cloud.

## RHEL upgrades

If you have an existing Connector running on a RHEL 7 host, we don't support upgrading the RHEL 7 operating system to RHEL 8 or 9. [Learn more about existing Connectors on RHEL 7 or CentOS 7.](#)

## End of support for RHEL 7 and CentOS 7

On June 30, 2024, RHEL 7 reached end of maintenance (EOM), while CentOS 7 reached end of life (EOL). NetApp stopped supporting the Connector on these Linux distributions on June 30, 2024.

[Red Hat: What to know about Red Hat Enterprise Linux 7 End of Maintenance](#)

## Existing Connectors on RHEL 7 or CentOS 7

If you have an existing Connector running on RHEL 7 or CentOS 7, we don't support upgrading or converting the operating system to RHEL 8 or 9. You need to create a new Connector on a supported operating system.

1. Set up a RHEL 8 or 9 host.
2. Install Podman.
3. Perform a *new* Connector installation.
4. Configure the Connector to discover the working environments that the old Connector was managing.

## Related information

### How to get started with RHEL 8 and 9

Refer to the following pages for details about host requirements, Podman requirements, and steps to install Podman and the Connector:

**Standard mode**

- [Install and set up a Connector on-premises](#)
- [Manually install the Connector in AWS](#)
- [Manually install the Connector in Azure](#)
- [Manually install the Connector in Google Cloud](#)

**Restricted mode**

[Prepare for deployment in restricted mode](#)

**Private mode**

[Prepare for deployment in private mode](#)

**How to rediscover your working environments**

Refer to the following pages to rediscover your working environments after a new Connector deployment.

- [Add existing Cloud Volumes ONTAP systems to BlueXP](#)
- [Discover on-premises ONTAP clusters](#)
- [Create or discover an FSx for ONTAP working environment](#)
- [Create an Azure NetApp Files working environment](#)
- [Discover E-Series systems](#)
- [Discover StorageGRID systems](#)

# Get started

## Learn the basics

### Learn about BlueXP

NetApp BlueXP provides your organization with a single control plane that helps you build, protect, and govern data across your on-premises and cloud environments. The BlueXP software as a service (SaaS) platform includes services that provide storage management, data mobility, data protection, and data analysis and control. Management capabilities are provided through a web-based console and APIs.

### Features

BlueXP provides unified control of storage across your hybrid multicloud and integrated data services to protect, secure, and optimize data.

#### Unified control of storage from the BlueXP canvas

The *BlueXP canvas* lets you discover, deploy, and manage cloud and on-premises storage. The canvas centralizes storage management.

#### Supported cloud and on-premises storage

BlueXP enables you to manage the following types of storage from the BlueXP canvas:

##### Cloud storage solutions

- Amazon FSx for NetApp ONTAP
- Azure NetApp Files
- Cloud Volumes ONTAP
- Google Cloud NetApp Volumes

##### On-premises flash and object storage

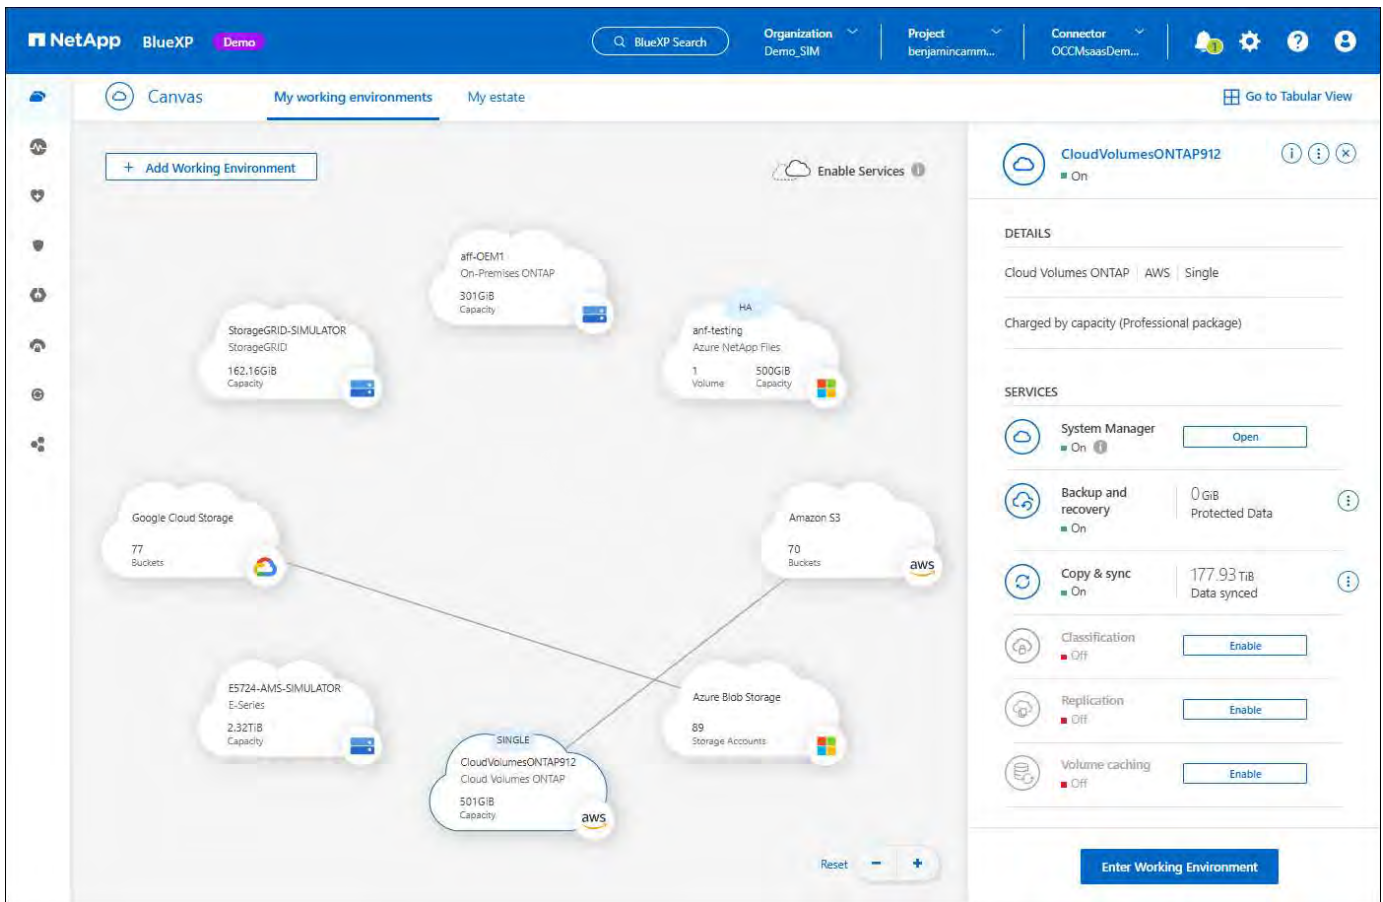
- E-Series systems
- ONTAP clusters
- StorageGRID systems

##### Cloud object storage

- Amazon S3 storage
- Azure Blob storage
- Google Cloud Storage

#### Storage management from working environments

On the BlueXP canvas, *working environments* represent discovered or deployed storage. You can select a *working environment* to integrate it with BlueXP data services or manage storage, such as adding volumes.



## Integrated services to protect, secure, and optimize data

BlueXP includes data services to secure and maintain data availability across storage.

### BlueXP alerts

View issues related to capacity, availability, performance, protection, and security in your ONTAP environment.

### BlueXP automation catalog

Use scripted solutions to automate the deployment and integration of NetApp products and services.

### BlueXP backup and recovery

Back up and restore cloud and on-premises data.

### BlueXP classification

Get your application data and cloud environments privacy ready.

### BlueXP copy and sync

Sync data between on-premises and cloud data stores.

### BlueXP digital advisor

Use predictive analytics and proactive support to optimize your data infrastructure.

### BlueXP digital wallet

Manage and monitor your licenses and subscriptions.

### **BlueXP disaster recovery**

Protect on-premises VMware workloads using VMware Cloud on Amazon FSx for ONTAP as a disaster recovery site.

### **BlueXP economic efficiency**

Identify clusters with current or forecasted low capacity and implement data tiering or additional capacity recommendations.

### **BlueXP ransomware protection**

Detect anomalies that might result in ransomware attacks. Protect and recover workloads.

### **BlueXP replication**

Replicate data between storage systems to support backup and disaster recovery.

### **BlueXP software updates**

Automate the assessment, planning, and execution of ONTAP upgrades.

### **BlueXP sustainability dashboard**

Analyze the sustainability of your storage systems.

### **BlueXP tiering**

Extend your on-premises ONTAP storage to the cloud.

### **BlueXP volume caching**

Create a writable cache volume to speed up access to data or offload traffic from heavily accessed volumes.

### **BlueXP workload factory**

Design, set up, and operate key workloads using Amazon FSx for NetApp ONTAP.

[Learn more about BlueXP and the available data services](#)

### **Supported cloud providers**

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

### **Cost**

Pricing for BlueXP depends on the services that you use.

[Learn about BlueXP pricing](#)

### **How BlueXP works**

BlueXP includes a web-based console that's provided through the SaaS layer, a resource and access management system, Connectors that manage working environments and enable BlueXP cloud services, and different deployment modes to meet your business requirements.

### **Software-as-a-service**

BlueXP is accessible through a [web-based console](#) and APIs. This SaaS experience enables you to automatically access the latest features as they're released and to easily switch between your BlueXP organizations, projects, and Connectors.



## BlueXP identity and access management (IAM)

BlueXP identity and access management (IAM) is a resource and access management model that provides granular management of resources and permissions:

- A top-level *organization* enables you to manage access across your various *projects*
- *Folders* enable you to group related projects together
- Resource management enables you to associate a resource with one or more folders or projects
- Access management enables you to assign a role to members at different levels of the organization hierarchy

BlueXP IAM is supported when using BlueXP in standard or restricted mode. If you're using BlueXP in private mode, then you use a BlueXP *account* to manage workspaces, users, and resources.

- [Learn more about BlueXP IAM](#)

## Connectors

You don't need a Connector to get started with BlueXP, but you'll need to create a Connector to unlock all BlueXP features and services. A Connector enables you to manage resources and processes across your on-premises and cloud environments. You need it to manage working environments (for example, Cloud Volumes ONTAP) and to use many BlueXP services.

[Learn more about Connectors.](#)

## Deployment modes

BlueXP offers three deployment modes. *Standard mode* leverages the BlueXP software as a service (SaaS) layer to provide full functionality. If your environment has security and connectivity restrictions, *restricted mode* and *private mode* limit outbound connectivity to the BlueXP SaaS layer.

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

## SOC 2 Type 2 certification

An independent certified public accountant firm and services auditor examined BlueXP and affirmed that BlueXP achieved SOC 2 Type 2 reports based on the applicable Trust Services criteria.

[View NetApp's SOC 2 reports](#)

## Learn about BlueXP Connectors

A *Connector* is NetApp software running in your cloud network or on-premises network. It's used to connect BlueXP's services to your storage environments.

### What you can do without a Connector

A Connector isn't required to get started with BlueXP. You can use several features and services within BlueXP without ever creating a Connector.

You can use the following BlueXP features and services without a Connector:

- Amazon FSx for NetApp ONTAP

Some actions require a Connector or a BlueXP workload factory link. [Learn which actions require a Connector or link](#)

- Automation catalog
- Azure NetApp Files

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

- Cloud Volumes Service for Google Cloud
- Copy and sync
- Digital advisor
- Digital wallet (licenses only, subscription monitoring requires a Connector)

In almost all cases, you can add a license to the digital wallet without a Connector.

The only time that a Connector is required to add a license to the digital wallet is for Cloud Volumes ONTAP *node-based* licenses. A Connector is required in this case because the data is taken from the licenses installed on Cloud Volumes ONTAP systems.

- Direct discovery of on-premises ONTAP clusters

While a Connector isn't required for direct discovery of an on-premises ONTAP cluster, a Connector is required if you want to take advantage of additional BlueXP features.

[Learn more about discovery and management options for on-premises ONTAP clusters](#)

- Software updates
- Sustainability
- Workload factory

## When a Connector is required

When you use BlueXP in standard mode, a Connector is required for the following features and services in BlueXP:

- Alerts
- Amazon FSx for ONTAP management features
- Amazon S3 storage
- Azure Blob storage
- Backup and recovery
- Classification
- Cloud Volumes ONTAP
- Disaster recovery
- E-Series systems
- Economic efficiency <sup>1</sup>
- Google Cloud Storage buckets

- On-premises ONTAP cluster integration with BlueXP data services
- Ransomware protection
- StorageGRID systems
- Tiering
- Volume caching

<sup>1</sup> While you can access these services without a Connector, a Connector is required to initiate actions from the services.

A Connector is required to use BlueXP in restricted mode or private mode.

### **Connectors must be operational at all times**

Connectors are a fundamental part of the BlueXP service architecture. It's your responsibility to ensure that relevant Connectors are up, operational, and accessible at all times. While the service is designed to overcome short outages of Connector availability, you must take immediate action when required to remedy infrastructure failures.

This documentation is governed by the EULA. If the product is not operated in accordance with the documentation, the functionality and operation of the product, as well as your rights under the EULA, might be adversely impacted.

### **Supported locations**

A Connector is supported in the following locations:

- Amazon Web Services
- Microsoft Azure

A Connector in Azure should be deployed in the same Azure region as the Cloud Volumes ONTAP systems that it manages, or in the [Azure region pair](#) for the Cloud Volumes ONTAP systems. This requirement ensures that an Azure Private Link connection is used between Cloud Volumes ONTAP and its associated storage accounts. [Learn how Cloud Volumes ONTAP uses an Azure Private Link](#)

- Google Cloud

If you want to use BlueXP services with Google Cloud, then you must use a Connector that's running in Google Cloud.

- On your premises

### **Communication with cloud providers**

The Connector uses TLS 1.3 for all communication to AWS, Azure, and Google Cloud.

### **Restricted mode and private mode**

To use BlueXP in restricted mode or private mode, you get started with BlueXP by installing the Connector and then accessing the user interface that's running locally on the Connector.

[Learn about BlueXP deployment modes.](#)

## How to install a Connector

You can install a Connector directly from BlueXP, from your cloud provider's marketplace, or by manually installing the software on your own Linux host. How you get started depends on whether you're using BlueXP in standard mode, restricted mode, or private mode.

- [Learn about BlueXP deployment modes](#)
- [Get started with BlueXP in standard mode](#)
- [Get started with BlueXP in restricted mode](#)
- [Get started with BlueXP in private mode](#)

## Permissions

Specific permissions are needed to create the Connector directly from BlueXP and another set of permissions are needed for the Connector instance itself. If you create the Connector in AWS or Azure directly from BlueXP, then BlueXP creates the Connector with the permissions that it needs.

When using BlueXP in standard mode, how you provide permissions depends on how you plan to create the Connector.

To learn how to set up permissions, refer to the following:

- Standard mode
  - [Connector installation options in AWS](#)
  - [Connector installation options in Azure](#)
  - [Connector installation options in Google Cloud](#)
  - [Set up cloud permissions for on-premises deployments](#)
- [Set up permissions for restricted mode](#)
- [Set up permissions for private mode](#)

To view the exact permissions that the Connector needs for day-to-day operations, refer to the following pages:

- [Learn how the Connector uses AWS permissions](#)
- [Learn how the Connector uses Azure permissions](#)
- [Learn how the Connector uses Google Cloud permissions](#)

It's your responsibility to update the Connector policies as new permissions are added in subsequent releases. If new permissions are required, they will be listed in the release notes.

## Connector upgrades

We typically update the Connector software each month to introduce new features and to provide stability improvements. While most of the services and features in the BlueXP platform are offered through SaaS-based software, a few features are dependent on the version of the Connector. That includes Cloud Volumes ONTAP management, on-premises ONTAP cluster management, settings, and help.

When you use BlueXP in standard mode or restricted mode, the Connector automatically updates its software to the latest version, as long as it has outbound internet access to obtain the software update. If you're using BlueXP in private mode, then you'll need to manually upgrade the Connector.

[Learn how to manually upgrade the Connector software when using private mode.](#)

## Operating system and VM maintenance

Maintaining the operating system on the Connector host is your (customer's) responsibility. For example, you (customer) should apply security updates to the operating system on the Connector host by following your company's standard procedures for operating system distribution.

Note that you (customer) don't need to stop any services on the Connector host when applying minor security updates.

If you (customer) need to stop and then start the Connector VM, you should do so from your cloud provider's console or by using the standard procedures for on-premises management.

[Be aware that the Connector must be operational at all times.](#)

## Multiple working environments and Connectors

A Connector can manage multiple working environments in BlueXP. The maximum number of working environments that a single Connector should manage varies. It depends on the type of working environments, the number of volumes, the amount of capacity being managed, and the number of users.

If you have a large-scale deployment, work with your NetApp representative to size your environment. If you experience any issues along the way, reach out to us by using the in-product chat.

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

Here are a few examples:

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

## Learn about BlueXP deployment modes

BlueXP offers *deployment modes* that enable you to meet your business and security requirements. *Standard mode* leverages a software as a service (SaaS) layer to provide full functionality, while *restricted mode* and *private mode* are available for organizations that have connectivity restrictions.

While BlueXP inhibits the flow of traffic, communication, and data when using restricted mode or private mode, it's your responsibility to ensure that your environment (on-premises and in the cloud) is in compliance with the required regulations for your business.

### Overview

Each deployment mode differs in outbound connectivity, location, installation, authentication, data services, and charging methods.

## Standard mode

You use a SaaS service from the web-based console. Depending on the data services and features that you plan to use, a BlueXP admin creates one or more Connectors to manage data within your hybrid cloud environment.

This mode uses encrypted data transmission over the public internet.

## Restricted mode

You install a BlueXP Connector in the cloud (in a government, sovereign, or commercial region), and it has limited outbound connectivity to the BlueXP SaaS layer.

This mode is typically used by state and local governments and regulated companies.

[Learn more about outbound connectivity to the SaaS layer.](#)

## Private mode

You install a BlueXP Connector on-premises or in the cloud (in a secure region, sovereign cloud region, or commercial region) and has *no* connectivity to the BlueXP SaaS layer. Users access the BlueXP console provided by the Connector locally, not the SaaS layer.

A secure region includes [AWS Secret Cloud](#), [AWS Top Secret Cloud](#), and [Azure IL6](#)

The following table provides a comparison of these modes.

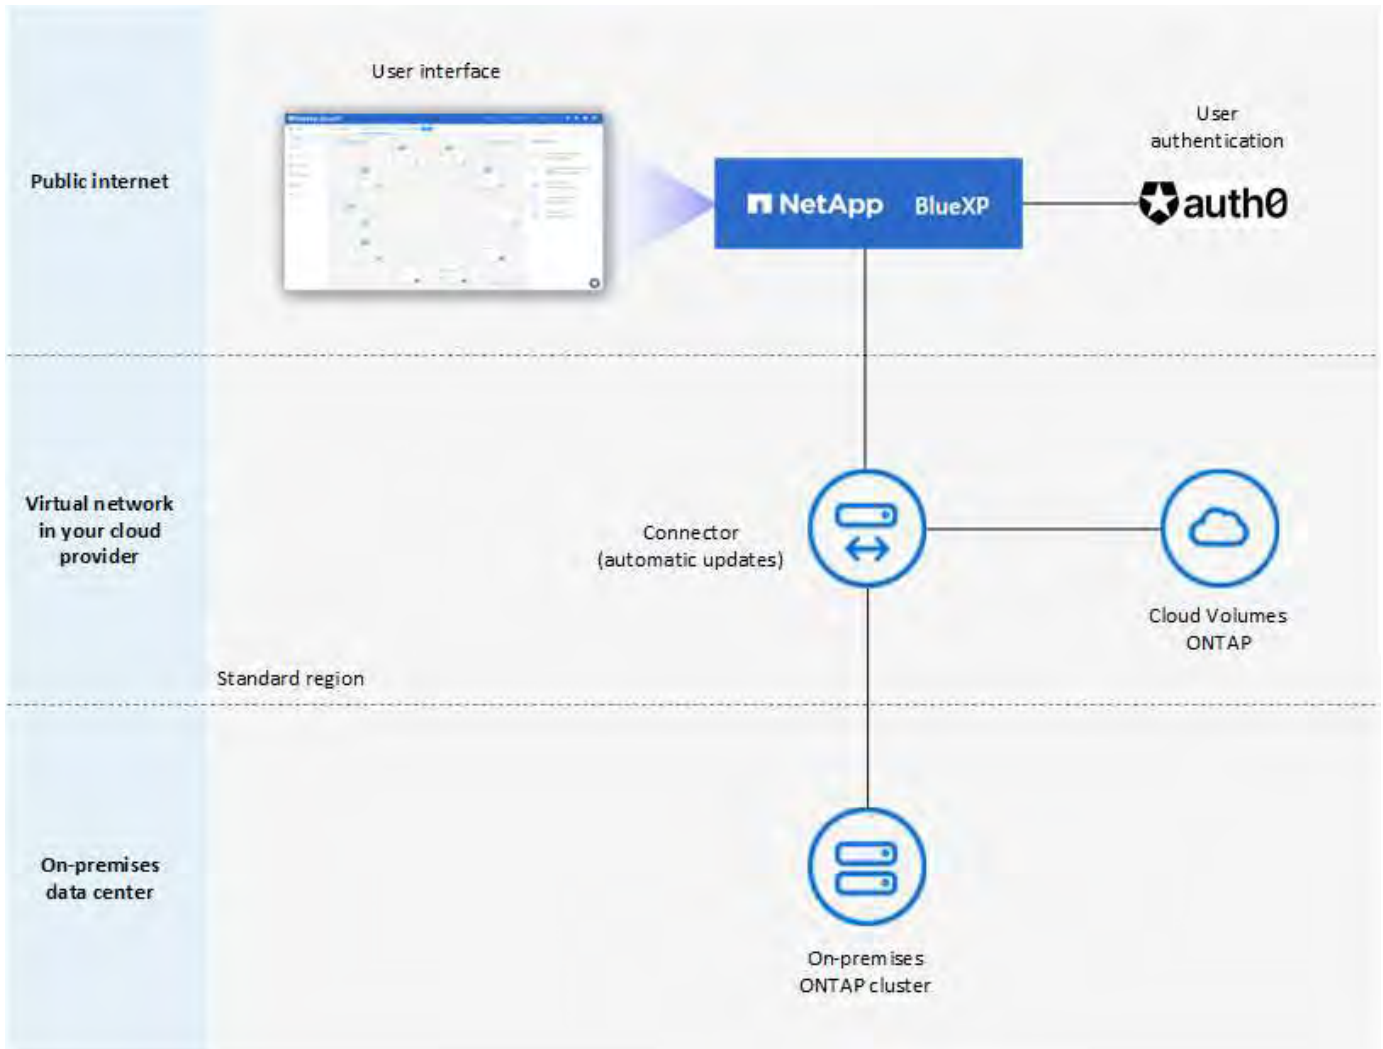
	Standard mode	Restricted mode	Private mode
<b>Connection required to BlueXP SaaS layer?</b>	Yes	Outbound only	No
<b>Connection required to your cloud provider?</b>	Yes	Yes, within the region	Yes, within the region (if using Cloud Volumes ONTAP)
<b>Connector installation</b>	From BlueXP, cloud marketplace, or manual install	Cloud marketplace or manual install	Manual install
<b>Connector upgrades</b>	Automatic upgrades of NetApp Connector software	Automatic upgrades of NetApp Connector software	Manual upgrade required
<b>UI access</b>	From the BlueXP SaaS layer	Locally from the Connector VM	Locally from the Connector VM
<b>API endpoint</b>	The BlueXP SaaS layer	The Connector	The Connector
<b>Authentication</b>	Through SaaS using auth0, NSS login, or identity federation	Through SaaS using auth0 or identity federation	Local user authentication
<b>Multi-factor authentication</b>	Available for local users	Not available	Not available
<b>Storage and data services</b>	All are supported	Many are supported	Several are supported

	Standard mode	Restricted mode	Private mode
Data service licensing options	Marketplace subscriptions and BYOL	Marketplace subscriptions and BYOL	BYOL

Read through the following sections to learn more about these modes, including which BlueXP features and services are supported.

### Standard mode

The following image is an example of a standard mode deployment.



BlueXP works as follows in standard mode:

### Outbound communication

Connectivity is required from the Connector to the BlueXP SaaS layer, to your cloud provider's publicly available resources, and to other essential components for day-to-day operations.

- [Endpoints that the Connector contacts in AWS](#)
- [Endpoints that the Connector contacts in Azure](#)

- [Endpoints that the Connector contacts in Google Cloud](#)

## Supported location for the Connector

In standard mode, the Connector is supported in the cloud or on your premises.

## Connector installation

You can install the Connector using the BlueXP setup wizard, AWS or Azure Marketplace, the Google Cloud SDK, or a manual installer on a Linux host in your data center or cloud.

## Connector upgrades

BlueXP provides automated upgrades of the Connector software with monthly updates.

## User interface access

The user interface is accessible from the web-based console that's provided through the SaaS layer.

## API endpoint

API calls are made to the following endpoint:  
<https://cloudmanager.cloud.netapp.com>

## Authentication

BlueXP provides authentication with auth0 or NetApp Support Site (NSS) logins. Identity federation is available.

## Supported BlueXP services

All BlueXP services are available to users.

## Supported licensing options

Marketplace subscriptions and BYOL are supported with standard mode; however, the supported licensing options depends on which BlueXP service you are using. Review the documentation for each service to learn more about the available licensing options.

## How to get started with standard mode

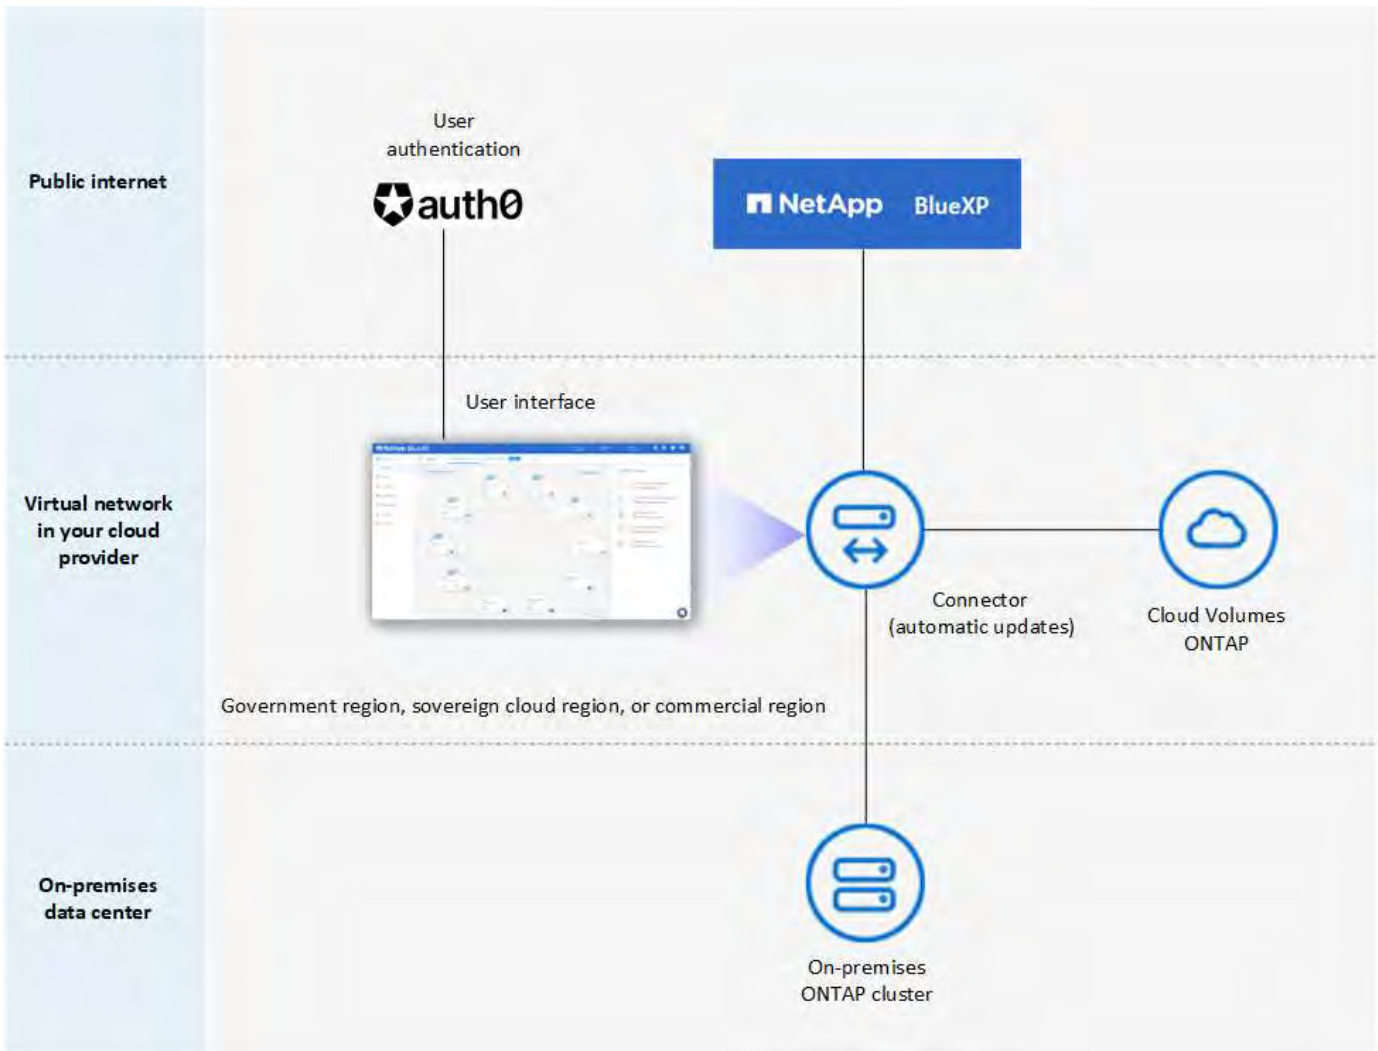
Go to the [BlueXP web-based console](#) and sign up.

[Learn how to get started with standard mode.](#)

## Restricted mode

The following image is an example of a restricted mode deployment.





BlueXP works as follows in restricted mode:

### Outbound communication

The Connector requires outbound connectivity to the BlueXP SaaS layer for data services, software upgrades, authentication, and metadata transmission.

The BlueXP SaaS layer does not initiate communication to the Connector. All communication is initiated by the Connector, which can pull or push data from or to the SaaS layer as required.

A connection is also required to cloud provider resources from within the region.

### Supported location for the Connector

In restricted mode, the Connector is supported in the cloud: in a government region, sovereign region, or commercial region.

### Connector installation

Connector installation is possible from the AWS or Azure Marketplace or a manual installation on your own Linux host.

### Connector upgrades

BlueXP provides automated upgrades of the Connector software with monthly updates.

## User interface access

The user interface is accessible from the Connector virtual machine that's deployed in your cloud region.

## API endpoint

API calls are made to the Connector virtual machine.

## Authentication

Authentication is provided through BlueXP's cloud service using auth0. Identity federation is also available.

## Supported BlueXP services

BlueXP supports the following storage and data services with restricted mode:

Supported services	Notes
Azure NetApp Files	Full support
Backup and recovery	Supported in Government regions and commercial regions with restricted mode. Not supported in sovereign regions with restricted mode.  In restricted mode, BlueXP backup and recovery supports back up and restore of ONTAP volume data only. <a href="#">View the list of supported backup destinations for ONTAP data</a> In restricted mode, BlueXP backup and recovery supports back up and restore of ONTAP volume data only. <a href="#">View the list of supported backup destinations for ONTAP data</a>  Back up and restore of application data and virtual machine data is not supported.
Classification	Supported in Government regions with restricted mode. Not supported in commercial regions or in sovereign regions with restricted mode.
Cloud Volumes ONTAP	Full support
Digital wallet	You can use the digital wallet with the supported licensing options listed below for restricted mode.
On-premises ONTAP clusters	Discovery with a Connector and discovery without a Connector (direct discovery) are both supported.  When you discover an on-premises cluster with a Connector, the Advanced view (System Manager) is not supported.
Replication	Supported in Government regions with restricted mode. Not supported in commercial regions or in sovereign regions with restricted mode.

## Supported licensing options

The following licensing options are supported with restricted mode:

- Marketplace subscriptions (hourly and annual contracts)

Note the following:

- For Cloud Volumes ONTAP, only capacity-based licensing is supported.
- In Azure, annual contracts are not supported with government regions.

- BYOL

For Cloud Volumes ONTAP, both capacity-based licensing and node-based licensing are supported with BYOL.

### How to get started with restricted mode

You need to enable restricted mode when you create your BlueXP account.

If you don't have an organization yet, you are prompted to create your organization and enable restricted mode when you log in to BlueXP for the first time from a Connector that you manually installed or that you created from your cloud provider's marketplace.

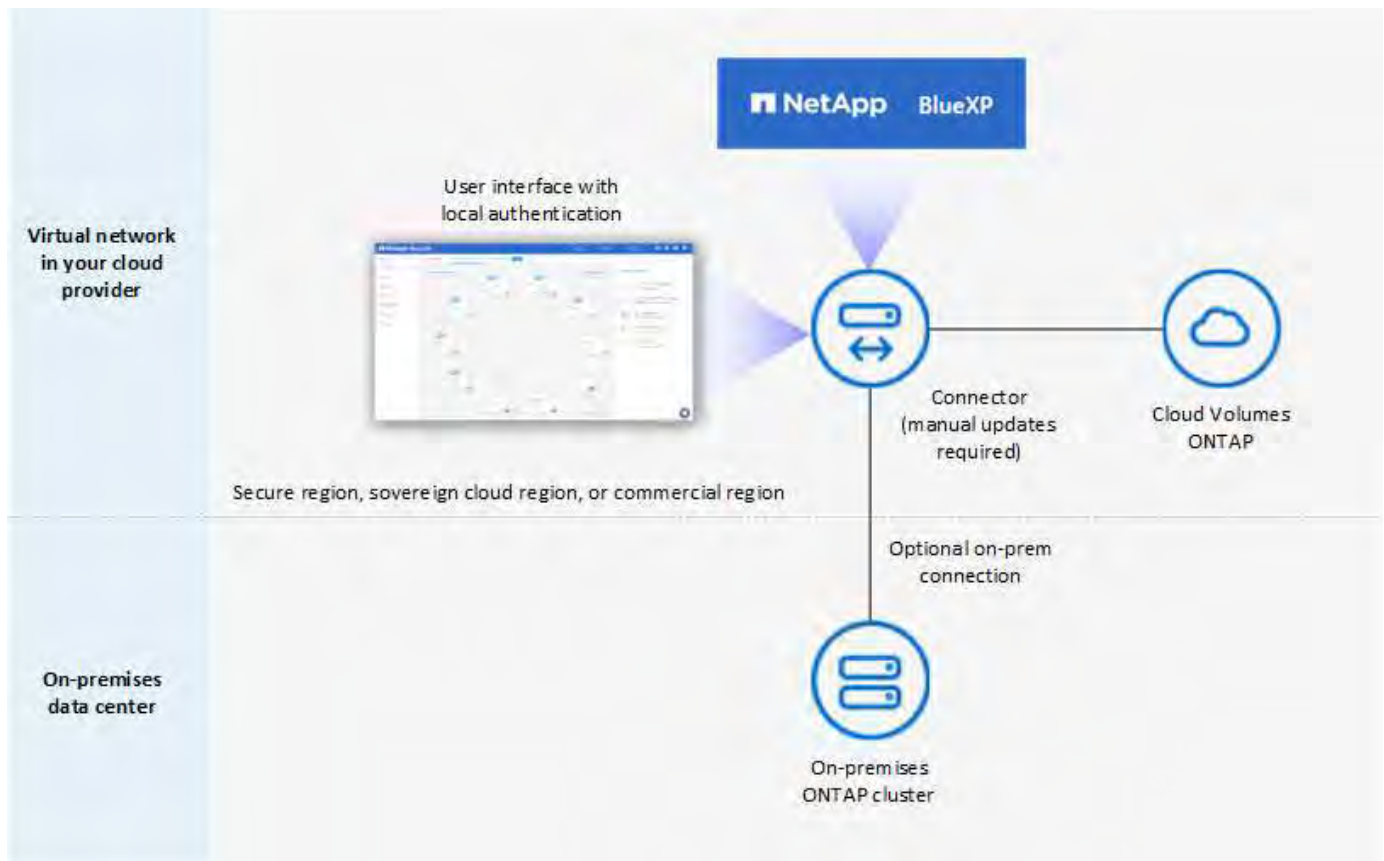
Note that you can't change the restricted mode setting after BlueXP creates the organization. You can't enable restricted mode later and you can't disable it later.

- [Learn how to get started with restricted mode.](#)

### Private mode

In private mode, you can install a Connector either on-premises or in the cloud and then use BlueXP to manage data across your hybrid cloud. There is no connectivity to the BlueXP SaaS layer.

The following image shows an example of a private mode deployment where the Connector is installed in the cloud and manages both Cloud Volumes ONTAP and an on-premises ONTAP cluster.



Meanwhile, the second image shows an example of a private mode deployment where the Connector is installed on-premises, manages an on-premises ONTAP cluster, and provides access to supported BlueXP data services.



BlueXP works as follows in private mode:

### Outbound communication

No outbound connectivity is required to the BlueXP SaaS layer. All packages, dependencies, and essential components are packaged with the Connector and served from the local machine. Connectivity to your cloud provider's publicly available resources is required only if you are deploying Cloud Volumes ONTAP.

### Supported location for the Connector

In private mode, the Connector is supported in the cloud or on-premises.

### Connector installation

Manual installations of the Connector are supported on your own Linux host in the cloud or on-premises.

### Connector upgrades

You need to upgrade the Connector software manually. The Connector software is published to the NetApp Support Site at undefined intervals.

### User interface access

The user interface is accessible from the Connector that's deployed in your cloud region or on-premises.

### API endpoint

API calls are made to the Connector virtual machine.

### Authentication

Authentication is provided through local user management and access. Authentication is not provided through BlueXP's cloud service.

### Supported BlueXP services in cloud deployments

BlueXP supports the following storage and data services with private mode when the Connector is installed in the cloud:

Supported services	Notes
Backup and recovery	<p>Supported in AWS and Azure commercial regions.</p> <p>Not supported in Google Cloud or in <a href="#">AWS Secret Cloud</a>, <a href="#">AWS Top Secret Cloud</a>, or <a href="#">Azure IL6</a></p> <p>In private mode, BlueXP backup and recovery supports back up and restore of ONTAP volume data only. <a href="#">View the list of supported backup destinations for ONTAP data</a></p> <p>Back up and restore of application data and virtual machine data is not supported.</p>
Cloud Volumes ONTAP	Because there's no internet access, the following features aren't available: automated software upgrades and AutoSupport.
Digital wallet	You can use the digital wallet with the supported licensing options listed below for private mode.
On-premises ONTAP clusters	<p>Requires connectivity from the cloud (where the Connector is installed) to the on-premises environment.</p> <p>Discovery without a Connector (direct discovery) is not supported.</p>

### Supported BlueXP services in on-premises deployments

BlueXP supports the following storage and data services with private mode when the Connector is installed on your premises:

Supported services	Notes
Backup and recovery	<p>In private mode, BlueXP backup and recovery supports back up and restore of ONTAP volume data only. <a href="#">View the list of supported backup destinations for ONTAP volume data</a></p> <p>Back up and restore of application data and virtual machine data is not supported.</p>
Classification	<ul style="list-style-type: none"> <li>The only supported data sources are the ones that you can discover locally. <a href="#">View the sources that you can discover locally</a></li> <li>Features that require outbound internet access are not supported. <a href="#">View the feature limitations</a></li> </ul>
Digital wallet	You can use the digital wallet with the supported licensing options listed below for private mode.
On-premises ONTAP clusters	Discovery without a Connector (direct discovery) is not supported.
Replication	Full support

### Supported licensing options

Only BYOL is supported with private mode.

For Cloud Volumes ONTAP BYOL, only node-based licensing is supported. Capacity-based licensing is not supported. Because an outbound internet connection isn't available, you need to manually upload your Cloud Volumes ONTAP licensing file in the BlueXP digital wallet.

[Learn how to add licenses to the BlueXP digital wallet](#)

### How to get started with private mode

Private mode is available by downloading the "offline" installer from the NetApp Support Site.

[Learn how to get started with private mode.](#)



If you want to use BlueXP in the [AWS Secret Cloud](#) or the [AWS Top Secret Cloud](#), then you should follow separate instructions to get started in those environments. [Learn how to get started with Cloud Volumes ONTAP in the AWS Secret Cloud or Top Secret Cloud](#)

### Service and feature comparison

The following table can help you quickly identify which BlueXP services and features are supported with restricted mode and private mode.

Note that some services might be supported with limitations. For more details about how these services are supported with restricted mode and private mode, refer to the sections above.

Product area	BlueXP service or feature	Restricted mode	Private mode
<b>Working environments</b>  This portion of the table lists support for working environment management from the BlueXP canvas. It does not indicate the supported backup destinations for BlueXP backup and recovery.	Amazon FSx for ONTAP	No	No
	Amazon S3	No	No
	Azure Blob	No	No
	Azure NetApp Files	Yes	No
	Cloud Volumes ONTAP	Yes	Yes
	Google Cloud NetApp Volumes	No	No
	Google Cloud Storage	No	No
	On-premises ONTAP clusters	Yes	Yes
	E-Series	No	No
	StorageGRID	No	No

Product area	BlueXP service or feature	Restricted mode	Private mode
<b>Services</b>	Alerts	No	No
	Backup and recovery	Yes <a href="#">View the list of supported backup destinations for ONTAP volume data</a>	Yes <a href="#">View the list of supported backup destinations for ONTAP volume data</a>
	Classification	Yes	Yes
	Copy and sync	No	No
	Digital advisor	No	No
	Digital wallet	Yes	Yes
	Disaster recovery	No	No
	Economic efficiency	No	No
	Ransomware protection	No	No
	Replication	Yes	Yes
	Software updates	No	No
	Sustainability	No	No
	Tiering	No	No
	Volume caching	No	No
	Workload factory	No	No
<b>Features</b>	Identity and access management	Yes	Yes
	Credentials	Yes	Yes
	Federation	Yes	No
	Multi-factor authentication	Yes	No
	NSS accounts	Yes	No
	Notifications	Yes	No
	Search	Yes	No
	Timeline	Yes	Yes

## Get started with standard mode

### Getting started workflow (standard mode)

Get started with BlueXP in standard mode by preparing networking for the BlueXP console, signing up and creating an account, optionally creating a Connector, and subscribing to NetApp Intelligent Services.

In standard mode, you access a web-based console that is hosted as a Software-as-a-service (SaaS) product



from NetApp. Before you get started, you should have an understanding of [deployment modes](#) and [Connectors](#).

1

### Prepare networking for using the BlueXP console

Computers that access the BlueXP console should have connections to specific endpoints to complete a few administrative tasks. If your network restricts outbound access, you should ensure that these endpoints are allowed.

2

### Sign up and create an organization

Go to the [BlueXP console](#) and sign up. You'll be given the option to create an organization, but you can skip that step if you're being invited to an existing organization.

At this point, you're logged in and can start using several BlueXP services like Digital Advisor, Amazon FSx for ONTAP, Azure NetApp Files, and more. [Learn what you can do without a Connector](#).

3

### Create a Connector

You don't need a Connector to get started with BlueXP, but you can create a Connector to unlock all BlueXP features and services. The Connector is NetApp software that enables BlueXP to manage resources and processes within your hybrid cloud environment.

You can create a Connector in your cloud or on-premises network.

- [Learn more about when Connectors are required and how they work](#)
- [Learn how to create a Connector in AWS](#)
- [Learn how to create a Connector in Azure](#)
- [Learn how to create a Connector in Google Cloud](#)
- [Learn how to create a Connector on-premises](#)

Note that if you want to use NetApp Intelligent Data Services to manage storage and data in Google Cloud, then the Connector must be running in Google Cloud.

Note that if you want to use NetApp's intelligent data services to manage storage and data in Google Cloud, then the Connector must be running in Google Cloud.

4

### Subscribe to NetApp Intelligent Services (optional)

Subscribe to NetApp Intelligent Services from your cloud provider's marketplace to pay for data services at an hourly rate (PAYGO) or through an annual contract. NetApp Intelligent Services include Backup and recovery, Cloud Volumes ONTAP, Tiering, Ransomware protection, and Disaster recovery. Classification is included with your subscription at no additional cost.

## Prepare networking for the BlueXP console

When you log in and use the web-based console, BlueXP contacts several endpoints to complete the actions that you initiate. Computers that access the console must have connections to these endpoints.



These endpoints are contacted in two scenarios:

- From a user's computer when completing sections from the [BlueXP web-based console](#) that's available as software as a service (SaaS).
- From a user's computer when opening a web browser, entering the IP address of the Connector host, and then logging in and setting up the Connector. These steps are required if you manually install the Connector.

Endpoints	Purpose
<a href="https://console.bluexp.netapp.com">https://console.bluexp.netapp.com</a> <a href="https://*.console.bluexp.netapp.com">https://*.console.bluexp.netapp.com</a>	This is the endpoint that you enter in your web browser to use the web-based console.
<a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a>	The web-based console contacts this endpoint to interact with the API for actions related to authorization, licensing, subscriptions, credentials, notifications, and more.
<a href="https://aiq.netapp.com">https://aiq.netapp.com</a>	Required to access digital advisor.
AWS services (amazonaws.com): <ul style="list-style-type: none"> <li>• CloudFormation</li> <li>• Elastic Compute Cloud (EC2)</li> <li>• Key Management Service (KMS)</li> <li>• Security Token Service (STS)</li> <li>• Simple Storage Service (S3)</li> </ul>	Required to deploy a Connector from BlueXP in AWS. The exact endpoint depends on the region in which you deploy the Connector. <a href="#">See AWS documentation for details.</a>  Suggestion: <a href="#">See AWS documentation for details.</a>
<a href="https://management.azure.com">https://management.azure.com</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a>	Required to deploy a Connector from BlueXP in most Azure regions.
<a href="https://management.microsoftazure.de">https://management.microsoftazure.de</a> <a href="https://login.microsoftonline.de">https://login.microsoftonline.de</a>	Required to deploy a Connector from BlueXP in Azure Germany regions.
<a href="https://management.usgovcloudapi.net">https://management.usgovcloudapi.net</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a>	Required to deploy a Connector from BlueXP in Azure US Gov regions.
<a href="https://www.googleapis.com">https://www.googleapis.com</a>	Required to deploy a Connector from BlueXP in Google Cloud.
<a href="https://signin.b2c.netapp.com">https://signin.b2c.netapp.com</a>	Required to update NetApp Support Site (NSS) credentials or to add new NSS credentials to BlueXP.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://cdn.auth0.com">https://cdn.auth0.com</a> <a href="https://services.cloud.netapp.com">https://services.cloud.netapp.com</a>	Your web browser connects to these endpoints for centralized user authentication through BlueXP.
<a href="https://widget.intercom.io">https://widget.intercom.io</a>	For in-product chat that enables you to talk to NetApp support.

Ensure the Connector has outbound internet access to contact endpoints for daily operations. Follow the links in the next section below to find the list of these endpoints.

### Related information

- Prepare networking for the Connector
  - [Set up AWS networking](#)
  - [Set up Azure networking](#)
  - [Set up Google Cloud networking](#)
  - [Set up on-premises networking](#)
- Prepare networking for BlueXP services

Refer to the documentation for each BlueXP service.

[BlueXP documentation](#)

## Sign up or log in to BlueXP

BlueXP is accessible from a web-based console. When you get started with BlueXP, your first step is to sign up or to log in using your NetApp Support Site credentials or SSO credentials from your corporate directory.

### About this task

When you access BlueXP for the first time, BlueXP enables you to sign up or log in using one of the following options:

#### BlueXP login

You can sign up by creating a BlueXP login. This authentication method requires you to specify your email address and a password. After you verify your email address, you can log in and then create a BlueXP organization, if you don't already belong to one.

#### NetApp Support Site (NSS) credentials

If you have existing NetApp Support Site credentials, you don't need to sign up to BlueXP. You log in using your NSS credentials and then BlueXP prompts you to create a BlueXP organization, if you don't already belong to one.

Note that the default password experience is a one-time passcode (OTP) to the registered email address. A new OTP is generated with each sign-in attempt.

#### Federated connection

You can use single sign-on to log in using credentials from your corporate directory (federated identity). The first user in your organization's account must sign up to BlueXP or log in using NSS credentials, and then set up identity federation. After that, you can add members from your corporate identity to your organization. Those users can then log in using their SSO credentials.

[Learn how to use identity federation with BlueXP.](#)

### Steps

1. Open a web browser and go to the [BlueXP console](#)
2. If you have a NetApp Support Site account or if you already set up identity federation, enter the email address associated with your account directly on the **Log in** page.

In both of these cases, BlueXP will sign you up as part of this initial login.

3. If you want to sign up by creating a BlueXP login, select **Sign up**.
  - a. On the **Sign up** page, enter the required information and select **Next**.

Note that only English characters are allowed in the sign up form.

- b. Check your inbox for an email from NetApp that includes instructions to verify your email address.

This step is required before you can log in to BlueXP.

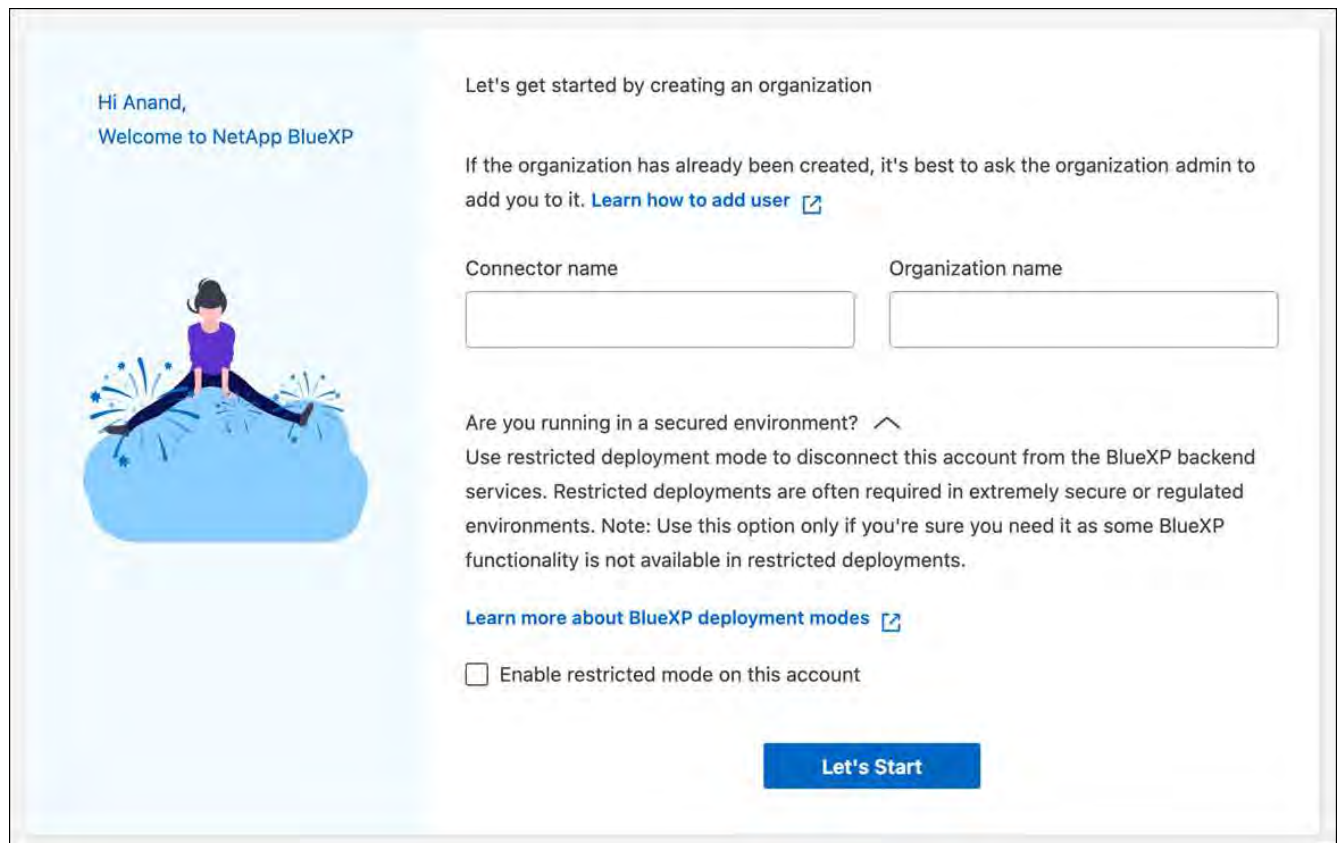
4. After you log in, review the End User License Agreement and accept the terms.

If your user account doesn't already belong to a BlueXP organization, you'll be prompted to create one.

5. On the **Welcome** page, enter a name for your BlueXP organization.

An organization is the top-level element in BlueXP identity and access management (IAM). [Learn about BlueXP IAM](#).

If your business already has a BlueXP organization and you want to join it, then you should close out of BlueXP and ask the owner to associate you with the organization. After the owner adds you, you can log in and you'll have access to the account. [Learn how to add members to an existing organization](#).



The screenshot shows a user interface for creating a BlueXP organization. On the left, there is a greeting: "Hi Anand, Welcome to NetApp BlueXP" and an illustration of a person sitting on a blue cloud. The main content area has the heading "Let's get started by creating an organization". Below this, there is a paragraph: "If the organization has already been created, it's best to ask the organization admin to add you to it. [Learn how to add user](#)". There are two input fields: "Connector name" and "Organization name". Below these fields, there is a section titled "Are you running in a secured environment?" with a caret icon. The text below reads: "Use restricted deployment mode to disconnect this account from the BlueXP backend services. Restricted deployments are often required in extremely secure or regulated environments. Note: Use this option only if you're sure you need it as some BlueXP functionality is not available in restricted deployments." There is a link: "[Learn more about BlueXP deployment modes](#)". At the bottom, there is a checkbox labeled "Enable restricted mode on this account" which is currently unchecked. A blue button labeled "Let's Start" is at the bottom right.

6. Select **Let's Start**.

## Result

You now have a BlueXP login and an organization. In most cases, the next step is to create a Connector, which connects BlueXP's services to your hybrid cloud environment.

# Create a Connector

## AWS

### Connector installation options in AWS

There are a few different ways to create a Connector in AWS. Directly from BlueXP is the most common way.

The following installation options are available:

- [Create the Connector directly from BlueXP](#) (this is the standard option)

This action launches an EC2 instance running Linux and the Connector software in a VPC of your choice.

- [Create a Connector from the AWS Marketplace](#)

This action also launches an EC2 instance running Linux and the Connector software, but the deployment is initiated directly from the AWS Marketplace, rather than from BlueXP.

- [Download and manually install the software on your own Linux host](#)

The installation option that you choose impacts how you prepare for installation. This includes how you provide BlueXP with the required permissions that it needs to authenticate and manage resources in AWS.

### Create a Connector in AWS from BlueXP

You can create a Connector in AWS directly from BlueXP. To create a Connector in AWS from BlueXP, you need to set up your networking, prepare AWS permissions, and then create the Connector.

#### Before you begin

- You should have an [understanding of Connectors](#).
- You should review [Connector limitations](#).

#### Step 1: Set up networking

Ensure that the network location where you plan to install the Connector supports the following requirements. Meeting these requirements enables the Connector to manage resources and processes within your hybrid cloud environment.

#### VPC and subnet

When you create the Connector, you need to specify the VPC and subnet where the Connector should reside.

#### Connections to target networks

A Connector requires a network connection to the location where you're planning to create and manage working environments. For example, the network where you plan to create Cloud Volumes ONTAP systems or a storage system in your on-premises environment.

#### Outbound internet access

The network location where you deploy the Connector must have an outbound internet connection to contact specific endpoints.

## Endpoints contacted from the Connector

The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment for day-to-day operations.

Note that the endpoints listed below are all CNAME entries.

Endpoints	Purpose
AWS services (amazonaws.com): <ul style="list-style-type: none"><li>• CloudFormation</li><li>• Elastic Compute Cloud (EC2)</li><li>• Identity and Access Management (IAM)</li><li>• Key Management Service (KMS)</li><li>• Security Token Service (STS)</li><li>• Simple Storage Service (S3)</li></ul>	To manage resources in AWS. The exact endpoint depends on the AWS region that you're using. <a href="#">Refer to AWS documentation for details</a>
<a href="https://support.netapp.com">https://support.netapp.com</a> <a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	To obtain licensing information and to send AutoSupport messages to NetApp support.
<a href="https://*.api.bluexp.netapp.com">https://*.api.bluexp.netapp.com</a> <a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a> <a href="https://*.cloudmanager.cloud.netapp.com">https://*.cloudmanager.cloud.netapp.com</a> <a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	To provide SaaS features and services within BlueXP.
Choose between two sets of endpoints: <ul style="list-style-type: none"><li>• Option 1 (recommended) <sup>1</sup>  <a href="https://bluexpinfraprod.eastus2.data.azurecr.io">https://bluexpinfraprod.eastus2.data.azurecr.io</a> <a href="https://bluexpinfraprod.azurecr.io">https://bluexpinfraprod.azurecr.io</a></li><li>• Option 2  <a href="https://*.blob.core.windows.net">https://*.blob.core.windows.net</a> <a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a></li></ul>	To obtain images for Connector upgrades.

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while

allowing the endpoints listed in option 2.

## Endpoints contacted from the BlueXP console

As you use the BlueXP web-based console that's provided through the SaaS layer, it contacts several endpoints to complete data management tasks. This includes endpoints that are contacted to deploy the Connector from the BlueXP console.

[View the list of endpoints contacted from the BlueXP console.](#)

## Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

## Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

## Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

You'll need to implement this networking requirement after you create the Connector.

## Step 2: Set up AWS permissions

BlueXP needs to authenticate with AWS before it can deploy the Connector instance in your VPC. You can choose one of these authentication methods:

- Let BlueXP assume an IAM role that has the required permissions
- Provide an AWS access key and secret key for an IAM user who has the required permissions

With either option, the first step is to create an IAM policy. This policy contains only the permissions needed to launch the Connector instance in AWS from BlueXP.

If needed, you can restrict the IAM policy by using the IAM `Condition` element. [AWS documentation: Condition element](#)

## Steps

1. Go to the AWS IAM console.
2. Select **Policies > Create policy**.
3. Select **JSON**.
4. Copy and paste the following policy:

This policy contains only the permissions needed to launch the Connector instance in AWS from BlueXP. When BlueXP creates the Connector, it applies a new set of permissions to the Connector instance that enables the Connector to manage AWS resources. [View permissions required for the Connector instance itself](#).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:CreateRole",
        "iam>DeleteRole",
        "iam:PutRolePolicy",
        "iam:CreateInstanceProfile",
        "iam>DeleteRolePolicy",
        "iam:AddRoleToInstanceProfile",
        "iam:RemoveRoleFromInstanceProfile",
        "iam>DeleteInstanceProfile",
        "iam:PassRole",
        "iam:ListRoles",
        "ec2:DescribeInstanceStatus",
        "ec2:RunInstances",
        "ec2:ModifyInstanceAttribute",
        "ec2:CreateSecurityGroup",
        "ec2>DeleteSecurityGroup",
        "ec2:DescribeSecurityGroups",
        "ec2:RevokeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:RevokeSecurityGroupIngress",
        "ec2:CreateNetworkInterface",
        "ec2:DescribeNetworkInterfaces",
        "ec2>DeleteNetworkInterface",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcs",
        "ec2:DescribeDhcpOptions",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeRegions",
```

```

    "ec2:DescribeInstances",
    "ec2:CreateTags",
    "ec2:DescribeImages",
    "ec2:DescribeAvailabilityZones",
    "ec2:DescribeLaunchTemplates",
    "ec2:CreateLaunchTemplate",
    "cloudformation:CreateStack",
    "cloudformation>DeleteStack",
    "cloudformation:DescribeStacks",
    "cloudformation:DescribeStackEvents",
    "cloudformation:ValidateTemplate",
    "ec2:AssociateIamInstanceProfile",
    "ec2:DescribeIamInstanceProfileAssociations",
    "ec2:DisassociateIamInstanceProfile",
    "iam:GetRole",
    "iam:TagRole",
    "kms:ListAliases",
    "cloudformation:ListStacks"
  ],
  "Resource": "*"
},
{
  "Effect": "Allow",
  "Action": [
    "ec2:TerminateInstances"
  ],
  "Condition": {
    "StringLike": {
      "ec2:ResourceTag/OCCMInstance": "*"
    }
  },
  "Resource": [
    "arn:aws:ec2:*:*:instance/*"
  ]
}
]
}

```

5. Select **Next** and add tags, if needed.
6. Select **Next** and enter a name and description.
7. Select **Create policy**.
8. Either attach the policy to an IAM role that BlueXP can assume or to an IAM user so that you can provide BlueXP with access keys:
  - (Option 1) Set up an IAM role that BlueXP can assume:
    - a. Go to the AWS IAM console in the target account.



- b. Under Access Management, select **Roles > Create Role** and follow the steps to create the role.
  - c. Under **Trusted entity type**, select **AWS account**.
  - d. Select **Another AWS account** and enter the ID of the BlueXP SaaS account: 952013314444
  - e. Select the policy that you created in the previous section.
  - f. After you create the role, copy the Role ARN so that you can paste it in BlueXP when you create the Connector.
- (Option 2) Set up permissions for an IAM user so that you can provide BlueXP with access keys:
    - a. From the AWS IAM console, select **Users** and then select the user name.
    - b. Select **Add permissions > Attach existing policies directly**.
    - c. Select the policy that you created.
    - d. Select **Next** and then select **Add permissions**.
    - e. Ensure that you have the access key and secret key for the IAM user.

## Result

You should now have an IAM role that has the required permissions or an IAM user that has the required permissions. When you create the Connector from BlueXP, you can provide information about the role or access keys.

## Step 3: Create the Connector

Create the Connector directly from the BlueXP web-based console.

### About this task

- Creating the Connector from BlueXP deploys an EC2 instance in AWS using a default configuration. After you create the Connector, you should not change to a smaller EC2 instance type that has less CPU or RAM. [Learn about the default configuration for the Connector](#).
- When BlueXP creates the Connector, it creates an IAM role and an instance profile for the instance. This role includes permissions that enables the Connector to manage AWS resources. You need to ensure that the role is kept up to date as new permissions are added in subsequent releases. [Learn more about the IAM policy for the Connector](#).

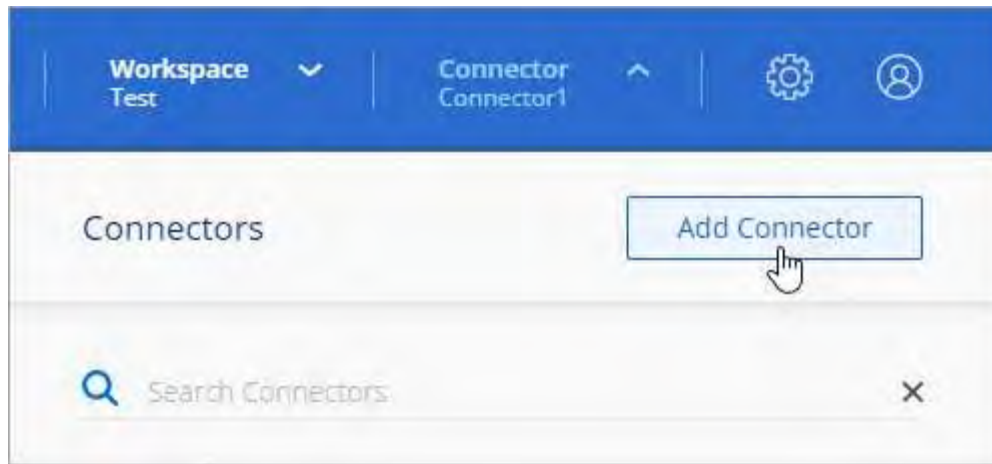
### Before you begin

You should have the following:

- An AWS authentication method: either an IAM role or access keys for an IAM user with the required permissions.
- A VPC and subnet that meets networking requirements.
- A key pair for the EC2 instance.
- Details about a proxy server, if a proxy is required for internet access from the Connector.

### Steps

1. Select the **Connector** drop-down and select **Add Connector**.



2. Choose **Amazon Web Services** as your cloud provider and select **Continue**.
3. On the **Deploying a Connector** page, review the details about what you'll need. You have two options:
  - a. Select **Continue** to prepare for deployment by using the in-product guide. Each step in the in-product guide includes the information that's contained on this page of the documentation.
  - b. Select **Skip to Deployment** if you already prepared by following the steps on this page.
4. Follow the steps in the wizard to create the Connector:
  - **Get Ready:** Review what you'll need.
  - **AWS Credentials:** Specify your AWS region and then choose an authentication method, which is either an IAM role that BlueXP can assume or an AWS access key and secret key.



If you choose **Assume Role**, you can create the first set of credentials from the Connector deployment wizard. Any additional set of credentials must be created from the Credentials page. They will then be available from the wizard in a drop-down list. [Learn how to add additional credentials.](#)

- **Details:** Provide details about the Connector.
  - Enter a name for the instance.
  - Add custom tags (metadata) to the instance.
  - Choose whether you want BlueXP to create a new role that has the required permissions, or if you want to select an existing role that you set up with [the required permissions](#).
  - Choose whether you want to encrypt the Connector's EBS disks. You have the option to use the default encryption key or to use a custom key.
- **Network:** Specify a VPC, subnet, and key pair for the instance, choose whether to enable a public IP address, and optionally specify a proxy configuration.

Make sure that you have the correct key pair to use with the Connector. Without a key pair, you will not be able to access the Connector virtual machine.

- **Security Group:** Choose whether to create a new security group or whether to select an existing security group that allows the required inbound and outbound rules.

[View security group rules for AWS.](#)

- **Review:** Review your selections to verify that your set up is correct.

## 5. Select **Add**.

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

### **Result**

After the process is complete, the Connector is available for use from BlueXP.

If you have Amazon S3 buckets in the same AWS account where you created the Connector, you'll see an Amazon S3 working environment appear on the BlueXP canvas automatically. [Learn how to manage S3 buckets from BlueXP](#)

### **Create a Connector from the AWS Marketplace**

You create a Connector in AWS directly from the AWS Marketplace. To create a Connector from the AWS Marketplace, you need to set up your networking, prepare AWS permissions, review instance requirements, and then create the Connector.

### **Before you begin**

- You should have an [understanding of Connectors](#).
- You should review [Connector limitations](#).

### **Step 1: Set up networking**

Ensure that the network location where you plan to install the Connector supports the following requirements. Meeting these requirements enables the Connector to manage resources and processes within your hybrid cloud environment.

#### **VPC and subnet**

When you create the Connector, you need to specify the VPC and subnet where the Connector should reside.

#### **Connections to target networks**

A Connector requires a network connection to the location where you're planning to create and manage working environments. For example, the network where you plan to create Cloud Volumes ONTAP systems or a storage system in your on-premises environment.

#### **Outbound internet access**

The network location where you deploy the Connector must have an outbound internet connection to contact specific endpoints.

#### **Endpoints contacted from the Connector**

The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment for day-to-day operations.

Note that the endpoints listed below are all CNAME entries.

Endpoints	Purpose
AWS services (amazonaws.com): <ul style="list-style-type: none"> <li>• CloudFormation</li> <li>• Elastic Compute Cloud (EC2)</li> <li>• Identity and Access Management (IAM)</li> <li>• Key Management Service (KMS)</li> <li>• Security Token Service (STS)</li> <li>• Simple Storage Service (S3)</li> </ul>	To manage resources in AWS. The exact endpoint depends on the AWS region that you're using. <a href="#">Refer to AWS documentation for details</a>
<a href="https://support.netapp.com">https://support.netapp.com</a> <a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	To obtain licensing information and to send AutoSupport messages to NetApp support.
<a href="https://*.api.bluexp.netapp.com">https://*.api.bluexp.netapp.com</a> <a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a> <a href="https://*.cloudmanager.cloud.netapp.com">https://*.cloudmanager.cloud.netapp.com</a> <a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	To provide SaaS features and services within BlueXP.
Choose between two sets of endpoints: <ul style="list-style-type: none"> <li>• Option 1 (recommended) <sup>1</sup> <p> <a href="https://bluexpinfraprod.eastus2.data.azurecr.io">https://bluexpinfraprod.eastus2.data.azurecr.io</a>  <a href="https://bluexpinfraprod.azurecr.io">https://bluexpinfraprod.azurecr.io</a> </p> </li> <li>• Option 2               <p> <a href="https://*.blob.core.windows.net">https://*.blob.core.windows.net</a>  <a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a> </p> </li> </ul>	To obtain images for Connector upgrades.

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

## Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

## Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

## Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

You'll need to implement this networking requirement after you create the Connector.

## Step 2: Set up AWS permissions

To prepare for a marketplace deployment, create IAM policies in AWS and attach them to an IAM role. When you create the Connector from the AWS Marketplace, you'll be prompted to select that IAM role.

### Steps

1. Log in to the AWS console and navigate to the IAM service.
2. Create a policy:
  - a. Select **Policies > Create policy**.
  - b. Select **JSON** and copy and paste the contents of the [IAM policy for the Connector](#).
  - c. Finish the remaining steps to create the policy.

Depending on the BlueXP services that you're planning to use, you might need to create a second policy. For standard regions, the permissions are spread across two policies. Two policies are required due to a maximum character size limit for managed policies in AWS. [Learn more about IAM policies for the Connector](#).

3. Create an IAM role:
  - a. Select **Roles > Create role**.
  - b. Select **AWS service > EC2**.
  - c. Add permissions by attaching the policy that you just created.
  - d. Finish the remaining steps to create the role.

## Result

You now have an IAM role that you can associate with the EC2 instance during deployment from the AWS Marketplace.

## Step 3: Review instance requirements

When you create the Connector, you need to choose an EC2 instance type that meets the following requirements.

### CPU

8 cores or 8 vCPUs

### RAM

32 GB

### AWS EC2 instance type

An instance type that meets the CPU and RAM requirements above. We recommend t3.2xlarge.

## Step 4: Create the Connector

Create the Connector directly from the AWS Marketplace.

### About this task

Creating the Connector from the AWS Marketplace deploys an EC2 instance in AWS using a default configuration. [Learn about the default configuration for the Connector.](#)

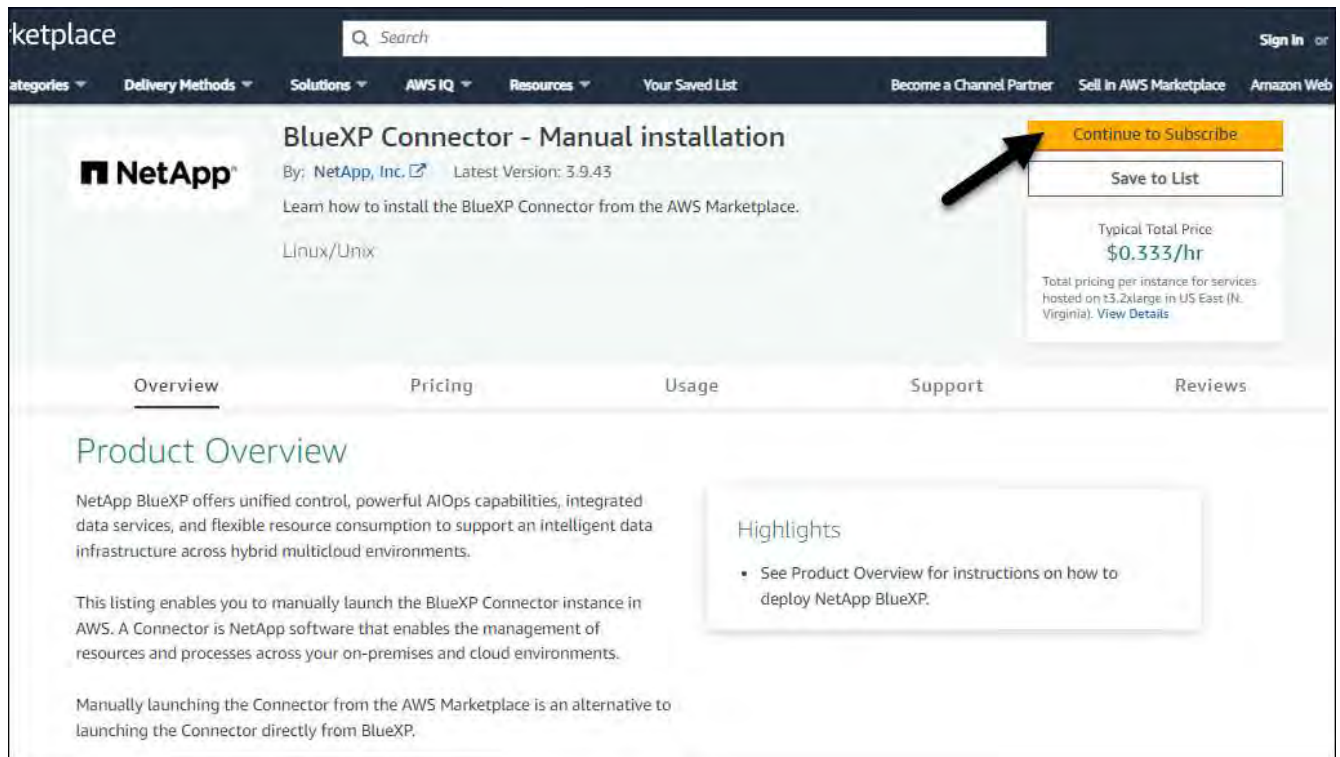
### Before you begin

You should have the following:

- A VPC and subnet that meets networking requirements.
- An IAM role with an attached policy that includes the required permissions for the Connector.
- Permissions to subscribe and unsubscribe from the AWS Marketplace for your IAM user.
- An understanding of CPU and RAM requirements for the instance.
- A key pair for the EC2 instance.

### Steps

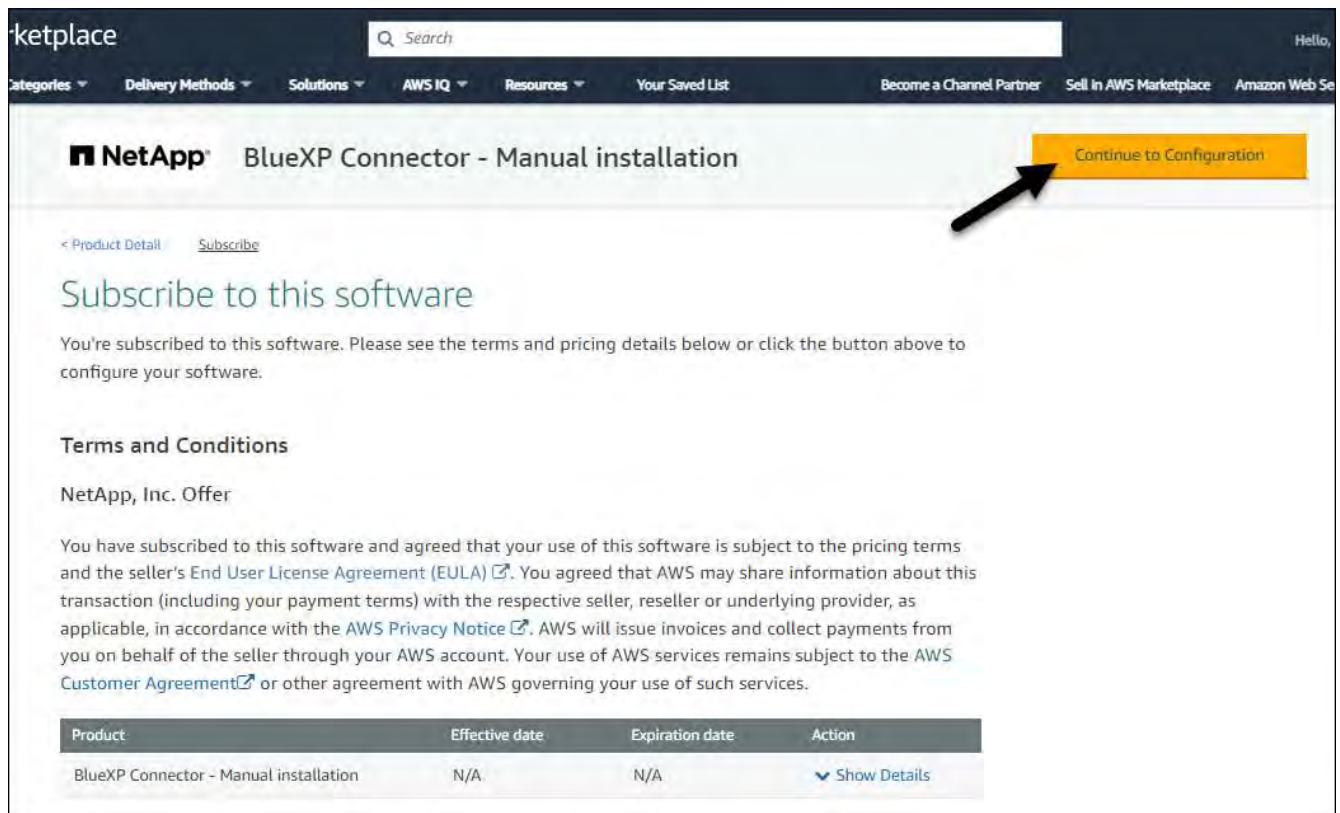
1. Go to the [BlueXP Connector listing on the AWS Marketplace](#)
2. On the Marketplace page, select **Continue to Subscribe**.



3. To subscribe to the software, select **Accept Terms**.

The subscription process can take a few minutes.

4. After the subscription process is complete, select **Continue to Configuration**.



5. On the **Configure this software** page, ensure that you've selected the correct region and then select



## Continue to Launch.

6. On the **Launch this software** page, under **Choose Action**, select **Launch through EC2** and then select **Launch**.

These steps describe how to launch the instance from the EC2 Console because the console enables you to attach an IAM role to the Connector instance. This isn't possible using the **Launch from Website** action.

7. Follow the prompts to configure and deploy the instance:

- **Name and tags:** Enter a name and tags for the instance.
- **Application and OS Images:** Skip this section. The Connector AMI is already selected.
- **Instance type:** Depending on region availability, choose an instance type that meets RAM and CPU requirements (t3.2xlarge is preselected and recommended).
- **Key pair (login):** Select the key pair that you want to use to securely connect to the instance.
- **Network settings:** Edit the network settings as needed:
  - Choose the desired VPC and subnet.
  - Specify whether the instance should have a public IP address.
  - Specify security group settings that enable the required connection methods for the Connector instance: SSH, HTTP, and HTTPS.

[View security group rules for AWS.](#)

- **Configure storage:** Keep the default size and disk type for the root volume.

If you want to enable Amazon EBS encryption on the root volume, select **Advanced**, expand **Volume 1**, select **Encrypted**, and then choose a KMS key.

- **Advanced details:** Under **IAM instance profile**, choose the IAM role that includes the required permissions for the Connector.
- **Summary:** Review the summary and select **Launch instance**.

AWS launches the software with the specified settings. The Connector instance and software should be running in approximately five minutes.

8. Open a web browser from a host that has a connection to the Connector virtual machine and enter the following URL:

`https://ipaddress`

9. After you log in, set up the Connector:

- a. Specify the BlueXP organization to associate with the Connector.
- b. Enter a name for the system.
- c. Under **Are you running in a secured environment?** keep restricted mode disabled.

You should keep restricted mode disabled because these steps describe how to use BlueXP in standard mode. You should enable restricted mode only if you have a secure environment and want to disconnect this account from BlueXP backend services. If that's the case, [follow steps to get started with BlueXP in restricted mode](#).

- d. Select **Let's start**.



## Result

The Connector is now installed and set up with your BlueXP organization.

Open a web browser and go to the [BlueXP console](#) to start using the Connector with BlueXP.

If you have Amazon S3 buckets in the same AWS account where you created the Connector, you'll see an Amazon S3 working environment appear on the BlueXP canvas automatically. [Learn how to manage S3 buckets from BlueXP](#)

## Manually install the Connector in AWS

You can manually install a Connector on a Linux host running in AWS. To manually install the Connector on your own Linux host, you need to review host requirements, set up your networking, prepare AWS permissions, install the Connector, and then provide the permissions that you prepared.

### Before you begin

- You should have an [understanding of Connectors](#).
- You should review [Connector limitations](#).

### Step 1: Review host requirements

The Connector software must run on a host that meets specific operating system requirements, RAM requirements, port requirements, and so on.



The Connector reserves the UID and GID range of 19000 to 19200. This range is fixed and cannot be modified. If any third-party software on your host is using UIDs or GIDs within this range, the Connector installation will fail. NetApp recommends using a host that is free of third-party software to avoid conflicts.

### Dedicated host

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

Host can be of any architecture that meets the following size requirements:

- CPU: 8 cores or 8 vCPUs
- RAM: 32 GB

### Hypervisor

A bare metal or hosted hypervisor that is certified to run a supported operating system is required.

### Operating system and container requirements

BlueXP supports the Connector with the following operating systems when using BlueXP in standard mode or restricted mode. A container orchestration tool is required before you install the Connector.

Operating system	Supported OS versions	Supported Connector versions	Required container tool	SELinux
Red Hat Enterprise Linux	9.1 to 9.4 8.6 to 8.10 7.9	3.9.40 or later with BlueXP in standard mode or restricted mode	Podman version 4.6.1 or 4.9.4  <a href="#">View Podman configuration requirements.</a>	Supported in enforcing mode or permissive mode <sup>1</sup>
Ubuntu	24.04 LTS	3.9.45 or later with BlueXP in standard mode or restricted mode	Docker Engine 23.06 to 28.0.0.	Not supported
	22.04 LTS	3.9.29 or later	Docker Engine 23.0.6 to 28.0.0.	Not supported

Notes:

1. Management of Cloud Volumes ONTAP systems is not supported by Connectors that have SELinux enabled on the operating system.
2. The Connector is supported on English-language versions of these operating systems.
3. For RHEL, the host must be registered with Red Hat Subscription Management. If it's not registered, the host can't access repositories to update required 3rd-party software during Connector installation.

### AWS EC2 instance type

An instance type that meets the CPU and RAM requirements above. We recommend t3.2xlarge.

### Key pair

When you create the Connector, you'll need to select an EC2 key pair to use with the instance.

### PUT response hop limit when using IMDSv2

If IMDSv2 is enabled on the EC2 instance (this is the default setting for new EC2 instances), you must change the PUT response hop limit on the instance to 3. If you don't change the limit on the EC2 instance, you'll receive a UI initialization error when you try to set up the Connector.

- [Require the use of IMDSv2 on Amazon EC2 instances](#)
- [AWS documentation: Change the PUT response hop limit](#)

### Disk space in /opt

100 GiB of space must be available

BlueXP uses /opt to install the /opt/application/netapp directory and its contents.

### Disk space in /var

20 GiB of space must be available

BlueXP requires this space in /var because Docker or Podman are architected to create the containers

within this directory. Specifically, they will create containers in the `/var/lib/containers/storage` directory. External mounts or symlinks do not work for this space.

## **Step 2: Install Podman or Docker Engine**

Depending on your operating system, either Podman or Docker Engine is required before you install the Connector.

- Podman is required for Red Hat Enterprise Linux 8 and 9.

[View the Podman versions that BlueXP supports.](#)

- Docker Engine is required for Ubuntu.

[View the Docker Engine versions that BlueXP supports.](#)

## Example 1. Steps

### Podman

Follow these steps to install Podman and configure it to meet the following requirements:

- Enable and start the podman.socket service
- Install python3
- Install the podman-compose package version 1.0.6
- Add podman-compose to the PATH environment variable



When using Podman, adjust the aardvark-dns service port (default: 53) after installing the Connector to avoid conflicts with the DNS port on the host. Follow the instructions to configure the port.

### Steps

1. Remove the podman-docker package if it's installed on the host.

```
dnf remove podman-docker
rm /var/run/docker.sock
```

2. Install Podman.

Podman is available from official Red Hat Enterprise Linux repositories.

For Red Hat Enterprise Linux 9:

```
sudo dnf install podman-2:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

For Red Hat Enterprise Linux 8:

```
sudo dnf install podman-3:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

3. Enable and start the podman.socket service.

```
sudo systemctl enable --now podman.socket
```

4. Install python3.

```
sudo dnf install python3
```

5. Install the EPEL repository package if it's not already available on your system.

This step is required because podman-compose is available from the Extra Packages for Enterprise Linux (EPEL) repository.

For Red Hat Enterprise Linux 9:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-9.noarch.rpm
```

For Red Hat Enterprise Linux 8:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

6. Install podman-compose package 1.0.6.

```
sudo dnf install podman-compose-1.0.6
```



Using the `dnf install` command meets the requirement for adding podman-compose to the PATH environment variable. The installation command adds podman-compose to `/usr/bin`, which is already included in the `secure_path` option on the host.

### Docker Engine

Follow the documentation from Docker to install Docker Engine.

#### Steps

1. [View installation instructions from Docker](#)

Be sure to follow the steps to install a specific version of Docker Engine. Installing the latest version will install a version of Docker that BlueXP doesn't support.

2. Verify that Docker is enabled and running.

```
sudo systemctl enable docker && sudo systemctl start docker
```

### Step 3: Set up networking

Ensure that the network location where you plan to install the Connector supports the following requirements. Meeting these requirements enables the Connector to manage resources and processes within your hybrid

cloud environment.

## Connections to target networks

A Connector requires a network connection to the location where you're planning to create and manage working environments. For example, the network where you plan to create Cloud Volumes ONTAP systems or a storage system in your on-premises environment.

## Outbound internet access

The network location where you deploy the Connector must have an outbound internet connection to contact specific endpoints.

## Endpoints contacted from computers when using the BlueXP web-based console

Computers that access the BlueXP console from a web browser must have the ability to contact several endpoints. You'll need to use the BlueXP console to set up the Connector and for day-to-day use of BlueXP.

[Prepare networking for the BlueXP console.](#)

## Endpoints contacted during manual installation

When you manually install the Connector on your own Linux host, the installer for the Connector requires access to the following URLs during the installation process:

- <https://mysupport.netapp.com>
- <https://signin.b2c.netapp.com> (this endpoint is the CNAME URL for <https://mysupport.netapp.com>)
- <https://cloudmanager.cloud.netapp.com/tenancy>
- <https://stream.cloudmanager.cloud.netapp.com>
- <https://production-artifacts.cloudmanager.cloud.netapp.com>
- To obtain images, the installer needs access to one of these two sets of endpoints:
  - Option 1 (recommended):
    - <https://bluexpinfraprod.eastus2.data.azurecr.io>
    - <https://bluexpinfraprod.azurecr.io>
  - Option 2:
    - [https://\\*.blob.core.windows.net](https://*.blob.core.windows.net)
    - <https://cloudmanagerinfraprod.azurecr.io>

The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

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

## Endpoints contacted from the Connector

The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment for day-to-day operations.

Note that the endpoints listed below are all CNAME entries.

Endpoints	Purpose
AWS services (amazonaws.com): <ul style="list-style-type: none"><li>• CloudFormation</li><li>• Elastic Compute Cloud (EC2)</li><li>• Identity and Access Management (IAM)</li><li>• Key Management Service (KMS)</li><li>• Security Token Service (STS)</li><li>• Simple Storage Service (S3)</li></ul>	To manage resources in AWS. The exact endpoint depends on the AWS region that you're using. <a href="#">Refer to AWS documentation for details</a>
<a href="https://support.netapp.com">https://support.netapp.com</a> <a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	To obtain licensing information and to send AutoSupport messages to NetApp support.
<a href="https://*.api.bluexp.netapp.com">https://*.api.bluexp.netapp.com</a> <a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a> <a href="https://*.cloudmanager.cloud.netapp.com">https://*.cloudmanager.cloud.netapp.com</a> <a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	To provide SaaS features and services within BlueXP.
Choose between two sets of endpoints: <ul style="list-style-type: none"><li>• Option 1 (recommended) <sup>1</sup>  <a href="https://bluexpinfraprod.eastus2.data.azurecr.io">https://bluexpinfraprod.eastus2.data.azurecr.io</a> <a href="https://bluexpinfraprod.azurecr.io">https://bluexpinfraprod.azurecr.io</a></li><li>• Option 2  <a href="https://*.blob.core.windows.net">https://*.blob.core.windows.net</a> <a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a></li></ul>	To obtain images for Connector upgrades.

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while

allowing the endpoints listed in option 2.

### Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

### Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

### Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

### Step 4: Set up permissions

You need to provide AWS permissions to BlueXP by using one of the following options:

- Option 1: Create IAM policies and attach the policies to an IAM role that you can associate with the EC2 instance.
- Option 2: Provide BlueXP with the AWS access key for an IAM user who has the required permissions.

Follow the steps to prepare permissions for BlueXP.



## IAM role

### Steps

1. Log in to the AWS console and navigate to the IAM service.
2. Create a policy:
  - a. Select **Policies > Create policy**.
  - b. Select **JSON** and copy and paste the contents of the [IAM policy for the Connector](#).
  - c. Finish the remaining steps to create the policy.

Depending on the BlueXP services that you're planning to use, you might need to create a second policy. For standard regions, the permissions are spread across two policies. Two policies are required due to a maximum character size limit for managed policies in AWS. [Learn more about IAM policies for the Connector](#).

3. Create an IAM role:
  - a. Select **Roles > Create role**.
  - b. Select **AWS service > EC2**.
  - c. Add permissions by attaching the policy that you just created.
  - d. Finish the remaining steps to create the role.

### Result

You now have an IAM role that you can associate with the EC2 instance after you install the Connector.

## AWS access key

### Steps

1. Log in to the AWS console and navigate to the IAM service.
2. Create a policy:
  - a. Select **Policies > Create policy**.
  - b. Select **JSON** and copy and paste the contents of the [IAM policy for the Connector](#).
  - c. Finish the remaining steps to create the policy.

Depending on the BlueXP services that you're planning to use, you might need to create a second policy.

For standard regions, the permissions are spread across two policies. Two policies are required due to a maximum character size limit for managed policies in AWS. [Learn more about IAM policies for the Connector](#).

3. Attach the policies to an IAM user.
  - [AWS Documentation: Creating IAM Roles](#)
  - [AWS Documentation: Adding and Removing IAM Policies](#)
4. Ensure that the user has an access key that you can add to BlueXP after you install the Connector.

### Result

You now have an IAM user that has the required permissions and an access key that you can provide to BlueXP.

## Step 5: Install the Connector

After the pre-requisites are complete, you can manually install the software on your own Linux host.

### Before you begin

You should have the following:

- Root privileges to install the Connector.
- Details about a proxy server, if a proxy is required for internet access from the Connector.

You have the option to configure a proxy server after installation but doing so requires restarting the Connector.

- A CA-signed certificate, if the proxy server uses HTTPS or if the proxy is an intercepting proxy.



You cannot set a certificate for a transparent proxy server when manually installing the Connector. If you need to set a certificate for a transparent proxy server, you must use the Maintenance Console after installation. Learn more about the [Connector Maintenance Console](#).

### About this task

The installer that is available on the NetApp Support Site might be an earlier version. After installation, the Connector automatically updates itself if a new version is available.

### Steps

1. If the *http\_proxy* or *https\_proxy* system variables are set on the host, remove them:

```
unset http_proxy
unset https_proxy
```

If you don't remove these system variables, the installation will fail.

2. Download the Connector software from the [NetApp Support Site](#), and then copy it to the Linux host.

You should download the "online" Connector installer that's meant for use in your network or in the cloud. A separate "offline" installer is available for the Connector, but it's only supported with private mode deployments.

3. Assign permissions to run the script.

```
chmod +x BlueXP-Connector-Cloud-<version>
```

Where <version> is the version of the Connector that you downloaded.

4. Run the installation script.

```
./BlueXP-Connector-Cloud-<version> --proxy <HTTP or HTTPS proxy server>
--cacert <path and file name of a CA-signed certificate>
```

You'll need to add proxy information if your network requires a proxy for internet access. You can add either a transparent or explicit proxy. The `--proxy` and `--cacert` parameters are optional and you won't be prompted to add them. If you have a proxy server, you will need to enter the parameters as shown.

Here is an example configuring an explicit proxy server with a CA-signed certificate:

```
./BlueXP-Connector-Cloud-v3.9.40--proxy
https://user:password@10.0.0.30:8080/ --cacert
/tmp/cacert/certificate.cer
```

`--proxy` configures the Connector to use an HTTP or HTTPS proxy server using one of the following formats:

- `http://address:port`
- `http://user-name:password@address:port`
- `http://domain-name%92user-name:password@address:port`
- `https://address:port`
- `https://user-name:password@address:port`
- `https://domain-name%92user-name:password@address:port`

Note the following:

- The user can be a local user or domain user.
- For a domain user, you must use the ASCII code for a `\` as shown above.
- BlueXP doesn't support user names or passwords that include the `@` character.
- If the password includes any of the following special characters, you must escape that special character by prepending it with a backslash: `&` or `!`

For example:

```
http://bxpproxyuser:netapp1!@address:3128
```

`--cacert` specifies a CA-signed certificate to use for HTTPS access between the Connector and the proxy server. This parameter is required for HTTPS proxy servers, intercepting proxy servers, and transparent proxy servers.

Here is an example configuring a transparent proxy server. When you configure a transparent proxy, you don't need to define the proxy server. You only add a CA-signed certificate to your Connector host:

```
./BlueXP-Connector-Cloud-v3.9.40 --cacert /tmp/cacert/certificate.cer
```

5. If you used Podman, you'll need to adjust the aardvark-dns port.
  - a. SSH to the BlueXP Connector virtual machine.
  - b. Open podman `/usr/share/containers/containers.conf` file and modify the chosen port for Aardvark DNS service. For example, change it to 54.

```
vi /usr/share/containers/containers.conf
...
# Port to use for dns forwarding daemon with netavark in rootful
bridge
# mode and dns enabled.
# Using an alternate port might be useful if other DNS services
should
# run on the machine.
#
dns_bind_port = 54
...
Esc:wq
```

- c. Reboot the Connector virtual machine.
6. Wait for the installation to complete.

At the end of the installation, the Connector service (occm) restarts twice if you specified a proxy server.

7. Open a web browser from a host that has a connection to the Connector virtual machine and enter the following URL:

`https://ipaddress`

8. After you log in, set up the Connector:
  - a. Specify the BlueXP organization to associate with the Connector.
  - b. Enter a name for the system.
  - c. Under **Are you running in a secured environment?** keep restricted mode disabled.

You should keep restricted mode disabled because these steps describe how to use BlueXP in standard mode. You should enable restricted mode only if you have a secure environment and want to disconnect this account from BlueXP backend services. If that's the case, [follow steps to get started with BlueXP in restricted mode](#).

- d. Select **Let's start**.

If you have Amazon S3 buckets in the same AWS account where you created the Connector, you'll see an Amazon S3 working environment appear on the BlueXP canvas automatically. [Learn how to manage S3 buckets from BlueXP](#)

## Step 6: Provide permissions to BlueXP

Now that you've installed the Connector, you need to provide BlueXP with the AWS permissions that you previously set up. Providing the permissions enables BlueXP to manage your data and storage infrastructure in AWS.

## IAM role

Attach the IAM role that you previously created to the Connector EC2 instance.

### Steps

1. Go to the Amazon EC2 console.
2. Select **Instances**.
3. Select the Connector instance.
4. Select **Actions > Security > Modify IAM role**.
5. Select the IAM role and select **Update IAM role**.

### Result

BlueXP now has the permissions that it needs to perform actions in AWS on your behalf.

Go to the [BlueXP console](#) to start using the Connector with BlueXP.

## AWS access key

Provide BlueXP with the AWS access key for an IAM user that has the required permissions.

### Steps

1. Ensure that the correct Connector is currently selected in BlueXP.
2. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



3. Select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location:** Select **Amazon Web Services > Connector**.
  - b. **Define Credentials:** Enter an AWS access key and secret key.
  - c. **Marketplace Subscription:** Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
  - d. **Review:** Confirm the details about the new credentials and select **Add**.

### Result

BlueXP now has the permissions that it needs to perform actions in AWS on your behalf.

Go to the [BlueXP console](#) to start using the Connector with BlueXP.

## Azure

### Connector installation options in Azure

There are a few different ways to create a Connector in Azure. Directly from BlueXP is the most common way.

The following installation options are available:

- [Create a Connector directly from BlueXP](#) (this is the standard option)

This action launches a VM running Linux and the Connector software in a VNet of your choice.

- [Create a Connector from the Azure Marketplace](#)

This action also launches a VM running Linux and the Connector software, but the deployment is initiated directly from the Azure Marketplace, rather than from BlueXP.

- [Download and manually install the software on your own Linux host](#)

The installation option that you choose impacts how you prepare for installation. This includes how you provide BlueXP with the required permissions that it needs to authenticate and manage resources in Azure.

### Create a Connector in Azure from BlueXP

You can install a Connector in Azure directly from BlueXP. To create a Connector in Azure from BlueXP, you need to set up your networking, prepare an Azure role to use to deploy the Connector, and then deploy the Connector.

#### Before you begin

- You should have an [understanding of Connectors](#).
- You should review [Connector limitations](#).

#### Step 1: Set up networking

Ensure that the network location where you plan to install the Connector supports the following requirements. Meeting these requirements enables the Connector to manage resources and processes within your hybrid cloud environment.

#### Azure region

If you use Cloud Volumes ONTAP, the Connector should be deployed in the same Azure region as the Cloud Volumes ONTAP systems that it manages, or in the [Azure region pair](#) for the Cloud Volumes ONTAP systems. This requirement ensures that an Azure Private Link connection is used between Cloud Volumes ONTAP and its associated storage accounts.

[Learn how Cloud Volumes ONTAP uses an Azure Private Link](#)

#### VNet and subnet

When you create the Connector, you need to specify the VNet and subnet where the Connector should reside.

#### Connections to target networks

A Connector requires a network connection to the location where you're planning to create and manage working environments. For example, the network where you plan to create Cloud Volumes ONTAP systems or a storage system in your on-premises environment.

#### Outbound internet access

The network location where you deploy the Connector must have an outbound internet connection to contact specific endpoints.

## Endpoints contacted from the Connector

The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment for day-to-day operations.

Note that the endpoints listed below are all CNAME entries.

Endpoints	Purpose
<a href="https://management.azure.com">https://management.azure.com</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a> <a href="https://blob.core.windows.net">https://blob.core.windows.net</a> <a href="https://core.windows.net">https://core.windows.net</a>	To manage resources in Azure public regions.
<a href="https://management.chinacloudapi.cn">https://management.chinacloudapi.cn</a> <a href="https://login.chinacloudapi.cn">https://login.chinacloudapi.cn</a> <a href="https://blob.core.chinacloudapi.cn">https://blob.core.chinacloudapi.cn</a> <a href="https://core.chinacloudapi.cn">https://core.chinacloudapi.cn</a>	To manage resources in Azure China regions.
<a href="https://support.netapp.com">https://support.netapp.com</a> <a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	To obtain licensing information and to send AutoSupport messages to NetApp support.
<a href="https://*.api.bluexp.netapp.com">https://*.api.bluexp.netapp.com</a> <a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a> <a href="https://*.cloudmanager.cloud.netapp.com">https://*.cloudmanager.cloud.netapp.com</a> <a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	To provide SaaS features and services within BlueXP.
Choose between two sets of endpoints: <ul style="list-style-type: none"><li>Option 1 (recommended) <sup>1</sup> <a href="https://bluexpinfraprod.eastus2.data.azurecr.io">https://bluexpinfraprod.eastus2.data.azurecr.io</a> <a href="https://bluexpinfraprod.azurecr.io">https://bluexpinfraprod.azurecr.io</a></li><li>Option 2 <a href="https://*.blob.core.windows.net">https://*.blob.core.windows.net</a> <a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a></li></ul>	To obtain images for Connector upgrades.

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

## Endpoints contacted from the BlueXP console

As you use the BlueXP web-based console that's provided through the SaaS layer, it contacts several endpoints to complete data management tasks. This includes endpoints that are contacted to deploy the Connector from the BlueXP console.

[View the list of endpoints contacted from the BlueXP console.](#)

## Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

## Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

## Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

You'll need to implement this networking requirement after you create the Connector.

## Step 2: Create a Connector deployment policy (custom role)

You need to create a custom role that has permissions to deploy the Connector in Azure.

Create an Azure custom role that you can assign to your Azure account or to a Microsoft Entra service principal. BlueXP authenticates with Azure and uses these permissions to create the Connector instance on your behalf.

After BlueXP deploys the Connector virtual machine in Azure, it enables a [system-assigned managed identity](#) on the virtual machine, automatically creates the role it needs, and assigns it to the virtual machine. The automatically created role provides BlueXP with the permissions required to manage resources and processes within that Azure subscription. [Review how BlueXP uses the permissions.](#)

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would prefer to use a different



method, refer to [Azure documentation](#)

## Steps

1. Copy the required permissions for a new custom role in Azure and save them in a JSON file.



This custom role contains only the permissions needed to launch the Connector VM in Azure from BlueXP. Don't use this policy for other situations. When BlueXP creates the Connector, it applies a new set of permissions to the Connector VM that enables the Connector to manage Azure resources.

```
{
  "Name": "Azure SetupAsService",
  "Actions": [
    "Microsoft.Compute/disks/delete",
    "Microsoft.Compute/disks/read",
    "Microsoft.Compute/disks/write",
    "Microsoft.Compute/locations/operations/read",
    "Microsoft.Compute/operations/read",
    "Microsoft.Compute/virtualMachines/instanceView/read",
    "Microsoft.Compute/virtualMachines/read",
    "Microsoft.Compute/virtualMachines/write",
    "Microsoft.Compute/virtualMachines/delete",
    "Microsoft.Compute/virtualMachines/extensions/write",
    "Microsoft.Compute/virtualMachines/extensions/read",
    "Microsoft.Compute/availabilitySets/read",
    "Microsoft.Network/locations/operationResults/read",
    "Microsoft.Network/locations/operations/read",
    "Microsoft.Network/networkInterfaces/join/action",
    "Microsoft.Network/networkInterfaces/read",
    "Microsoft.Network/networkInterfaces/write",
    "Microsoft.Network/networkInterfaces/delete",
    "Microsoft.Network/networkSecurityGroups/join/action",
    "Microsoft.Network/networkSecurityGroups/read",
    "Microsoft.Network/networkSecurityGroups/write",

    "Microsoft.Network/virtualNetworks/checkIpAddressAvailability/read",
    "Microsoft.Network/virtualNetworks/read",
    "Microsoft.Network/virtualNetworks/subnets/join/action",
    "Microsoft.Network/virtualNetworks/subnets/read",

    "Microsoft.Network/virtualNetworks/subnets/virtualMachines/read",
    "Microsoft.Network/virtualNetworks/virtualMachines/read",
    "Microsoft.Network/publicIPAddresses/write",
    "Microsoft.Network/publicIPAddresses/read",
    "Microsoft.Network/publicIPAddresses/delete",
    "Microsoft.Network/networkSecurityGroups/securityRules/read",
```

```

    "Microsoft.Network/networkSecurityGroups/securityRules/write",
    "Microsoft.Network/networkSecurityGroups/securityRules/delete",
    "Microsoft.Network/publicIPAddresses/join/action",

"Microsoft.Network/locations/virtualNetworkAvailableEndpointServices/read",
    "Microsoft.Network/networkInterfaces/ipConfigurations/read",
    "Microsoft.Resources/deployments/operations/read",
    "Microsoft.Resources/deployments/read",
    "Microsoft.Resources/deployments/delete",
    "Microsoft.Resources/deployments/cancel/action",
    "Microsoft.Resources/deployments/validate/action",
    "Microsoft.Resources/resources/read",
    "Microsoft.Resources/subscriptions/operationresults/read",
    "Microsoft.Resources/subscriptions/resourceGroups/delete",
    "Microsoft.Resources/subscriptions/resourceGroups/read",

"Microsoft.Resources/subscriptions/resourcegroups/resources/read",
    "Microsoft.Resources/subscriptions/resourceGroups/write",
    "Microsoft.Authorization/roleDefinitions/write",
    "Microsoft.Authorization/roleAssignments/write",

"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/read",

"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/write",
    "Microsoft.Network/networkSecurityGroups/delete",
    "Microsoft.Storage/storageAccounts/delete",
    "Microsoft.Storage/storageAccounts/write",
    "Microsoft.Resources/deployments/write",
    "Microsoft.Resources/deployments/operationStatuses/read",
    "Microsoft.Authorization/roleAssignments/read"
],
  "NotActions": [],
  "AssignableScopes": [],
  "Description": "Azure SetupAsService",
  "IsCustom": "true"
}

```

2. Modify the JSON by adding your Azure subscription ID to the assignable scope.


### Example

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

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- a. Start [Azure Cloud Shell](#) and choose the Bash environment.
- b. Upload the JSON file.



c. Enter the following Azure CLI command:

```
az role definition create --role-definition
Policy_for_Setup_As_Service_Azure.json
```

You should now have a custom role called *Azure SetupAsService*. You can now apply this custom role to your user account or to a service principal.

### Step 3: Set up authentication

When creating the Connector from BlueXP, you need to provide a login that enables BlueXP to authenticate with Azure and deploy the VM. You have two options:

1. Sign in with your Azure account when prompted. This account must have specific Azure permissions. This is the default option.
2. Provide details about a Microsoft Entra service principal. This service principal also requires specific permissions.

Follow the steps to prepare one of these authentication methods for use with BlueXP.

## Azure account

Assign the custom role to the user who will deploy the Connector from BlueXP.

### Steps

1. In the Azure portal, open the **Subscriptions** service and select the user's subscription.
2. Click **Access control (IAM)**.
3. Click **Add > Add role assignment** and then add the permissions:
  - a. Select the **Azure SetupAsService** role and click **Next**.



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

- b. Keep **User, group, or service principal** selected.
- c. Click **Select members**, choose your user account, and click **Select**.
- d. Click **Next**.
- e. Click **Review + assign**.

### Result

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

### Service principal

Rather than logging in with your Azure account, you can provide BlueXP with the credentials for an Azure service principal that has the required permissions.

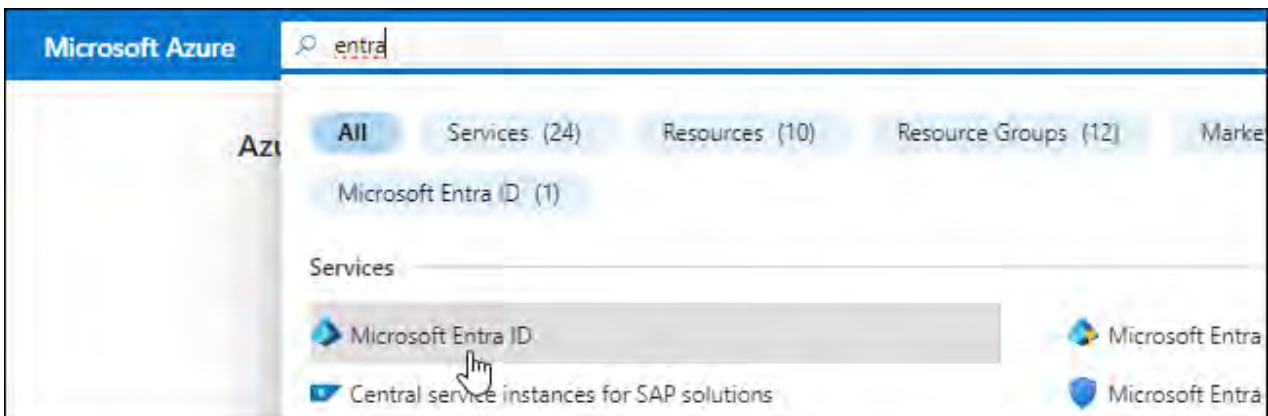
Create and set up a service principal in Microsoft Entra ID and obtain the Azure credentials that BlueXP needs.

### Create a Microsoft Entra application for role-based access control

1. Ensure that you have permissions in Azure to create an Active Directory application and to assign the application to a role.

For details, refer to [Microsoft Azure Documentation: Required permissions](#)

2. From the Azure portal, open the **Microsoft Entra ID** service.

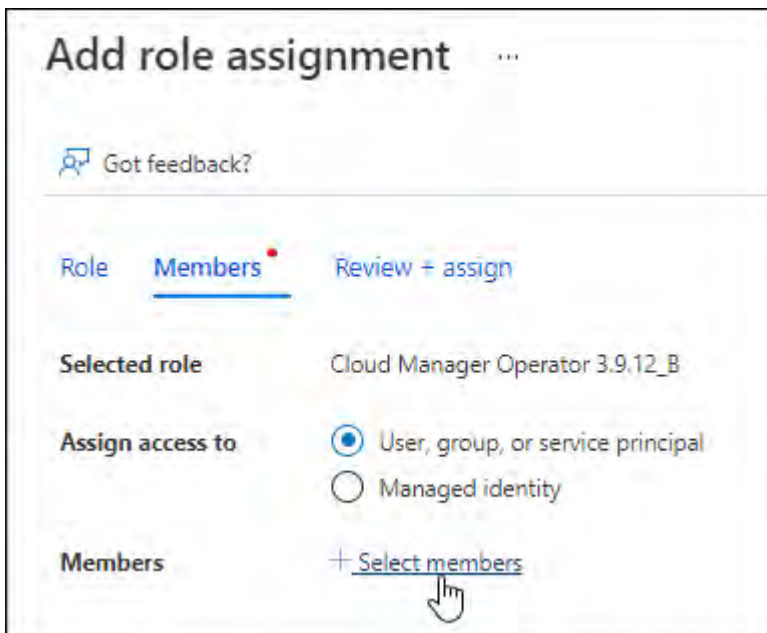


3. In the menu, select **App registrations**.
4. Select **New registration**.
5. Specify details about the application:
  - **Name**: Enter a name for the application.
  - **Account type**: Select an account type (any will work with BlueXP).
  - **Redirect URI**: You can leave this field blank.
6. Select **Register**.

You've created the AD application and service principal.

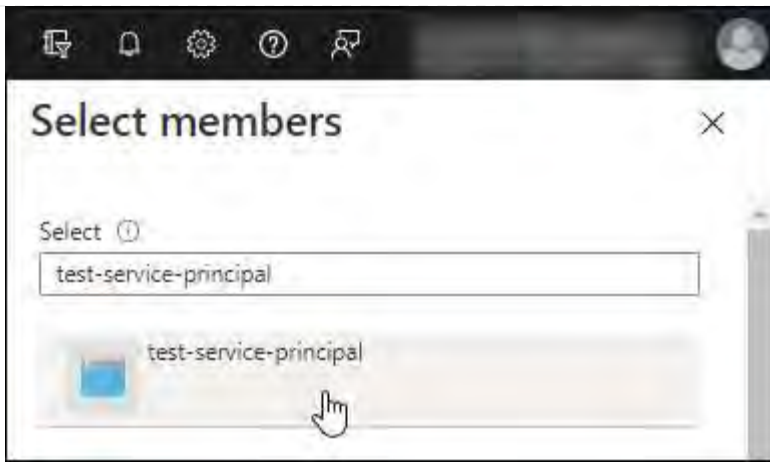
### Assign the custom role to the application

1. From the Azure portal, open the **Subscriptions** service.
2. Select the subscription.
3. Click **Access control (IAM) > Add > Add role assignment**.
4. In the **Role** tab, select the **BlueXP Operator** role and click **Next**.
5. In the **Members** tab, complete the following steps:
  - a. Keep **User, group, or service principal** selected.
  - b. Click **Select members**.



- c. Search for the name of the application.

Here's an example:



- d. Select the application and click **Select**.
  - e. Click **Next**.
6. Click **Review + assign**.

The service principal now has the required Azure permissions to deploy the Connector.

If you want to manage resources in multiple Azure subscriptions, then you must bind the service principal to each of those subscriptions. For example, BlueXP enables you to select the subscription that you want to use when deploying Cloud Volumes ONTAP.

#### **Add Windows Azure Service Management API permissions**

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Select **API permissions > Add a permission**.
3. Under **Microsoft APIs**, select **Azure Service Management**.

## Request API permissions

Select an API

Microsoft APIs   APIs my organization uses   My APIs

### Commonly used Microsoft APIs

<b>Microsoft Graph</b> Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10. Access Azure AD, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint.		
<b>Azure Batch</b> Schedule large-scale parallel and HPC applications in the cloud	<b>Azure Data Catalog</b> Programmatic access to Data Catalog resources to register, annotate and search data assets	<b>Azure Data Explorer</b> Perform ad-hoc queries on terabytes of data to build near real-time and complex analytics solutions
<b>Azure Data Lake</b> Access to storage and compute for big data analytic scenarios	<b>Azure DevOps</b> Integrate with Azure DevOps and Azure DevOps server	<b>Azure Import/Export</b> Programmatic control of import/export jobs
<b>Azure Key Vault</b> Manage your key vaults as well as the keys, secrets, and certificates within your Key Vaults	<b>Azure Rights Management Services</b> Allow validated users to read and write protected content	<b>Azure Service Management</b> Programmatic access to much of the functionality available through the Azure portal
<b>Azure Storage</b> Secure, massively scalable object and data lake storage for unstructured and semi-structured data	<b>Customer Insights</b> Create profile and interaction models for your products	<b>Data Export Service for Microsoft Dynamics 365</b> Export data from Microsoft Dynamics CRM organization to an external destination

4. Select **Access Azure Service Management as organization users** and then select **Add permissions**.



## Request API permissions

< All APIs



Azure Service Management

<https://management.azure.com/> [Docs](#)

What type of permissions does your application require?

### Delegated permissions

Your application needs to access the API as the signed-in user.

### Application permissions

Your application runs as a background service or daemon without a signed-in user.

Select permissions

[expand all](#)

Type to search

PERMISSION

ADMIN CONSENT REQUIRED




user\_impersonation

Access Azure Service Management as organization users (preview)

## Get the application ID and directory ID for the application

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Copy the **Application (client) ID** and the **Directory (tenant) ID**.



When you add the Azure account to BlueXP, you need to provide the application (client) ID and the directory (tenant) ID for the application. BlueXP uses the IDs to programmatically sign in.

## Create a client secret

1. Open the **Microsoft Entra ID** service.
2. Select **App registrations** and select your application.
3. Select **Certificates & secrets > New client secret**.
4. Provide a description of the secret and a duration.
5. Select **Add**.
6. Copy the value of the client secret.



## Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

[+ New client secret](#)

DESCRIPTION	EXPIRES	VALUE	Copy to clipboard
test secret	8/16/2020	*sZ1jSe2By:D*-ZRov4NLfdAcY7:+0vA	

You now have a client secret that BlueXP can use it to authenticate with Microsoft Entra ID.

### Result

Your service principal is now setup and you should have copied the application (client) ID, the directory (tenant) ID, and the value of the client secret. You need to enter this information in BlueXP when you create the Connector.

## Step 4: Create the Connector

Create the Connector directly from the BlueXP web-based console.

### About this task

- Creating the Connector from BlueXP deploys a virtual machine in Azure using a default configuration. After you create the Connector, you should not change to a smaller VM type that has less CPU or RAM. [Learn about the default configuration for the Connector.](#)
- When BlueXP deploys the Connector, it creates a custom role and assigns it to the Connector VM. This role includes permissions that enables the Connector to manage Azure resources. You need to ensure that the role is kept up to date as new permissions are added in subsequent releases. [Learn more about the custom role for the Connector.](#)

### Before you begin

You should have the following:

- An Azure subscription.
- A VNet and subnet in your Azure region of choice.
- Details about a proxy server, if your organization requires a proxy for all outgoing internet traffic:
  - IP address
  - Credentials
  - HTTPS certificate
- An SSH public key, if you want to use that authentication method for the Connector virtual machine. The other option for the authentication method is to use a password.

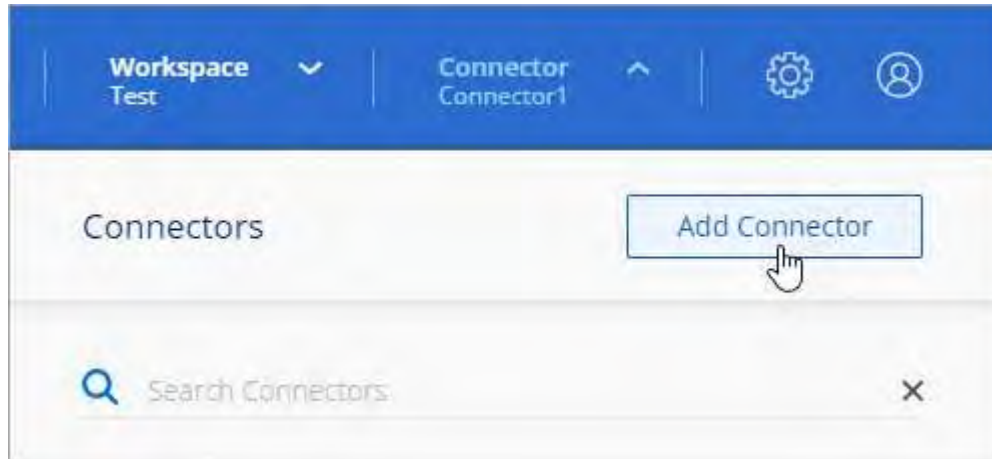
[Learn about connecting to a Linux VM in Azure](#)

- If you don't want BlueXP to automatically create an Azure role for the Connector, then you'll need to create your own [using the policy on this page.](#)

These permissions are for the Connector instance itself. It's a different set of permissions than what you previously set up to deploy the Connector VM.

## Steps

1. Select the **Connector** drop-down and select **Add Connector**.



2. Choose **Microsoft Azure** as your cloud provider.
3. On the **Deploying a Connector** page:
  - a. Under **Authentication**, select the authentication option that matches how you set up Azure permissions:
    - Select **Azure user account** to log in to your Microsoft account, which should have the required permissions.

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



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

- Select **Active Directory service principal** to enter information about the Microsoft Entra service principal that grants the required permissions:
  - Application (client) ID
  - Directory (tenant) ID
  - Client Secret

[Learn how to obtain these values for a service principal.](#)

4. Follow the steps in the wizard to create the Connector:
  - **VM Authentication:** Choose an Azure subscription, a location, a new resource group or an existing resource group, and then choose an authentication method for the Connector virtual machine that you're creating.

The authentication method for the virtual machine can be a password or an SSH public key.

[Learn about connecting to a Linux VM in Azure](#)

- **Details:** Enter a name for the instance, specify tags, and choose whether you want BlueXP to create a new role that has the required permissions, or if you want to select an existing role that you set up with [the required permissions](#).

Note that you can choose the Azure subscriptions associated with this role. Each subscription that you choose provides the Connector permissions to manage resources in that subscription (for example, Cloud Volumes ONTAP).

- **Network:** Choose a VNet and subnet, whether to enable a public IP address, and optionally specify a proxy configuration.
- **Security Group:** Choose whether to create a new security group or whether to select an existing security group that allows the required inbound and outbound rules.

[View security group rules for Azure.](#)

- **Review:** Review your selections to verify that your set up is correct.

5. Click **Add**.

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

## Result

After the process is complete, the Connector is available for use from BlueXP.

If you have Azure Blob storage in the same Azure subscription where you created the Connector, you'll see an Azure Blob storage working environment appear on the BlueXP canvas automatically. [Learn how to manage Azure Blob storage from BlueXP](#)

## Create a Connector from the Azure Marketplace

You can create a Connector in Azure directly from the Azure Marketplace. To create a Connector from the Azure Marketplace, you need to set up your networking, prepare Azure permissions, review instance requirements, and then create the Connector.

### Before you begin

- You should have an [understanding of Connectors](#).
- Review [Connector limitations](#).

### Step 1: Set up networking

Ensure that the network location where you plan to install the Connector supports the following requirements. These requirements enable the Connector to manage resources in your hybrid cloud.

#### Azure region

If you use Cloud Volumes ONTAP, the Connector should be deployed in the same Azure region as the Cloud Volumes ONTAP systems that it manages, or in the [Azure region pair](#) for the Cloud Volumes ONTAP systems. This requirement ensures that an Azure Private Link connection is used between Cloud Volumes ONTAP and its associated storage accounts.

[Learn how Cloud Volumes ONTAP uses an Azure Private Link](#)

#### VNet and subnet

When you create the Connector, you need to specify the VNet and subnet where the Connector should reside.

## Connections to target networks

A Connector requires a network connection to the location where you're planning to create and manage working environments. For example, the network where you plan to create Cloud Volumes ONTAP systems or a storage system in your on-premises environment.

## Outbound internet access

The network location where you deploy the Connector must have an outbound internet connection to contact specific endpoints.

## Endpoints contacted from the Connector

The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment for day-to-day operations.

Note that the endpoints listed below are all CNAME entries.

Endpoints	Purpose
<a href="https://management.azure.com">https://management.azure.com</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a> <a href="https://blob.core.windows.net">https://blob.core.windows.net</a> <a href="https://core.windows.net">https://core.windows.net</a>	To manage resources in Azure public regions.
<a href="https://management.chinacloudapi.cn">https://management.chinacloudapi.cn</a> <a href="https://login.chinacloudapi.cn">https://login.chinacloudapi.cn</a> <a href="https://blob.core.chinacloudapi.cn">https://blob.core.chinacloudapi.cn</a> <a href="https://core.chinacloudapi.cn">https://core.chinacloudapi.cn</a>	To manage resources in Azure China regions.
<a href="https://support.netapp.com">https://support.netapp.com</a> <a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	To obtain licensing information and to send AutoSupport messages to NetApp support.
<a href="https://*.api.bluexp.netapp.com">https://*.api.bluexp.netapp.com</a> <a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a> <a href="https://*.cloudmanager.cloud.netapp.com">https://*.cloudmanager.cloud.netapp.com</a> <a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	To provide SaaS features and services within BlueXP.
Choose between two sets of endpoints: <ul style="list-style-type: none"><li>Option 1 (recommended) <sup>1</sup> <a href="https://bluexpinfraprod.eastus2.data.azurecr.io">https://bluexpinfraprod.eastus2.data.azurecr.io</a> <a href="https://bluexpinfraprod.azurecr.io">https://bluexpinfraprod.azurecr.io</a></li><li>Option 2 <a href="https://*.blob.core.windows.net">https://*.blob.core.windows.net</a> <a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a></li></ul>	To obtain images for Connector upgrades.

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

### Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

### Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

### Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

Implement the networking requirements after creating the Connector.

### Step 2: Review VM requirements

When you create the Connector, choose a virtual machine type that meets the following requirements.

#### CPU

8 cores or 8 vCPUs

#### RAM

32 GB

#### Azure VM size

An instance type that meets the CPU and RAM requirements above. We recommend Standard\_D8s\_v3.

### **Step 3: Set up permissions**

You can provide permissions in the following ways:

- Option 1: Assign a custom role to the Azure VM using a system-assigned managed identity.
- Option 2: Provide BlueXP with the credentials for an Azure service principal that has the required permissions.

Follow these steps to set up permissions for BlueXP.

## Custom role

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would prefer to use a different method, refer to [Azure documentation](#)

## Steps

1. If you're planning to manually install the software on your own host, enable a system-assigned managed identity on the VM so that you can provide the required Azure permissions through a custom role.

[Microsoft Azure documentation: Configure managed identities for Azure resources on a VM using the Azure portal](#)

2. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
3. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription that you want to use with BlueXP.


## Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- a. Start [Azure Cloud Shell](#) and choose the Bash environment.
- b. Upload the JSON file.



c. Use the Azure CLI to create the custom role:

```
az role definition create --role-definition Connector_Policy.json
```

## Result

You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.

## Service principal


Create and set up a service principal in Microsoft Entra ID and obtain the Azure credentials that BlueXP needs.

### Create a Microsoft Entra application for role-based access control

1. Ensure that you have permissions in Azure to create an Active Directory application and to assign the application to a role.

For details, refer to [Microsoft Azure Documentation: Required permissions](#)

2. From the Azure portal, open the **Microsoft Entra ID** service.



3. In the menu, select **App registrations**.
4. Select **New registration**.
5. Specify details about the application:
  - **Name**: Enter a name for the application.
  - **Account type**: Select an account type (any will work with BlueXP).
  - **Redirect URI**: You can leave this field blank.
6. Select **Register**.

You've created the AD application and service principal.

### Assign the application to a role

1. Create a custom role:

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would



prefer to use a different method, refer to [Azure documentation](#)

- a. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
- b. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription from which users will create Cloud Volumes ONTAP systems.


### Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- Start [Azure Cloud Shell](#) and choose the Bash environment.
- Upload the JSON file.



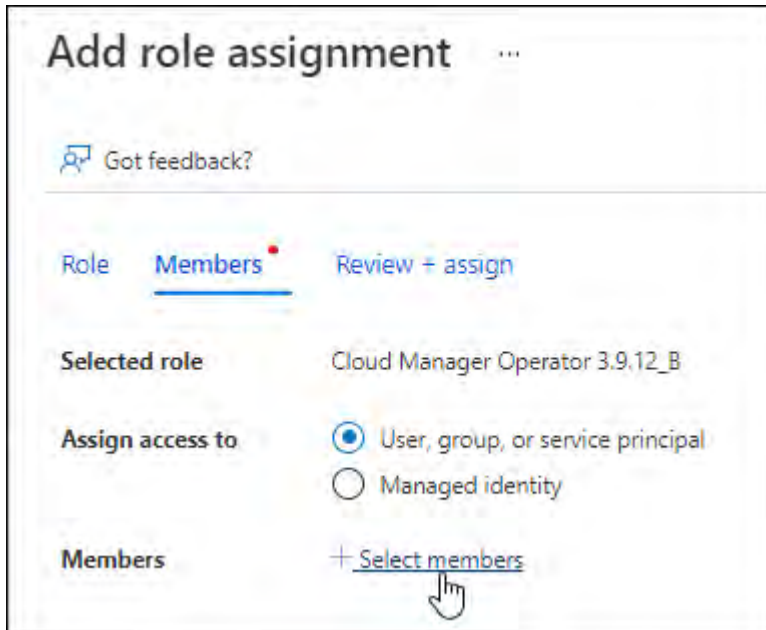
- Use the Azure CLI to create the custom role:

```
az role definition create --role-definition  
Connector_Policy.json
```

You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.


2. Assign the application to the role:

- a. From the Azure portal, open the **Subscriptions** service.
- b. Select the subscription.
- c. Select **Access control (IAM) > Add > Add role assignment**.
- d. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.
- e. In the **Members** tab, complete the following steps:
  - Keep **User, group, or service principal** selected.
  - Select **Select members**.



- Search for the name of the application.

Here's an example:



- Select the application and select **Select**.
  - Select **Next**.
- f. Select **Review + assign**.

The service principal now has the required Azure permissions to deploy the Connector.

If you want to deploy Cloud Volumes ONTAP from multiple Azure subscriptions, then you must bind the service principal to each of those subscriptions. BlueXP enables you to select the subscription that you want to use when deploying Cloud Volumes ONTAP.

### Add Windows Azure Service Management API permissions

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Select **API permissions > Add a permission**.
3. Under **Microsoft APIs**, select **Azure Service Management**.

#### Request API permissions

Select an API

Microsoft APIs   APIs my organization uses   My APIs

#### Commonly used Microsoft APIs

<b>Microsoft Graph</b> Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10. Access Azure AD, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint.		
<b>Azure Batch</b> Schedule large-scale parallel and HPC applications in the cloud	<b>Azure Data Catalog</b> Programmatic access to Data Catalog resources to register, annotate and search data assets	<b>Azure Data Explorer</b> Perform ad-hoc queries on terabytes of data to build near real-time and complex analytics solutions
<b>Azure Data Lake</b> Access to storage and compute for big data analytic scenarios	<b>Azure DevOps</b> Integrate with Azure DevOps and Azure DevOps server	<b>Azure Import/Export</b> Programmatic control of import/export jobs
<b>Azure Key Vault</b> Manage your key vaults as well as the keys, secrets, and certificates within your Key Vaults	<b>Azure Rights Management Services</b> Allow validated users to read and write protected content	<b>Azure Service Management</b> Programmatic access to much of the functionality available through the Azure portal
<b>Azure Storage</b> Secure, massively scalable object and data lake storage for unstructured and semi-structured data	<b>Customer Insights</b> Create profile and interaction models for your products	<b>Data Export Service for Microsoft Dynamics 365</b> Export data from Microsoft Dynamics CRM organization to an external destination

4. Select **Access Azure Service Management as organization users** and then select **Add permissions**.

## Request API permissions

< All APIs



Azure Service Management

<https://management.azure.com/> [Docs](#)

What type of permissions does your application require?

### Delegated permissions

Your application needs to access the API as the signed-in user.

### Application permissions

Your application runs as a background service or daemon without a signed-in user.

Select permissions

[expand all](#)

Type to search

PERMISSION

ADMIN CONSENT REQUIRED




user\_impersonation

Access Azure Service Management as organization users (preview)

## Get the application ID and directory ID for the application

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Copy the **Application (client) ID** and the **Directory (tenant) ID**.



When you add the Azure account to BlueXP, you need to provide the application (client) ID and the directory (tenant) ID for the application. BlueXP uses the IDs to programmatically sign in.

## Create a client secret

1. Open the **Microsoft Entra ID** service.
2. Select **App registrations** and select your application.
3. Select **Certificates & secrets > New client secret**.
4. Provide a description of the secret and a duration.
5. Select **Add**.
6. Copy the value of the client secret.

Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

[+ New client secret](#)

DESCRIPTION	EXPIRES	VALUE
test secret	8/16/2020	*sZ1jSe2By:D*-ZRov4NLfdAcY7:+0vA

Copy to clipboard

You now have a client secret that BlueXP can use it to authenticate with Microsoft Entra ID.

## Step 4: Create the Connector

Launch the Connector directly from the Azure Marketplace.

### About this task

Creating the Connector from the Azure Marketplace sets up a virtual machine with a default configuration. [Learn about the default configuration for the Connector.](#)

### Before you begin

You should have the following:

- An Azure subscription.
- A VNet and subnet in your Azure region of choice.
- Details about a proxy server, if your organization requires a proxy for all outgoing internet traffic:
  - IP address
  - Credentials
  - HTTPS certificate
- An SSH public key, if you want to use that authentication method for the Connector virtual machine. The other option for the authentication method is to use a password.

[Learn about connecting to a Linux VM in Azure](#)

- If you don't want BlueXP to automatically create an Azure role for the Connector, then you'll need to create your own [using the policy on this page](#).

These permissions are for the Connector instance itself. It's a different set of permissions than what you previously set up to deploy the Connector VM.

### Steps

1. Go to the NetApp Connector VM page in the Azure Marketplace.

[Azure Marketplace page for commercial regions](#)

2. Select **Get it now** and then select **Continue**.
3. From the Azure portal, select **Create** and follow the steps to configure the virtual machine.

Note the following as you configure the VM:

- **VM size:** Choose a VM size that meets CPU and RAM requirements. We recommend Standard\_D8s\_v3.
- **Disks:** The Connector can perform optimally with either HDD or SSD disks.
- **Network security group:** The Connector requires inbound connections using SSH, HTTP, and HTTPS.

[View security group rules for Azure.](#)

- **Identity:** Under **Management**, select **Enable system assigned managed identity**.

This setting is important because a managed identity allows the Connector virtual machine to identify itself to Microsoft Entra ID without providing any credentials. [Learn more about managed identities for Azure resources.](#)

4. On the **Review + create** page, review your selections and select **Create** to start the deployment.

Azure deploys the virtual machine with the specified settings. You should see the virtual machine and Connector software running in about five minutes.

5. Open a web browser from a host that has a connection to the Connector virtual machine and enter the following URL:

`https://ipaddress`

6. After you log in, set up the Connector:
  - a. Specify the BlueXP organization to associate with the Connector.
  - b. Enter a name for the system.
  - c. Under **Are you running in a secured environment?** keep restricted mode disabled.

Keep restricted mode disabled to use BlueXP in standard mode. You should enable restricted mode only if you have a secure environment and want to disconnect this account from BlueXP backend services. If that's the case, [follow steps to get started with BlueXP in restricted mode.](#)

- d. Select **Let's start**.

## Result

You have now installed the Connector and set it up with your BlueXP organization.

If you have Azure Blob storage in the same Azure subscription where you created the Connector, you'll see an Azure Blob storage working environment appear on the BlueXP canvas automatically. [Learn how to manage Azure Blob storage from BlueXP](#)

## Step 5: Provide permissions to BlueXP

Now that you've created the Connector, you need to provide BlueXP with the permissions that you previously set up. Providing the permissions enables BlueXP to manage your data and storage infrastructure in Azure.

## Custom role

Go to the Azure portal and assign the Azure custom role to the Connector virtual machine for one or more subscriptions.

### Steps

1. From the Azure Portal, open the **Subscriptions** service and select your subscription.

It's important to assign the role from the **Subscriptions** service because this specifies the scope of the role assignment at the subscription level. The *scope* defines the set of resources that the access applies to. If you specify a scope at a different level (for example, at the virtual machine level), your ability to complete actions from within BlueXP will be affected.

[Microsoft Azure documentation: Understand scope for Azure RBAC](#)

2. Select **Access control (IAM) > Add > Add role assignment**.
3. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.



BlueXP Operator is the default name provided in the BlueXP policy. If you chose a different name for the role, then select that name instead.

4. In the **Members** tab, complete the following steps:
  - a. Assign access to a **Managed identity**.
  - b. Select **Select members**, select the subscription in which the Connector virtual machine was created, under **Managed identity**, choose **Virtual machine**, and then select the Connector virtual machine.
  - c. Select **Select**.
  - d. Select **Next**.
  - e. Select **Review + assign**.
  - f. If you want to manage resources in additional Azure subscriptions, switch to that subscription and then repeat these steps.

## Result

BlueXP now has the permissions that it needs to perform actions in Azure on your behalf.

## What's next?

Go to the [BlueXP console](#) to start using the Connector with BlueXP.

## Service principal

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



2. Select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location**: Select **Microsoft Azure > Connector**.



- b. **Define Credentials:** Enter information about the Microsoft Entra service principal that grants the required permissions:
  - Application (client) ID
  - Directory (tenant) ID
  - Client Secret
- c. **Marketplace Subscription:** Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
- d. **Review:** Confirm the details about the new credentials and select **Add**.

### Result

BlueXP now has the permissions that it needs to perform actions in Azure on your behalf.

## Manually install the Connector in Azure

A Connector is NetApp software running in your cloud network or on-premises network that gives you the ability to use all BlueXP features and services. One of the available installation options is to manually install the Connector software on a Linux host running in Azure. To manually install the Connector on your own Linux host, you need to review host requirements, set up your networking, prepare Azure permissions, install the Connector, and then provide the permissions that you prepared.

### Before you begin

- You should have an [understanding of Connectors](#).
- You should review [Connector limitations](#).

### Step 1: Review host requirements

The Connector software must run on a host that meets specific operating system requirements, RAM requirements, port requirements, and so on.



The Connector reserves the UID and GID range of 19000 to 19200. This range is fixed and cannot be modified. If any third-party software on your host is using UIDs or GIDs within this range, the Connector installation will fail. NetApp recommends using a host that is free of third-party software to avoid conflicts.

### Dedicated host

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

Host can be of any architecture that meets the following size requirements:

- CPU: 8 cores or 8 vCPUs
- RAM: 32 GB

### Hypervisor

A bare metal or hosted hypervisor that is certified to run a supported operating system is required.



## Operating system and container requirements

BlueXP supports the Connector with the following operating systems when using BlueXP in standard mode or restricted mode. A container orchestration tool is required before you install the Connector.

Operating system	Supported OS versions	Supported Connector versions	Required container tool	SELinux
Red Hat Enterprise Linux	9.1 to 9.4 8.6 to 8.10 7.9	3.9.40 or later with BlueXP in standard mode or restricted mode	Podman version 4.6.1 or 4.9.4 <a href="#">View Podman configuration requirements.</a>	Supported in enforcing mode or permissive mode <sup>1</sup>
Ubuntu	24.04 LTS	3.9.45 or later with BlueXP in standard mode or restricted mode	Docker Engine 23.06 to 28.0.0.	Not supported
	22.04 LTS	3.9.29 or later	Docker Engine 23.0.6 to 28.0.0.	Not supported

Notes:

1. Management of Cloud Volumes ONTAP systems is not supported by Connectors that have SELinux enabled on the operating system.
2. The Connector is supported on English-language versions of these operating systems.
3. For RHEL, the host must be registered with Red Hat Subscription Management. If it's not registered, the host can't access repositories to update required 3rd-party software during Connector installation.

### Azure VM size

An instance type that meets the CPU and RAM requirements above. We recommend Standard\_D8s\_v3.

### Disk space in /opt

100 GiB of space must be available

BlueXP uses /opt to install the /opt/application/netapp directory and its contents.

### Disk space in /var

20 GiB of space must be available

BlueXP requires this space in /var because Docker or Podman are architected to create the containers within this directory. Specifically, they will create containers in the /var/lib/containers/storage directory. External mounts or symlinks do not work for this space.

## Step 2: Install Podman or Docker Engine

Depending on your operating system, either Podman or Docker Engine is required before you install the Connector.

- Podman is required for Red Hat Enterprise Linux 8 and 9.

[View the Podman versions that BlueXP supports.](#)

- Docker Engine is required for Ubuntu.

[View the Docker Engine versions that BlueXP supports.](#)

## Example 2. Steps

### Podman

Follow these steps to install Podman and configure it to meet the following requirements:

- Enable and start the podman.socket service
- Install python3
- Install the podman-compose package version 1.0.6
- Add podman-compose to the PATH environment variable



When using Podman, adjust the aardvark-dns service port (default: 53) after installing the Connector to avoid conflicts with the DNS port on the host. Follow the instructions to configure the port.

### Steps

1. Remove the podman-docker package if it's installed on the host.

```
dnf remove podman-docker
rm /var/run/docker.sock
```

2. Install Podman.

Podman is available from official Red Hat Enterprise Linux repositories.

For Red Hat Enterprise Linux 9:

```
sudo dnf install podman-2:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

For Red Hat Enterprise Linux 8:

```
sudo dnf install podman-3:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

3. Enable and start the podman.socket service.

```
sudo systemctl enable --now podman.socket
```

4. Install python3.

```
sudo dnf install python3
```

5. Install the EPEL repository package if it's not already available on your system.

This step is required because podman-compose is available from the Extra Packages for Enterprise Linux (EPEL) repository.

For Red Hat Enterprise Linux 9:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-9.noarch.rpm
```

For Red Hat Enterprise Linux 8:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

6. Install podman-compose package 1.0.6.

```
sudo dnf install podman-compose-1.0.6
```



Using the `dnf install` command meets the requirement for adding podman-compose to the PATH environment variable. The installation command adds podman-compose to `/usr/bin`, which is already included in the `secure_path` option on the host.

### Docker Engine

Follow the documentation from Docker to install Docker Engine.

#### Steps

1. [View installation instructions from Docker](#)

Be sure to follow the steps to install a specific version of Docker Engine. Installing the latest version will install a version of Docker that BlueXP doesn't support.

2. Verify that Docker is enabled and running.

```
sudo systemctl enable docker && sudo systemctl start docker
```

### Step 3: Set up networking

Ensure that the network location where you plan to install the Connector supports the following requirements. Meeting these requirements enables the Connector to manage resources and processes within your hybrid

cloud environment.

## Azure region

If you use Cloud Volumes ONTAP, the Connector should be deployed in the same Azure region as the Cloud Volumes ONTAP systems that it manages, or in the [Azure region pair](#) for the Cloud Volumes ONTAP systems. This requirement ensures that an Azure Private Link connection is used between Cloud Volumes ONTAP and its associated storage accounts.

[Learn how Cloud Volumes ONTAP uses an Azure Private Link](#)

## Connections to target networks

A Connector requires a network connection to the location where you're planning to create and manage working environments. For example, the network where you plan to create Cloud Volumes ONTAP systems or a storage system in your on-premises environment.

## Outbound internet access

The network location where you deploy the Connector must have an outbound internet connection to contact specific endpoints.

## Endpoints contacted from computers when using the BlueXP web-based console

Computers that access the BlueXP console from a web browser must have the ability to contact several endpoints. You'll need to use the BlueXP console to set up the Connector and for day-to-day use of BlueXP.

[Prepare networking for the BlueXP console.](#)

## Endpoints contacted during manual installation

When you manually install the Connector on your own Linux host, the installer for the Connector requires access to the following URLs during the installation process:

- <https://mysupport.netapp.com>
- <https://signin.b2c.netapp.com> (this endpoint is the CNAME URL for <https://mysupport.netapp.com>)
- <https://cloudmanager.cloud.netapp.com/tenancy>
- <https://stream.cloudmanager.cloud.netapp.com>
- <https://production-artifacts.cloudmanager.cloud.netapp.com>
- To obtain images, the installer needs access to one of these two sets of endpoints:
  - Option 1 (recommended):
    - <https://bluexpinfraprod.eastus2.data.azurecr.io>
    - <https://bluexpinfraprod.azurecr.io>
  - Option 2:
    - [https://\\*.blob.core.windows.net](https://*.blob.core.windows.net)
    - <https://cloudmanagerinfraprod.azurecr.io>

The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.

- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

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

### Endpoints contacted from the Connector

The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment for day-to-day operations.

Note that the endpoints listed below are all CNAME entries.

Endpoints	Purpose
<a href="https://management.azure.com">https://management.azure.com</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a> <a href="https://blob.core.windows.net">https://blob.core.windows.net</a> <a href="https://core.windows.net">https://core.windows.net</a>	To manage resources in Azure public regions.
<a href="https://management.chinacloudapi.cn">https://management.chinacloudapi.cn</a> <a href="https://login.chinacloudapi.cn">https://login.chinacloudapi.cn</a> <a href="https://blob.core.chinacloudapi.cn">https://blob.core.chinacloudapi.cn</a> <a href="https://core.chinacloudapi.cn">https://core.chinacloudapi.cn</a>	To manage resources in Azure China regions.
<a href="https://support.netapp.com">https://support.netapp.com</a> <a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	To obtain licensing information and to send AutoSupport messages to NetApp support.
<a href="https://*.api.bluexp.netapp.com">https://*.api.bluexp.netapp.com</a> <a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a> <a href="https://*.cloudmanager.cloud.netapp.com">https://*.cloudmanager.cloud.netapp.com</a> <a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	To provide SaaS features and services within BlueXP.
Choose between two sets of endpoints: <ul style="list-style-type: none"> <li>• Option 1 (recommended) <sup>1</sup>  <a href="https://bluexpinfraprod.eastus2.data.azurecr.io">https://bluexpinfraprod.eastus2.data.azurecr.io</a>  <a href="https://bluexpinfraprod.azurecr.io">https://bluexpinfraprod.azurecr.io</a> </li> <li>• Option 2            <a href="https://*.blob.core.windows.net">https://*.blob.core.windows.net</a>  <a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a> </li> </ul>	To obtain images for Connector upgrades.

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

### Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

### Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

### Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

### Step 4: Set up Connector deployment permissions

You need to provide Azure permissions to BlueXP by using one of the following options:

- Option 1: Assign a custom role to the Azure VM using a system-assigned managed identity.
- Option 2: Provide BlueXP with the credentials for an Azure service principal that has the required permissions.

Follow the steps to prepare permissions for BlueXP.

## Create a custom role for Connector deployment

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would prefer to use a different method, refer to [Azure documentation](#)

### Steps

1. If you're planning to manually install the software on your own host, enable a system-assigned managed identity on the VM so that you can provide the required Azure permissions through a custom role.

[Microsoft Azure documentation: Configure managed identities for Azure resources on a VM using the Azure portal](#)

2. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
3. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription that you want to use with BlueXP.


### Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- a. Start [Azure Cloud Shell](#) and choose the Bash environment.
- b. Upload the JSON file.





c. Use the Azure CLI to create the custom role:

```
az role definition create --role-definition Connector_Policy.json
```

## Result

You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.

## Service principal

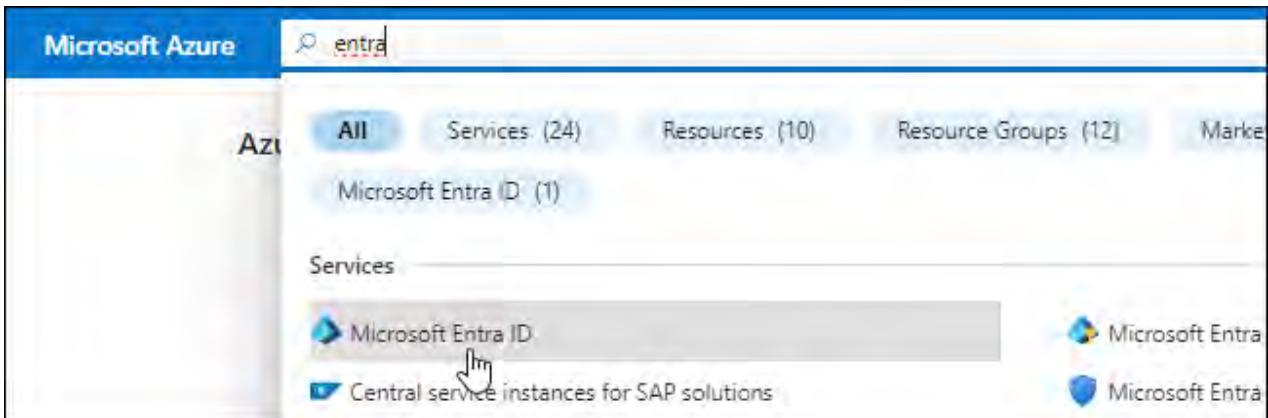
Create and set up a service principal in Microsoft Entra ID and obtain the Azure credentials that BlueXP needs.

### Create a Microsoft Entra application for role-based access control

1. Ensure that you have permissions in Azure to create an Active Directory application and to assign the application to a role.

For details, refer to [Microsoft Azure Documentation: Required permissions](#)

2. From the Azure portal, open the **Microsoft Entra ID** service.



3. In the menu, select **App registrations**.
4. Select **New registration**.
5. Specify details about the application:
  - **Name:** Enter a name for the application.
  - **Account type:** Select an account type (any will work with BlueXP).
  - **Redirect URI:** You can leave this field blank.
6. Select **Register**.

You've created the AD application and service principal.

### Assign the application to a role

1. Create a custom role:

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would

prefer to use a different method, refer to [Azure documentation](#)

- a. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
- b. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription from which users will create Cloud Volumes ONTAP systems.


### Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- Start [Azure Cloud Shell](#) and choose the Bash environment.
- Upload the JSON file.




- Use the Azure CLI to create the custom role:

```
az role definition create --role-definition  
Connector_Policy.json
```

You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.

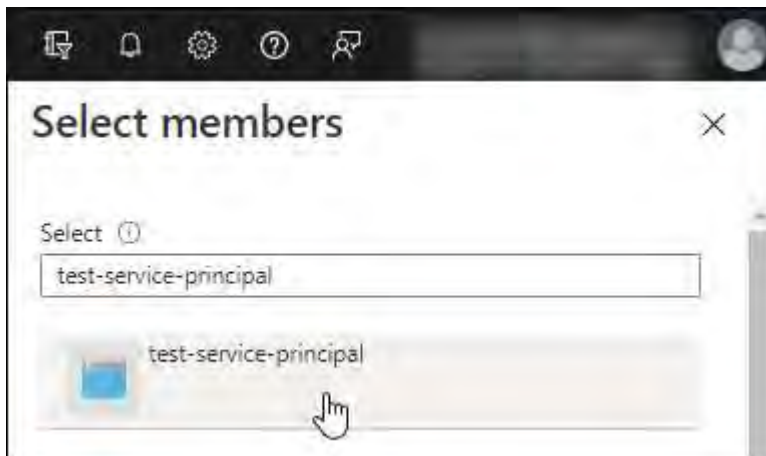
2. Assign the application to the role:

- a. From the Azure portal, open the **Subscriptions** service.
- b. Select the subscription.
- c. Select **Access control (IAM) > Add > Add role assignment**.
- d. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.
- e. In the **Members** tab, complete the following steps:
  - Keep **User, group, or service principal** selected.
  - Select **Select members**.



- Search for the name of the application.

Here's an example:



- Select the application and select **Select**.
  - Select **Next**.
- f. Select **Review + assign**.

The service principal now has the required Azure permissions to deploy the Connector.

If you want to deploy Cloud Volumes ONTAP from multiple Azure subscriptions, then you must bind the service principal to each of those subscriptions. BlueXP enables you to select the subscription that you want to use when deploying Cloud Volumes ONTAP.

### Add Windows Azure Service Management API permissions

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Select **API permissions > Add a permission**.
3. Under **Microsoft APIs**, select **Azure Service Management**.

#### Request API permissions

Select an API

Microsoft APIs   APIs my organization uses   My APIs

#### Commonly used Microsoft APIs

<b>Microsoft Graph</b> Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10. Access Azure AD, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint.		
<b>Azure Batch</b> Schedule large-scale parallel and HPC applications in the cloud	<b>Azure Data Catalog</b> Programmatic access to Data Catalog resources to register, annotate and search data assets	<b>Azure Data Explorer</b> Perform ad-hoc queries on terabytes of data to build near real-time and complex analytics solutions
<b>Azure Data Lake</b> Access to storage and compute for big data analytic scenarios	<b>Azure DevOps</b> Integrate with Azure DevOps and Azure DevOps server	<b>Azure Import/Export</b> Programmatic control of import/export jobs
<b>Azure Key Vault</b> Manage your key vaults as well as the keys, secrets, and certificates within your Key Vaults	<b>Azure Rights Management Services</b> Allow validated users to read and write protected content	<b>Azure Service Management</b> Programmatic access to much of the functionality available through the Azure portal
<b>Azure Storage</b> Secure, massively scalable object and data lake storage for unstructured and semi-structured data	<b>Customer Insights</b> Create profile and interaction models for your products	<b>Data Export Service for Microsoft Dynamics 365</b> Export data from Microsoft Dynamics CRM organization to an external destination

4. Select **Access Azure Service Management as organization users** and then select **Add permissions**.

## Request API permissions

< All APIs



Azure Service Management

<https://management.azure.com/> [Docs](#)

What type of permissions does your application require?

### Delegated permissions

Your application needs to access the API as the signed-in user.

### Application permissions

Your application runs as a background service or daemon without a signed-in user.

Select permissions

[expand all](#)

Type to search

PERMISSION

ADMIN CONSENT REQUIRED

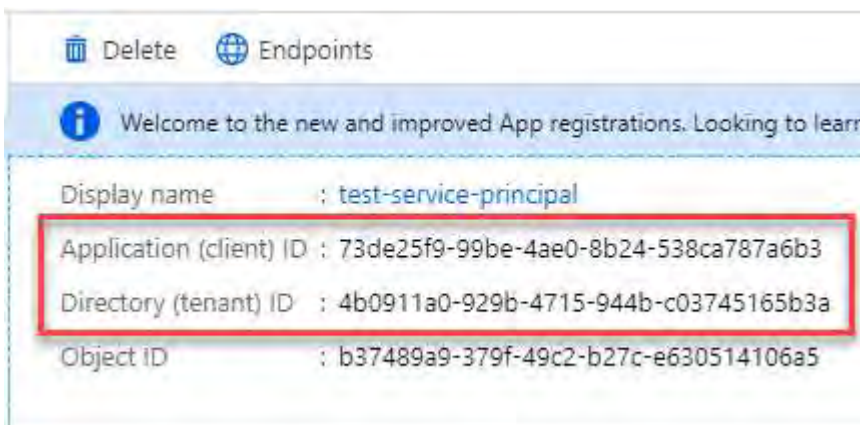


user\_impersonation

Access Azure Service Management as organization users (preview)

## Get the application ID and directory ID for the application

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Copy the **Application (client) ID** and the **Directory (tenant) ID**.



When you add the Azure account to BlueXP, you need to provide the application (client) ID and the directory (tenant) ID for the application. BlueXP uses the IDs to programmatically sign in.

## Create a client secret

1. Open the **Microsoft Entra ID** service.
2. Select **App registrations** and select your application.
3. Select **Certificates & secrets > New client secret**.
4. Provide a description of the secret and a duration.
5. Select **Add**.
6. Copy the value of the client secret.

## Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

[+ New client secret](#)

DESCRIPTION	EXPIRES	VALUE	Copy to clipboard
test secret	8/16/2020	*sZ1jSe2By:D*-ZR0V4NLfdAcY7:+0vA	

You now have a client secret that BlueXP can use it to authenticate with Microsoft Entra ID.

### Result

Your service principal is now setup and you should have copied the application (client) ID, the directory (tenant) ID, and the value of the client secret. You need to enter this information in BlueXP when you add an Azure account.

## Step 5: Install the Connector

After the pre-requisites are complete, you can manually install the software on your own Linux host.

### Before you begin

You should have the following:

- Root privileges to install the Connector.
- Details about a proxy server, if a proxy is required for internet access from the Connector.

You have the option to configure a proxy server after installation but doing so requires restarting the Connector.

- A CA-signed certificate, if the proxy server uses HTTPS or if the proxy is an intercepting proxy.



You cannot set a certificate for a transparent proxy server when manually installing the Connector. If you need to set a certificate for a transparent proxy server, you must use the Maintenance Console after installation. Learn more about the [Connector Maintenance Console](#).

- A managed identity enabled on the VM in Azure so that you can provide the required Azure permissions through a custom role.

[Microsoft Azure documentation: Configure managed identities for Azure resources on a VM using the Azure portal](#)

### About this task

The installer that is available on the NetApp Support Site might be an earlier version. After installation, the Connector automatically updates itself if a new version is available.

### Steps

1. If the `http_proxy` or `https_proxy` system variables are set on the host, remove them:

```
unset http_proxy
unset https_proxy
```

If you don't remove these system variables, the installation will fail.

2. Download the Connector software from the [NetApp Support Site](#), and then copy it to the Linux host.

You should download the "online" Connector installer that's meant for use in your network or in the cloud. A separate "offline" installer is available for the Connector, but it's only supported with private mode deployments.

3. Assign permissions to run the script.

```
chmod +x BlueXP-Connector-Cloud-<version>
```

Where <version> is the version of the Connector that you downloaded.

4. Run the installation script.

```
./BlueXP-Connector-Cloud-<version> --proxy <HTTP or HTTPS proxy server>
--cacert <path and file name of a CA-signed certificate>
```

You'll need to add proxy information if your network requires a proxy for internet access. You can add either a transparent or explicit proxy. The `--proxy` and `--cacert` parameters are optional and you won't be prompted to add them. If you have a proxy server, you will need to enter the parameters as shown.

Here is an example configuring an explicit proxy server with a CA-signed certificate:

```
./BlueXP-Connector-Cloud-v3.9.40--proxy
https://user:password@10.0.0.30:8080/ --cacert
/tmp/cacert/certificate.cer
```

`--proxy` configures the Connector to use an HTTP or HTTPS proxy server using one of the following formats:

- `http://address:port`
- `http://user-name:password@address:port`
- `http://domain-name%92user-name:password@address:port`
- `https://address:port`
- `https://user-name:password@address:port`
- `https://domain-name%92user-name:password@address:port`

Note the following:

- The user can be a local user or domain user.



- For a domain user, you must use the ASCII code for a \ as shown above.
- BlueXP doesn't support user names or passwords that include the @ character.
- If the password includes any of the following special characters, you must escape that special character by prepending it with a backslash: & or !

For example:

```
http://bxpproxyuser:netapp1!\@address:3128
```

--cacert specifies a CA-signed certificate to use for HTTPS access between the Connector and the proxy server. This parameter is required for HTTPS proxy servers, intercepting proxy servers, and transparent proxy servers.

Here is an example configuring a transparent proxy server. When you configure a transparent proxy, you don't need to define the proxy server. You only add a CA-signed certificate to your Connector host:

```
./BlueXP-Connector-Cloud-v3.9.40 --cacert /tmp/cacert/certificate.cer
```

5. If you used Podman, you'll need to adjust the aardvark-dns port.
  - a. SSH to the BlueXP Connector virtual machine.
  - b. Open podman `/usr/share/containers/containers.conf` file and modify the chosen port for Aardvark DNS service. For example, change it to 54.

```
vi /usr/share/containers/containers.conf
...
# Port to use for dns forwarding daemon with netavark in rootful
bridge
# mode and dns enabled.
# Using an alternate port might be useful if other DNS services
should
# run on the machine.
#
dns_bind_port = 54
...
Esc:wq
```

- c. Reboot the Connector virtual machine.
6. Wait for the installation to complete.

At the end of the installation, the Connector service (occm) restarts twice if you specified a proxy server.

7. Open a web browser from a host that has a connection to the Connector virtual machine and enter the following URL:

`https://ipaddress`

8. After you log in, set up the Connector:



- a. Specify the BlueXP organization to associate with the Connector.
- b. Enter a name for the system.
- c. Under **Are you running in a secured environment?** keep restricted mode disabled.

You should keep restricted mode disabled because these steps describe how to use BlueXP in standard mode. You should enable restricted mode only if you have a secure environment and want to disconnect this account from BlueXP backend services. If that's the case, [follow steps to get started with BlueXP in restricted mode](#).

- d. Select **Let's start**.

If you have Azure Blob storage in the same Azure subscription where you created the Connector, you'll see an Azure Blob storage working environment appear on the BlueXP canvas automatically. [Learn how to manage Azure Blob storage from BlueXP](#)

### **Step 6: Provide permissions to BlueXP**

Now that you've installed the Connector, you need to provide BlueXP with the Azure permissions that you previously set up. Providing the permissions enables BlueXP to manage your data and storage infrastructure in Azure.

## Custom role

Go to the Azure portal and assign the Azure custom role to the Connector virtual machine for one or more subscriptions.

### Steps

1. From the Azure Portal, open the **Subscriptions** service and select your subscription.

It's important to assign the role from the **Subscriptions** service because this specifies the scope of the role assignment at the subscription level. The *scope* defines the set of resources that the access applies to. If you specify a scope at a different level (for example, at the virtual machine level), your ability to complete actions from within BlueXP will be affected.

[Microsoft Azure documentation: Understand scope for Azure RBAC](#)

2. Select **Access control (IAM) > Add > Add role assignment**.
3. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.



BlueXP Operator is the default name provided in the BlueXP policy. If you chose a different name for the role, then select that name instead.

4. In the **Members** tab, complete the following steps:
  - a. Assign access to a **Managed identity**.
  - b. Select **Select members**, select the subscription in which the Connector virtual machine was created, under **Managed identity**, choose **Virtual machine**, and then select the Connector virtual machine.
  - c. Select **Select**.
  - d. Select **Next**.
  - e. Select **Review + assign**.
  - f. If you want to manage resources in additional Azure subscriptions, switch to that subscription and then repeat these steps.

## Result

BlueXP now has the permissions that it needs to perform actions in Azure on your behalf.

## What's next?

Go to the [BlueXP console](#) to start using the Connector with BlueXP.

## Service principal

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



2. Select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location**: Select **Microsoft Azure > Connector**.

- b. **Define Credentials:** Enter information about the Microsoft Entra service principal that grants the required permissions:
  - Application (client) ID
  - Directory (tenant) ID
  - Client Secret
- c. **Marketplace Subscription:** Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
- d. **Review:** Confirm the details about the new credentials and select **Add**.

### Result

BlueXP now has the permissions that it needs to perform actions in Azure on your behalf.

## Google Cloud

### Connector installation options in Google Cloud

There are a few different ways to create a Connector in Google Cloud. Directly from BlueXP is the most common way.

The following installation options are available:

- [Create the Connector directly from BlueXP](#) (this is the standard option)

This action launches a VM instance running Linux and the Connector software in a VPC of your choice.

- [Create the Connector using gcloud](#)

This action also launches a VM instance running Linux and the Connector software, but the deployment is initiated directly from Google Cloud, rather than from BlueXP.

- [Download and manually install the software on your own Linux host](#)

The installation option that you choose impacts how you prepare for installation. This includes how you provide BlueXP with the required permissions that it needs to authenticate and manage resources in Google Cloud.

### Create a Connector in Google Cloud from BlueXP or gcloud

You can create a Connector in Google Cloud from BlueXP or by using Google Cloud. You need to set up your networking, prepare Google Cloud permissions, enable Google Cloud APIs, and then create the Connector.

### Before you begin

- You should have an [understanding of Connectors](#).
- You should review [Connector limitations](#).

### Step 1: Set up networking

Set up networking to ensure the Connector can manage resources, with connections to target networks and outbound internet access.

## VPC and subnet

When you create the Connector, you need to specify the VPC and subnet where the Connector should reside.

## Connections to target networks

A Connector requires a network connection to the location where you're planning to create and manage working environments. For example, the network where you plan to create Cloud Volumes ONTAP systems or a storage system in your on-premises environment.

## Outbound internet access

The network location where you deploy the Connector must have an outbound internet connection to contact specific endpoints.

## Endpoints contacted from the Connector

The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment for day-to-day operations.

Note that the endpoints listed below are all CNAME entries.

Endpoints	Purpose
<a href="https://www.googleapis.com/compute/v1/">https://www.googleapis.com/compute/v1/</a> <a href="https://compute.googleapis.com/compute/v1">https://compute.googleapis.com/compute/v1</a> <a href="https://cloudresourcemanager.googleapis.com/v1/projects">https://cloudresourcemanager.googleapis.com/v1/projects</a> <a href="https://www.googleapis.com/compute/beta">https://www.googleapis.com/compute/beta</a> <a href="https://storage.googleapis.com/storage/v1">https://storage.googleapis.com/storage/v1</a> <a href="https://www.googleapis.com/storage/v1">https://www.googleapis.com/storage/v1</a> <a href="https://iam.googleapis.com/v1">https://iam.googleapis.com/v1</a> <a href="https://cloudkms.googleapis.com/v1">https://cloudkms.googleapis.com/v1</a> <a href="https://www.googleapis.com/deploymentmanager/v2/projects">https://www.googleapis.com/deploymentmanager/v2/projects</a>	To manage resources in Google Cloud.
<a href="https://support.netapp.com">https://support.netapp.com</a> <a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	To obtain licensing information and to send AutoSupport messages to NetApp support.
<a href="https://*.api.bluexp.netapp.com">https://*.api.bluexp.netapp.com</a> <a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a> <a href="https://*.cloudmanager.cloud.netapp.com">https://*.cloudmanager.cloud.netapp.com</a> <a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	To provide SaaS features and services within BlueXP.

Endpoints	Purpose
<p>Choose between two sets of endpoints:</p> <ul style="list-style-type: none"> <li>• Option 1 (recommended) <sup>1</sup></li> </ul> <p>https://bluexpinfraprod.eastus2.data.azurecr.io https://bluexpinfraprod.azurecr.io</p> <ul style="list-style-type: none"> <li>• Option 2</li> </ul> <p>https://*.blob.core.windows.net https://cloudmanagerinfraprod.azurecr.io</p>	<p>To obtain images for Connector upgrades.</p>

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

### Endpoints contacted from the BlueXP console

As you use the BlueXP web-based console that's provided through the SaaS layer, it contacts several endpoints to complete data management tasks. This includes endpoints that are contacted to deploy the Connector from the BlueXP console.

[View the list of endpoints contacted from the BlueXP console.](#)

### Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

### Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

## Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

Implement this networking requirement after creating the Connector.

## Step 2: Set up permissions to create the Connector

Before you can deploy a Connector from BlueXP or by using gcloud, you need to set up permissions for the Google Cloud user who will deploy the Connector VM.

### Steps

1. Create a custom role in Google Cloud:
  - a. Create a YAML file that includes the following permissions:

```
title: Connector deployment policy
description: Permissions for the user who deploys the Connector from
BlueXP
stage: GA
includedPermissions:
- compute.disks.create
- compute.disks.get
- compute.disks.list
- compute.disks.setLabels
- compute.disks.use
- compute.firewalls.create
- compute.firewalls.delete
- compute.firewalls.get
- compute.firewalls.list
- compute.globalOperations.get
- compute.images.get
- compute.images.getFromFamily
- compute.images.list
- compute.images.useReadOnly
- compute.instances.attachDisk
- compute.instances.create
- compute.instances.get
- compute.instances.list
- compute.instances.setDeletionProtection
- compute.instances.setLabels
- compute.instances.setMachineType
- compute.instances.setMetadata
```

```

- compute.instances.setTags
- compute.instances.start
- compute.instances.updateDisplayDevice
- compute.machineTypes.get
- compute.networks.get
- compute.networks.list
- compute.networks.updatePolicy
- compute.projects.get
- compute.regions.get
- compute.regions.list
- compute.subnetworks.get
- compute.subnetworks.list
- compute.zoneOperations.get
- compute.zones.get
- compute.zones.list
- deploymentmanager.compositeTypes.get
- deploymentmanager.compositeTypes.list
- deploymentmanager.deployments.create
- deploymentmanager.deployments.delete
- deploymentmanager.deployments.get
- deploymentmanager.deployments.list
- deploymentmanager.manifests.get
- deploymentmanager.manifests.list
- deploymentmanager.operations.get
- deploymentmanager.operations.list
- deploymentmanager.resources.get
- deploymentmanager.resources.list
- deploymentmanager.typeProviders.get
- deploymentmanager.typeProviders.list
- deploymentmanager.types.get
- deploymentmanager.types.list
- resourcemanager.projects.get
- compute.instances.setServiceAccount
- iam.serviceAccounts.list

```

- b. From Google Cloud, activate cloud shell.
- c. Upload the YAML file that includes the required permissions.
- d. Create a custom role by using the `gcloud iam roles create` command.

The following example creates a role named "connectorDeployment" at the project level:

```
gcloud iam roles create connectorDeployment --project=myproject --file=connector-deployment.yaml
```

[Google Cloud docs: Creating and managing custom roles](#)

2. Assign this custom role to the user who will deploy the Connector from BlueXP or by using `gcloud`.

### Step 3: Set up permissions for the Connector

A Google Cloud service account is required to provide the Connector with the permissions that BlueXP needs to manage resources in Google Cloud. When you create the Connector, you'll need to associate this service account with the Connector VM.

It's your responsibility to update the custom role as new permissions are added in subsequent releases. If new permissions are required, they will be listed in the release notes.

#### Steps

1. Create a custom role in Google Cloud:
  - a. Create a YAML file that includes the contents of the [service account permissions for the Connector](#).
  - b. From Google Cloud, activate cloud shell.
  - c. Upload the YAML file that includes the required permissions.
  - d. Create a custom role by using the `gcloud iam roles create` command.

The following example creates a role named "connector" at the project level:

```
gcloud iam roles create connector --project=myproject --file=connector.yaml
```

[Google Cloud docs: Creating and managing custom roles](#)

2. Create a service account in Google Cloud and assign the role to the service account:
  - a. From the IAM & Admin service, select **Service Accounts > Create Service Account**.
  - b. Enter service account details and select **Create and Continue**.
  - c. Select the role that you just created.
  - d. Finish the remaining steps to create the role.

[Google Cloud docs: Creating a service account](#)

3. If you plan to deploy Cloud Volumes ONTAP systems in different projects than the project where the Connector resides, then you'll need to provide the Connector's service account with access to those projects.

For example, let's say the Connector is in project 1 and you want to create Cloud Volumes ONTAP systems in project 2. You'll need to grant access to the service account in project 2.

- a. From the IAM & Admin service, select the Google Cloud project where you want to create Cloud Volumes ONTAP systems.
- b. On the **IAM** page, select **Grant Access** and provide the required details.
  - Enter the email of the Connector's service account.
  - Select the Connector's custom role.
  - Select **Save**.

For more details, refer to [Google Cloud documentation](#)

#### Result



The service account for the Connector VM is set up.

#### **Step 4: Set up shared VPC permissions**

If you are using a shared VPC to deploy resources into a service project, then you'll need to prepare your permissions.

This table is for reference and your environment should reflect the permissions table when IAM configuration is complete.

## View shared VPC permissions

Identity	Creator	Hosted in	Service project permissions	Host project permissions	Purpose
Google account to deploy the Connector	Custom	Service Project	<a href="#">Connector deployment policy</a>	compute.network User	Deploying the Connector in the service project
Connector service account	Custom	Service project	<a href="#">Connector service account policy</a>	compute.network User deploymentmanager.editor	Deploying and maintaining Cloud Volumes ONTAP and services in the service project
Cloud Volumes ONTAP service account	Custom	Service project	storage.admin member: BlueXP service account as serviceAccount.user	N/A	(Optional) For data tiering and BlueXP backup and recovery
Google APIs service agent	Google Cloud	Service project	(Default) Editor	compute.network User	Interacts with Google Cloud APIs on behalf of deployment. Allows BlueXP to use the shared network.
Google Compute Engine default service account	Google Cloud	Service project	(Default) Editor	compute.network User	Deploys Google Cloud instances and compute infrastructure on behalf of deployment. Allows BlueXP to use the shared network.

### Notes:

1. deploymentmanager.editor is only required at the host project if you are not passing firewall rules to the deployment and are choosing to let BlueXP create them for you. BlueXP will create a deployment in the host project which contains the VPC0 firewall rule if no rule is specified.
2. firewall.create and firewall.delete are only required if you are not passing firewall rules to the deployment and are choosing to let BlueXP create them for you. These permissions reside in the BlueXP account .yaml file. If you are deploying an HA pair using a shared VPC, these permissions will be used to create the firewall rules for VPC1, 2 and 3. For all other deployments, these permissions will also be used to create rules for VPC0.
3. For data tiering, the tiering service account must have the serviceAccount.user role on the service account, not just at the project level. Currently if you assign serviceAccount.user at the project level, the permissions don't show when you query the service account with getIAMPolicy.

## Step 5: Enable Google Cloud APIs

You must enable several Google Cloud APIs before deploying the Connector and Cloud Volumes ONTAP.

### Step

1. Enable the following Google Cloud APIs in your project:

- Cloud Deployment Manager V2 API
- Cloud Logging API
- Cloud Resource Manager API
- Compute Engine API
- Identity and Access Management (IAM) API
- Cloud Key Management Service (KMS) API

(Required only if you are planning to use BlueXP backup and recovery with customer-managed encryption keys (CMEK))

[Google Cloud documentation: Enabling APIs](#)

## Step 6: Create the Connector

Create a Connector directly from the BlueXP web-based console or by using gcloud.

### About this task

Creating the Connector deploys a virtual machine instance in Google Cloud using a default configuration. Do not change the Connector to a smaller VM instance with less CPU or RAM after creation. [Learn about the default configuration for the Connector.](#)

## BlueXP


### Before you begin

You should have the following:

- The required Google Cloud permissions to create the Connector and a service account for the Connector VM.
- A VPC and subnet that meets networking requirements.
- Details about a proxy server, if a proxy is required for internet access from the Connector.

### Steps

1. Select the **Connector** drop-down and select **Add Connector**.



2. Choose **Google Cloud Platform** as your cloud provider.
3. On the **Deploying a Connector** page, review the details about what you'll need. You have two options:
  - a. Select **Continue** to prepare for deployment by using the in-product guide. Each step in the in-product guide includes the information that's contained on this page of the documentation.
  - b. Select **Skip to Deployment** if you already prepared by following the steps on this page.
4. Follow the steps in the wizard to create the Connector:
  - If you're prompted, log in to your Google account, which should have the required permissions to create the virtual machine instance.

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

- **Details:** Enter a name for the virtual machine instance, specify tags, select a project, and then select the service account that has the required permissions (refer to the section above for details).
- **Location:** Specify a region, zone, VPC, and subnet for the instance.
- **Network:** Choose whether to enable a public IP address and optionally specify a proxy configuration.
- **Network tags:** Add a network tag to the Connector instance if using a transparent proxy. Network tags must start with a lowercase letter and can contain lowercase letters, numbers, and hyphens. Tags must end with a lowercase letter or number. For example, you might use the tag "connector-proxy".

- **Firewall Policy:** Choose whether to create a new firewall policy or whether to select an existing firewall policy that allows the required inbound and outbound rules.

#### [Firewall rules in Google Cloud](#)

- **Review:** Review your selections to verify that your set up is correct.

#### 5. Select **Add**.

The instance is ready in approximately 7 minutes; stay on the page until the process completes.

### Result

After the process completes, the Connector is available for use from BlueXP.

If you have Google Cloud Storage buckets in the same Google Cloud account where you created the Connector, you'll see a Google Cloud Storage working environment appear on the BlueXP canvas automatically. [Learn how to manage Google Cloud Storage from BlueXP](#)

### gcloud

#### Before you begin

You should have the following:

- The required Google Cloud permissions to create the Connector and a service account for the Connector VM.
- A VPC and subnet that meets networking requirements.
- An understanding of VM instance requirements.
  - **CPU:** 8 cores or 8 vCPUs
  - **RAM:** 32 GB
  - **Machine type:** We recommend n2-standard-8.

The Connector is supported in Google Cloud on a VM instance with an OS that supports Shielded VM features.

### Steps

1. Log in to the gcloud SDK using your preferred method.

In our examples, we'll use a local shell with the gcloud SDK installed, but you could use the native Google Cloud Shell in the Google Cloud console.

For more information about the Google Cloud SDK, visit the [Google Cloud SDK documentation page](#).

2. Verify that you are logged in as a user who has the required permissions that are defined in the section above:

```
gcloud auth list
```

The output should show the following where the \* user account is the desired user account to be logged in as:

## Credentialed Accounts

ACTIVE ACCOUNT

```
some_user_account@domain.com
```

```
* desired_user_account@domain.com
```

To set the active account, run:

```
$ gcloud config set account `ACCOUNT`
```

Updates are available for some Cloud SDK components. To install them,

please run:

```
$ gcloud components update
```

### 3. Run the `gcloud compute instances create` command:

```
gcloud compute instances create <instance-name>
  --machine-type=n2-standard-8
  --image-project=netapp-cloudmanager
  --image-family=cloudmanager
  --scopes=cloud-platform
  --project=<project>
  --service-account=<service-account>
  --zone=<zone>
  --no-address
  --tags <network-tag>
  --network <network-path>
  --subnet <subnet-path>
  --boot-disk-kms-key <kms-key-path>
```

#### **instance-name**

The desired instance name for the VM instance.

#### **project**

(Optional) The project where you want to deploy the VM.

#### **service-account**

The service account specified in the output from step 2.

#### **zone**

The zone where you want to deploy the VM

#### **no-address**

(Optional) No external IP address is used (you need a cloud NAT or proxy to route traffic to the public internet)

#### **network-tag**

(Optional) Add network tagging to link a firewall rule using tags to the Connector instance

**network-path**

(Optional) Add the name of the network to deploy the Connector into (for a Shared VPC, you need the full path)

**subnet-path**

(Optional) Add the name of the subnet to deploy the Connector into (for a Shared VPC, you need the full path)

**kms-key-path**

(Optional) Add a KMS key to encrypt the Connector's disks (IAM permissions also need to be applied)

For more information about these flags, visit the [Google Cloud compute SDK documentation](#).

Running the command deploys the Connector using the NetApp golden image. The Connector instance and software should be running in approximately five minutes.

4. Open a web browser from a host that has a connection to the Connector instance and enter the following URL:

`https://ipaddress`

5. After you log in, set up the Connector:
  - a. Specify the BlueXP organization to associate with the Connector.

[Learn about BlueXP identity and access management.](#)

- b. Enter a name for the system.

**Result**

The Connector is now installed and set up with your BlueXP organization.

Open a web browser and go to the [BlueXP console](#) to start using the Connector with BlueXP.

**Manually install the Connector in Google Cloud**

You can manually install the Connector on a Linux host running in Google Cloud. To manually install the Connector on your own Linux host, you need to review host requirements, set up your networking, prepare Google Cloud permissions, enable APIs, install the Connector, and then provide the permissions that you prepared.

**Before you begin**

- You should have an [understanding of Connectors](#).
- You should review [Connector limitations](#).

**Step 1: Review host requirements**

The Connector software must run on a host that meets specific operating system requirements, RAM requirements, port requirements, and so on.



The Connector reserves the UID and GID range of 19000 to 19200. This range is fixed and cannot be modified. If any third-party software on your host is using UIDs or GIDs within this range, the Connector installation will fail. NetApp recommends using a host that is free of third-party software to avoid conflicts.

### Dedicated host

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

Host can be of any architecture that meets the following size requirements:

- CPU: 8 cores or 8 vCPUs
- RAM: 32 GB

### Hypervisor

A bare metal or hosted hypervisor that is certified to run a supported operating system is required.

### Operating system and container requirements

BlueXP supports the Connector with the following operating systems when using BlueXP in standard mode or restricted mode. A container orchestration tool is required before you install the Connector.

Operating system	Supported OS versions	Supported Connector versions	Required container tool	SELinux
Red Hat Enterprise Linux	9.1 to 9.4 8.6 to 8.10 7.9	3.9.40 or later with BlueXP in standard mode or restricted mode	Podman version 4.6.1 or 4.9.4 <a href="#">View Podman configuration requirements.</a>	Supported in enforcing mode or permissive mode <sup>1</sup>
Ubuntu	24.04 LTS	3.9.45 or later with BlueXP in standard mode or restricted mode	Docker Engine 23.06 to 28.0.0.	Not supported
	22.04 LTS	3.9.29 or later	Docker Engine 23.0.6 to 28.0.0.	Not supported

Notes:

1. Management of Cloud Volumes ONTAP systems is not supported by Connectors that have SELinux enabled on the operating system.
2. The Connector is supported on English-language versions of these operating systems.
3. For RHEL, the host must be registered with Red Hat Subscription Management. If it's not registered, the host can't access repositories to update required 3rd-party software during Connector installation.

### Google Cloud machine type

An instance type that meets the CPU and RAM requirements above. We recommend n2-standard-8.



The Connector is supported in Google Cloud on a VM instance with an OS that supports [Shielded VM features](#)

### Disk space in /opt

100 GiB of space must be available

BlueXP uses /opt to install the /opt/application/netapp directory and its contents.

### Disk space in /var

20 GiB of space must be available

BlueXP requires this space in /var because Docker or Podman are architected to create the containers within this directory. Specifically, they will create containers in the /var/lib/containers/storage directory. External mounts or symlinks do not work for this space.

### Step 2: Install Podman or Docker Engine

Depending on your operating system, either Podman or Docker Engine is required before you install the Connector.

- Podman is required for Red Hat Enterprise Linux 8 and 9.

[View the Podman versions that BlueXP supports.](#)

- Docker Engine is required for Ubuntu.

[View the Docker Engine versions that BlueXP supports.](#)

### Example 3. Steps

#### Podman

Follow these steps to install Podman and configure it to meet the following requirements:

- Enable and start the podman.socket service
- Install python3
- Install the podman-compose package version 1.0.6
- Add podman-compose to the PATH environment variable



When using Podman, adjust the aardvark-dns service port (default: 53) after installing the Connector to avoid conflicts with the DNS port on the host. Follow the instructions to configure the port.

#### Steps

1. Remove the podman-docker package if it's installed on the host.

```
dnf remove podman-docker
rm /var/run/docker.sock
```

2. Install Podman.

Podman is available from official Red Hat Enterprise Linux repositories.

For Red Hat Enterprise Linux 9:

```
sudo dnf install podman-2:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

For Red Hat Enterprise Linux 8:

```
sudo dnf install podman-3:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

3. Enable and start the podman.socket service.

```
sudo systemctl enable --now podman.socket
```

4. Install python3.

```
sudo dnf install python3
```

5. Install the EPEL repository package if it's not already available on your system.

This step is required because podman-compose is available from the Extra Packages for Enterprise Linux (EPEL) repository.

For Red Hat Enterprise Linux 9:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-9.noarch.rpm
```

For Red Hat Enterprise Linux 8:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

6. Install podman-compose package 1.0.6.

```
sudo dnf install podman-compose-1.0.6
```



Using the `dnf install` command meets the requirement for adding podman-compose to the PATH environment variable. The installation command adds podman-compose to `/usr/bin`, which is already included in the `secure_path` option on the host.

### Docker Engine

Follow the documentation from Docker to install Docker Engine.

#### Steps

1. [View installation instructions from Docker](#)

Be sure to follow the steps to install a specific version of Docker Engine. Installing the latest version will install a version of Docker that BlueXP doesn't support.

2. Verify that Docker is enabled and running.

```
sudo systemctl enable docker && sudo systemctl start docker
```

### Step 3: Set up networking

Set up your networking so the Connector can manage resources and processes within your hybrid cloud environment. For example, you need to ensure that connections are available to target networks and that

outbound internet access is available.

## Connections to target networks

A Connector requires a network connection to the location where you're planning to create and manage working environments. For example, the network where you plan to create Cloud Volumes ONTAP systems or a storage system in your on-premises environment.

## Outbound internet access

The network location where you deploy the Connector must have an outbound internet connection to contact specific endpoints.

## Endpoints contacted from computers when using the BlueXP web-based console

Computers that access the BlueXP console from a web browser must have the ability to contact several endpoints. You'll need to use the BlueXP console to set up the Connector and for day-to-day use of BlueXP.

[Prepare networking for the BlueXP console.](#)

## Endpoints contacted during manual installation

When you manually install the Connector on your own Linux host, the installer for the Connector requires access to the following URLs during the installation process:

- <https://mysupport.netapp.com>
- <https://signin.b2c.netapp.com> (this endpoint is the CNAME URL for <https://mysupport.netapp.com>)
- <https://cloudmanager.cloud.netapp.com/tenancy>
- <https://stream.cloudmanager.cloud.netapp.com>
- <https://production-artifacts.cloudmanager.cloud.netapp.com>
- To obtain images, the installer needs access to one of these two sets of endpoints:
  - Option 1 (recommended):
    - <https://bluexpinfraprod.eastus2.data.azurecr.io>
    - <https://bluexpinfraprod.azurecr.io>
  - Option 2:
    - [https://\\*.blob.core.windows.net](https://*.blob.core.windows.net)
    - <https://cloudmanagerinfraprod.azurecr.io>

The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

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

## Endpoints contacted from the Connector

The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment for day-to-day operations.

Note that the endpoints listed below are all CNAME entries.

Endpoints	Purpose
<a href="https://www.googleapis.com/compute/v1/">https://www.googleapis.com/compute/v1/</a> <a href="https://compute.googleapis.com/compute/v1">https://compute.googleapis.com/compute/v1</a> <a href="https://cloudresourcemanager.googleapis.com/v1/projects">https://cloudresourcemanager.googleapis.com/v1/projects</a> <a href="https://www.googleapis.com/compute/beta">https://www.googleapis.com/compute/beta</a> <a href="https://storage.googleapis.com/storage/v1">https://storage.googleapis.com/storage/v1</a> <a href="https://www.googleapis.com/storage/v1">https://www.googleapis.com/storage/v1</a> <a href="https://iam.googleapis.com/v1">https://iam.googleapis.com/v1</a> <a href="https://cloudkms.googleapis.com/v1">https://cloudkms.googleapis.com/v1</a> <a href="https://www.googleapis.com/deploymentmanager/v2/projects">https://www.googleapis.com/deploymentmanager/v2/projects</a>	To manage resources in Google Cloud.
<a href="https://support.netapp.com">https://support.netapp.com</a> <a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	To obtain licensing information and to send AutoSupport messages to NetApp support.
<a href="https://*.api.bluexp.netapp.com">https://*.api.bluexp.netapp.com</a> <a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a> <a href="https://*.cloudmanager.cloud.netapp.com">https://*.cloudmanager.cloud.netapp.com</a> <a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	To provide SaaS features and services within BlueXP.
Choose between two sets of endpoints: <ul style="list-style-type: none"><li>Option 1 (recommended) <sup>1</sup> <a href="https://bluexpinfraprod.eastus2.data.azurecr.io">https://bluexpinfraprod.eastus2.data.azurecr.io</a> <a href="https://bluexpinfraprod.azurecr.io">https://bluexpinfraprod.azurecr.io</a></li><li>Option 2 <a href="https://*.blob.core.windows.net">https://*.blob.core.windows.net</a> <a href="https://cloudmanagerinfraprod.azurecr.io">https://cloudmanagerinfraprod.azurecr.io</a></li></ul>	To obtain images for Connector upgrades.

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while

allowing the endpoints listed in option 2.

## Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

## Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

## Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

## Step 4: Set up permissions for the Connector

A Google Cloud service account is required to provide the Connector with the permissions that BlueXP needs to manage resources in Google Cloud. When you create the Connector, you'll need to associate this service account with the Connector VM.

It's your responsibility to update the custom role as new permissions are added in subsequent releases. If new permissions are required, they will be listed in the release notes.

### Steps

1. Create a custom role in Google Cloud:
  - a. Create a YAML file that includes the contents of the [service account permissions for the Connector](#).
  - b. From Google Cloud, activate cloud shell.
  - c. Upload the YAML file that includes the required permissions.
  - d. Create a custom role by using the `gcloud iam roles create` command.

The following example creates a role named "connector" at the project level:

```
gcloud iam roles create connector --project=myproject --file=connector.yaml
```

[Google Cloud docs: Creating and managing custom roles](#)

2. Create a service account in Google Cloud and assign the role to the service account:
  - a. From the IAM & Admin service, select **Service Accounts > Create Service Account**.
  - b. Enter service account details and select **Create and Continue**.
  - c. Select the role that you just created.
  - d. Finish the remaining steps to create the role.

[Google Cloud docs: Creating a service account](#)

3. If you plan to deploy Cloud Volumes ONTAP systems in different projects than the project where the Connector resides, then you'll need to provide the Connector's service account with access to those projects.

For example, let's say the Connector is in project 1 and you want to create Cloud Volumes ONTAP systems in project 2. You'll need to grant access to the service account in project 2.

- a. From the IAM & Admin service, select the Google Cloud project where you want to create Cloud Volumes ONTAP systems.
- b. On the **IAM** page, select **Grant Access** and provide the required details.
  - Enter the email of the Connector's service account.
  - Select the Connector's custom role.
  - Select **Save**.

For more details, refer to [Google Cloud documentation](#)

## Result

The service account for the Connector VM is set up.

## Step 5: Set up shared VPC permissions

If you are using a shared VPC to deploy resources into a service project, then you'll need to prepare your permissions.

This table is for reference and your environment should reflect the permissions table when IAM configuration is complete.

## View shared VPC permissions

Identity	Creator	Hosted in	Service project permissions	Host project permissions	Purpose
Google account to deploy the Connector	Custom	Service Project	<a href="#">Connector deployment policy</a>	compute.network User	Deploying the Connector in the service project
Connector service account	Custom	Service project	<a href="#">Connector service account policy</a>	compute.network User deploymentmanager.editor	Deploying and maintaining Cloud Volumes ONTAP and services in the service project
Cloud Volumes ONTAP service account	Custom	Service project	storage.admin member: BlueXP service account as serviceAccount.user	N/A	(Optional) For data tiering and BlueXP backup and recovery
Google APIs service agent	Google Cloud	Service project	(Default) Editor	compute.network User	Interacts with Google Cloud APIs on behalf of deployment. Allows BlueXP to use the shared network.
Google Compute Engine default service account	Google Cloud	Service project	(Default) Editor	compute.network User	Deploys Google Cloud instances and compute infrastructure on behalf of deployment. Allows BlueXP to use the shared network.

### Notes:

1. deploymentmanager.editor is only required at the host project if you are not passing firewall rules to the deployment and are choosing to let BlueXP create them for you. BlueXP will create a deployment in the host project which contains the VPC0 firewall rule if no rule is specified.
2. firewall.create and firewall.delete are only required if you are not passing firewall rules to the deployment and are choosing to let BlueXP create them for you. These permissions reside in the BlueXP account .yaml file. If you are deploying an HA pair using a shared VPC, these permissions will be used to create the firewall rules for VPC1, 2 and 3. For all other deployments, these permissions will also be used to create rules for VPC0.
3. For data tiering, the tiering service account must have the serviceAccount.user role on the service account, not just at the project level. Currently if you assign serviceAccount.user at the project level, the permissions don't show when you query the service account with getIAMPolicy.



## Step 6: Enable Google Cloud APIs

Several Google Cloud APIs must be enabled before you can deploy Cloud Volumes ONTAP systems in Google Cloud.

### Step

1. Enable the following Google Cloud APIs in your project:

- Cloud Deployment Manager V2 API
- Cloud Logging API
- Cloud Resource Manager API
- Compute Engine API
- Identity and Access Management (IAM) API
- Cloud Key Management Service (KMS) API

(Required only if you are planning to use BlueXP backup and recovery with customer-managed encryption keys (CMEK))

[Google Cloud documentation: Enabling APIs](#)

## Step 7: Install the Connector

After the pre-requisites are complete, you can manually install the software on your own Linux host.

### Before you begin

You should have the following:

- Root privileges to install the Connector.
- Details about a proxy server, if a proxy is required for internet access from the Connector.

You have the option to configure a proxy server after installation but doing so requires restarting the Connector.

- A CA-signed certificate, if the proxy server uses HTTPS or if the proxy is an intercepting proxy.



You cannot set a certificate for a transparent proxy server when manually installing the Connector. If you need to set a certificate for a transparent proxy server, you must use the Maintenance Console after installation. Learn more about the [Connector Maintenance Console](#).

### About this task

The installer that is available on the NetApp Support Site might be an earlier version. After installation, the Connector automatically updates itself if a new version is available.

### Steps

1. If the `http_proxy` or `https_proxy` system variables are set on the host, remove them:

```
unset http_proxy
unset https_proxy
```

If you don't remove these system variables, the installation will fail.

2. Download the Connector software from the [NetApp Support Site](#), and then copy it to the Linux host.

You should download the "online" Connector installer that's meant for use in your network or in the cloud. A separate "offline" installer is available for the Connector, but it's only supported with private mode deployments.

3. Assign permissions to run the script.

```
chmod +x BlueXP-Connector-Cloud-<version>
```

Where <version> is the version of the Connector that you downloaded.

4. Run the installation script.

```
./BlueXP-Connector-Cloud-<version> --proxy <HTTP or HTTPS proxy server>  
--cacert <path and file name of a CA-signed certificate>
```

You'll need to add proxy information if your network requires a proxy for internet access. You can add either a transparent or explicit proxy. The `--proxy` and `--cacert` parameters are optional and you won't be prompted to add them. If you have a proxy server, you will need to enter the parameters as shown.

Here is an example configuring an explicit proxy server with a CA-signed certificate:

```
./BlueXP-Connector-Cloud-v3.9.40--proxy  
https://user:password@10.0.0.30:8080/ --cacert  
/tmp/cacert/certificate.cer
```

`--proxy` configures the Connector to use an HTTP or HTTPS proxy server using one of the following formats:

- `http://address:port`
- `http://user-name:password@address:port`
- `http://domain-name%92user-name:password@address:port`
- `https://address:port`
- `https://user-name:password@address:port`
- `https://domain-name%92user-name:password@address:port`

Note the following:

- The user can be a local user or domain user.
- For a domain user, you must use the ASCII code for a \ as shown above.
- BlueXP doesn't support user names or passwords that include the @ character.
- If the password includes any of the following special characters, you must escape that special character by prepending it with a backslash: & or !

For example:

```
http://bxpproxyuser:netapp1!@address:3128
```

--cacert specifies a CA-signed certificate to use for HTTPS access between the Connector and the proxy server. This parameter is required for HTTPS proxy servers, intercepting proxy servers, and transparent proxy servers.

Here is an example configuring a transparent proxy server. When you configure a transparent proxy, you don't need to define the proxy server. You only add a CA-signed certificate to your Connector host:

```
./BlueXP-Connector-Cloud-v3.9.40 --cacert /tmp/cacert/certificate.cer
```

5. If you used Podman, you'll need to adjust the aardvark-dns port.
  - a. SSH to the BlueXP Connector virtual machine.
  - b. Open podman `/usr/share/containers/containers.conf` file and modify the chosen port for Aardvark DNS service. For example, change it to 54.

```
vi /usr/share/containers/containers.conf
...
# Port to use for dns forwarding daemon with netavark in rootful
bridge
# mode and dns enabled.
# Using an alternate port might be useful if other DNS services
should
# run on the machine.
#
dns_bind_port = 54
...
Esc:wq
```

- c. Reboot the Connector virtual machine.
6. Wait for the installation to complete.

At the end of the installation, the Connector service (occm) restarts twice if you specified a proxy server.

7. Open a web browser from a host that has a connection to the Connector virtual machine and enter the following URL:

```
https://ipaddress
```

8. After you log in, set up the Connector:
  - a. Specify the BlueXP organization to associate with the Connector.
  - b. Enter a name for the system.
  - c. Under **Are you running in a secured environment?** keep restricted mode disabled.

You should keep restricted mode disabled because these steps describe how to use BlueXP in

standard mode. You should enable restricted mode only if you have a secure environment and want to disconnect this account from BlueXP backend services. If that's the case, [follow steps to get started with BlueXP in restricted mode](#).

d. Select **Let's start**.

If you have Google Cloud Storage buckets in the same Google Cloud account where you created the Connector, you'll see a Google Cloud Storage working environment appear on the BlueXP canvas automatically. [Learn how to manage Google Cloud Storage from BlueXP](#)

## Step 8: Provide permissions to BlueXP

You need to provide BlueXP with the Google Cloud permissions that you previously set up. Providing the permissions enables BlueXP to manage your data and storage infrastructure in Google Cloud.

### Steps

1. Go to the Google Cloud portal and assign the service account to the Connector VM instance.

[Google Cloud documentation: Changing the service account and access scopes for an instance](#)

2. If you want to manage resources in other Google Cloud projects, grant access by adding the service account with the BlueXP role to that project. You'll need to repeat this step for each project.

### Result

BlueXP now has the permissions that it needs to perform actions in Google Cloud on your behalf.

## Install and set up a Connector on-premises

You can install a Connector on one of your on-premises machines. To run the Connector on-premises, you need to review host requirements, set up your networking, prepare cloud permissions, install the Connector, set up the Connector, and then provide the permissions that you prepared.

### Before you begin

- Review information about [Connectors](#).
- You should review [Connector limitations](#).

### Step 1: Review host requirements

Run the Connector software on a host that meets operating system, RAM, and port requirements. Ensure that your host meets these requirements before you install the Connector.



The Connector reserves the UID and GID range of 19000 to 19200. This range is fixed and cannot be modified. If any third-party software on your host is using UIDs or GIDs within this range, the Connector installation will fail. NetApp recommends using a host that is free of third-party software to avoid conflicts.

### Dedicated host

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

Host can be of any architecture that meets the following size requirements:

- CPU: 8 cores or 8 vCPUs
- RAM: 32 GB

## Hypervisor

A bare metal or hosted hypervisor that is certified to run a supported operating system is required.

## Operating system and container requirements

BlueXP supports the Connector with the following operating systems when using BlueXP in standard mode or restricted mode. A container orchestration tool is required before you install the Connector.

Operating system	Supported OS versions	Supported Connector versions	Required container tool	SELinux
Red Hat Enterprise Linux	9.1 to 9.4 8.6 to 8.10 7.9	3.9.40 or later with BlueXP in standard mode or restricted mode	Podman version 4.6.1 or 4.9.4  <a href="#">View Podman configuration requirements.</a>	Supported in enforcing mode or permissive mode <sup>1</sup>
Ubuntu	24.04 LTS	3.9.45 or later with BlueXP in standard mode or restricted mode	Docker Engine 23.06 to 28.0.0.	Not supported
	22.04 LTS	3.9.29 or later	Docker Engine 23.0.6 to 28.0.0.	Not supported

Notes:

1. Management of Cloud Volumes ONTAP systems is not supported by Connectors that have SELinux enabled on the operating system.
2. The Connector is supported on English-language versions of these operating systems.
3. For RHEL, the host must be registered with Red Hat Subscription Management. If it's not registered, the host can't access repositories to update required 3rd-party software during Connector installation.

## Disk space in /opt

100 GiB of space must be available

BlueXP uses /opt to install the /opt/application/netapp directory and its contents.

## Disk space in /var

20 GiB of space must be available

BlueXP requires this space in /var because Docker or Podman are architected to create the containers within this directory. Specifically, they will create containers in the /var/lib/containers/storage directory. External mounts or symlinks do not work for this space.

## Step 2: Install Podman or Docker Engine

Depending on your operating system, either Podman or Docker Engine is required before you install the Connector.

- Podman is required for Red Hat Enterprise Linux 8 and 9.

[View the Podman versions that BlueXP supports.](#)

- Docker Engine is required for Ubuntu.

[View the Docker Engine versions that BlueXP supports.](#)

## Example 4. Steps

### Podman

Follow these steps to install Podman and configure it to meet the following requirements:

- Enable and start the podman.socket service
- Install python3
- Install the podman-compose package version 1.0.6
- Add podman-compose to the PATH environment variable



When using Podman, adjust the aardvark-dns service port (default: 53) after installing the Connector to avoid conflicts with the DNS port on the host. Follow the instructions to configure the port.

### Steps

1. Remove the podman-docker package if it's installed on the host.

```
dnf remove podman-docker
rm /var/run/docker.sock
```

2. Install Podman.

Podman is available from official Red Hat Enterprise Linux repositories.

For Red Hat Enterprise Linux 9:

```
sudo dnf install podman-2:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

For Red Hat Enterprise Linux 8:

```
sudo dnf install podman-3:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

3. Enable and start the podman.socket service.

```
sudo systemctl enable --now podman.socket
```

4. Install python3.

```
sudo dnf install python3
```

5. Install the EPEL repository package if it's not already available on your system.

This step is required because podman-compose is available from the Extra Packages for Enterprise Linux (EPEL) repository.

For Red Hat Enterprise Linux 9:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-9.noarch.rpm
```

For Red Hat Enterprise Linux 8:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

6. Install podman-compose package 1.0.6.

```
sudo dnf install podman-compose-1.0.6
```



Using the `dnf install` command meets the requirement for adding podman-compose to the PATH environment variable. The installation command adds podman-compose to `/usr/bin`, which is already included in the `secure_path` option on the host.

## Docker Engine

Follow the documentation from Docker to install Docker Engine.

### Steps

1. [View installation instructions from Docker](#)

Be sure to follow the steps to install a specific version of Docker Engine. Installing the latest version will install a version of Docker that BlueXP doesn't support.

2. Verify that Docker is enabled and running.

```
sudo systemctl enable docker && sudo systemctl start docker
```

## Step 3: Set up networking

Set up networking to ensure the Connector can manage resources, with connections to target networks and outbound internet access.



## Connections to target networks

A Connector requires a network connection to the location where you're planning to create and manage working environments. For example, the network where you plan to create Cloud Volumes ONTAP systems or a storage system in your on-premises environment.

## Outbound internet access

The network location where you deploy the Connector must have an outbound internet connection to contact specific endpoints.

## Endpoints contacted from computers when using the BlueXP web-based console

Computers that access the BlueXP console from a web browser must have the ability to contact several endpoints. You'll need to use the BlueXP console to set up the Connector and for day-to-day use of BlueXP.

[Prepare networking for the BlueXP console.](#)

## Endpoints contacted during manual installation

When you manually install the Connector on your own Linux host, the installer for the Connector requires access to the following URLs during the installation process:

- <https://mysupport.netapp.com>
- <https://signin.b2c.netapp.com> (this endpoint is the CNAME URL for <https://mysupport.netapp.com>)
- <https://cloudmanager.cloud.netapp.com/tenancy>
- <https://stream.cloudmanager.cloud.netapp.com>
- <https://production-artifacts.cloudmanager.cloud.netapp.com>
- To obtain images, the installer needs access to one of these two sets of endpoints:
  - Option 1 (recommended):
    - <https://bluexpinfraprod.eastus2.data.azurecr.io>
    - <https://bluexpinfraprod.azurecr.io>
  - Option 2:
    - [https://\\*.blob.core.windows.net](https://*.blob.core.windows.net)
    - <https://cloudmanagerinfraprod.azurecr.io>

The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

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

## Endpoints contacted from the Connector

The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment for day-to-day operations.

Note that the endpoints listed below are all CNAME entries.

Endpoints	Purpose
AWS services (amazonaws.com): <ul style="list-style-type: none"><li>• CloudFormation</li><li>• Elastic Compute Cloud (EC2)</li><li>• Identity and Access Management (IAM)</li><li>• Key Management Service (KMS)</li><li>• Security Token Service (STS)</li><li>• Simple Storage Service (S3)</li></ul>	To manage resources in AWS. The exact endpoint depends on the AWS region that you're using. <a href="#">Refer to AWS documentation for details</a>
https://management.azure.com https://login.microsoftonline.com https://blob.core.windows.net https://core.windows.net	To manage resources in Azure public regions.
https://management.chinacloudapi.cn https://login.chinacloudapi.cn https://blob.core.chinacloudapi.cn https://core.chinacloudapi.cn	To manage resources in Azure China regions.
https://www.googleapis.com/compute/v1/ https://compute.googleapis.com/compute/v1 https://cloudresourcemanager.googleapis.com/v1/projects https://www.googleapis.com/compute/beta https://storage.googleapis.com/storage/v1 https://www.googleapis.com/storage/v1 https://iam.googleapis.com/v1 https://cloudkms.googleapis.com/v1 https://www.googleapis.com/deploymentmanager/v2/projects	To manage resources in Google Cloud.
https://support.netapp.com https://mysupport.netapp.com	To obtain licensing information and to send AutoSupport messages to NetApp support.
https://*.api.bluexp.netapp.com https://api.bluexp.netapp.com https://*.cloudmanager.cloud.netapp.com https://cloudmanager.cloud.netapp.com https://netapp-cloud-account.auth0.com	To provide SaaS features and services within BlueXP.

Endpoints	Purpose
<p>Choose between two sets of endpoints:</p> <ul style="list-style-type: none"> <li>• Option 1 (recommended) <sup>1</sup></li> </ul> <p>https://bluexpinfraprod.eastus2.data.azurecr.io https://bluexpinfraprod.azurecr.io</p> <ul style="list-style-type: none"> <li>• Option 2</li> </ul> <p>https://*.blob.core.windows.net https://cloudmanagerinfraprod.azurecr.io</p>	<p>To obtain images for Connector upgrades.</p>

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

### Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

### Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

## Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

### Step 4: Set up cloud permissions

If you want to use BlueXP services in AWS or Azure with an on-premises Connector, then you need to set up permissions in your cloud provider so that you can add the credentials to the Connector after you install it.



Why not Google Cloud? When the Connector is installed on your premises, it can't manage your resources in Google Cloud. You must install the Connector in Google Cloud to manage any resources that reside there.

## AWS

When the Connector is installed on-premises, you need to provide BlueXP with AWS permissions by adding access keys for an IAM user who has the required permissions.

You must use this authentication method if the Connector is installed on-premises. You can't use an IAM role.

### Steps

1. Log in to the AWS console and navigate to the IAM service.
2. Create a policy:
  - a. Select **Policies > Create policy**.
  - b. Select **JSON** and copy and paste the contents of the [IAM policy for the Connector](#).
  - c. Finish the remaining steps to create the policy.

Depending on the BlueXP services that you're planning to use, you might need to create a second policy.

For standard regions, the permissions are spread across two policies. Two policies are required due to a maximum character size limit for managed policies in AWS. [Learn more about IAM policies for the Connector](#).

3. Attach the policies to an IAM user.
  - [AWS Documentation: Creating IAM Roles](#)
  - [AWS Documentation: Adding and Removing IAM Policies](#)
4. Ensure that the user has an access key that you can add to BlueXP after you install the Connector.

### Result

You should now have access keys for an IAM user who has the required permissions. After you install the Connector, associate these credentials with the Connector from BlueXP.

## Azure

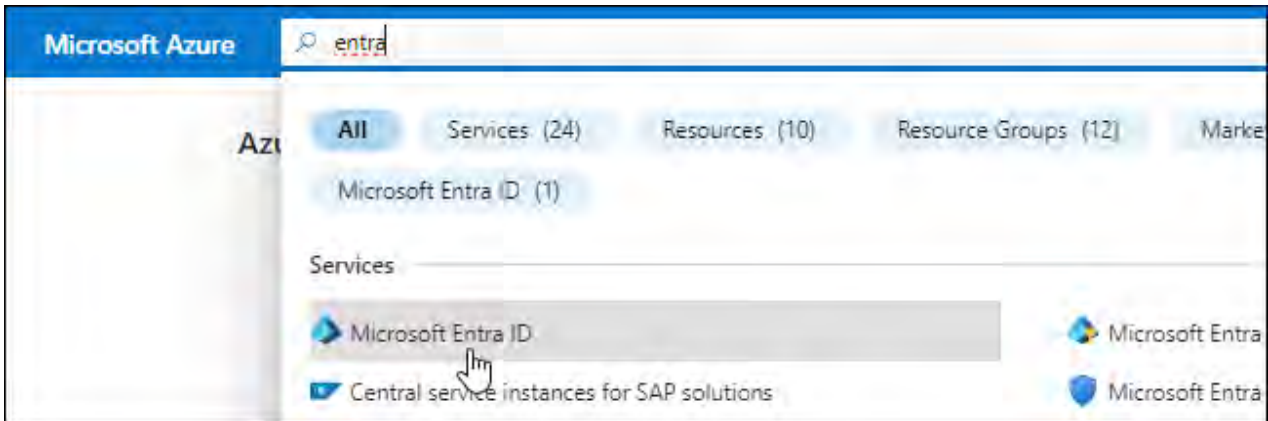
When the Connector is installed on-premises, you need to provide BlueXP with Azure permissions by setting up a service principal in Microsoft Entra ID and obtaining the Azure credentials that BlueXP needs.

### Create a Microsoft Entra application for role-based access control

1. Ensure that you have permissions in Azure to create an Active Directory application and to assign the application to a role.

For details, refer to [Microsoft Azure Documentation: Required permissions](#)

2. From the Azure portal, open the **Microsoft Entra ID** service.



3. In the menu, select **App registrations**.
4. Select **New registration**.
5. Specify details about the application:
  - **Name**: Enter a name for the application.
  - **Account type**: Select an account type (any will work with BlueXP).
  - **Redirect URI**: You can leave this field blank.
6. Select **Register**.

You've created the AD application and service principal.

### Assign the application to a role

1. Create a custom role:

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would prefer to use a different method, refer to [Azure documentation](#)

- a. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
- b. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription from which users will create Cloud Volumes ONTAP systems.

### Example


```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- Start [Azure Cloud Shell](#) and choose the Bash environment.

- Upload the JSON file.

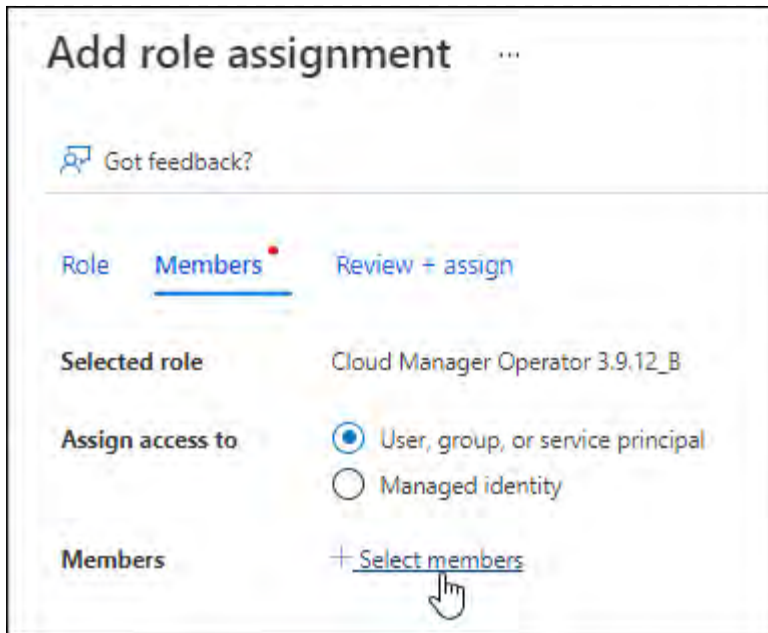


- Use the Azure CLI to create the custom role:

```
az role definition create --role-definition  
Connector_Policy.json
```


You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.

2. Assign the application to the role:
  - a. From the Azure portal, open the **Subscriptions** service.
  - b. Select the subscription.
  - c. Select **Access control (IAM) > Add > Add role assignment**.
  - d. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.
  - e. In the **Members** tab, complete the following steps:
    - Keep **User, group, or service principal** selected.
    - Select **Select members**.



- Search for the name of the application.

Here's an example:



- Select the application and select **Select**.
  - Select **Next**.
- f. Select **Review + assign**.

The service principal now has the required Azure permissions to deploy the Connector.

If you want to deploy Cloud Volumes ONTAP from multiple Azure subscriptions, then you must bind the service principal to each of those subscriptions. BlueXP enables you to select the subscription that you want to use when deploying Cloud Volumes ONTAP.

#### Add Windows Azure Service Management API permissions

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Select **API permissions > Add a permission**.
3. Under **Microsoft APIs**, select **Azure Service Management**.
















## Request API permissions

Select an API

Microsoft APIs APIs my organization uses My APIs

### Commonly used Microsoft APIs

<b>Microsoft Graph</b> Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10. Access Azure AD, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint. 		
 <b>Azure Batch</b> Schedule large-scale parallel and HPC applications in the cloud	 <b>Azure Data Catalog</b> Programmatic access to Data Catalog resources to register, annotate and search data assets	 <b>Azure Data Explorer</b> Perform ad-hoc queries on terabytes of data to build near real-time and complex analytics solutions
 <b>Azure Data Lake</b> Access to storage and compute for big data analytic scenarios	 <b>Azure DevOps</b> Integrate with Azure DevOps and Azure DevOps server	 <b>Azure Import/Export</b> Programmatic control of import/export jobs
 <b>Azure Key Vault</b> Manage your key vaults as well as the keys, secrets, and certificates within your Key Vaults	 <b>Azure Rights Management Services</b> Allow validated users to read and write protected content	 <b>Azure Service Management</b> Programmatic access to much of the functionality available through the Azure portal
 <b>Azure Storage</b> Secure, massively scalable object and data lake storage for unstructured and semi-structured data	 <b>Customer Insights</b> Create profile and interaction models for your products	 <b>Data Export Service for Microsoft Dynamics 365</b> Export data from Microsoft Dynamics CRM organization to an external destination

4. Select **Access Azure Service Management as organization users** and then select **Add permissions**.

## Request API permissions

< All APIs



Azure Service Management

<https://management.azure.com/> [Docs](#)

What type of permissions does your application require?

### Delegated permissions

Your application needs to access the API as the signed-in user.

### Application permissions

Your application runs as a background service or daemon without a signed-in user.

Select permissions

[expand all](#)

Type to search

PERMISSION

ADMIN CONSENT REQUIRED

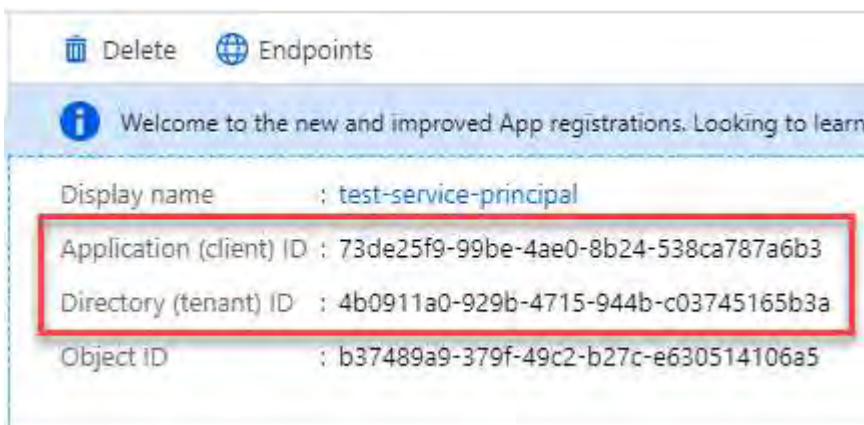


user\_impersonation

Access Azure Service Management as organization users (preview)

## Get the application ID and directory ID for the application

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Copy the **Application (client) ID** and the **Directory (tenant) ID**.



When you add the Azure account to BlueXP, you need to provide the application (client) ID and the directory (tenant) ID for the application. BlueXP uses the IDs to programmatically sign in.

## Create a client secret

1. Open the **Microsoft Entra ID** service.
2. Select **App registrations** and select your application.
3. Select **Certificates & secrets > New client secret**.
4. Provide a description of the secret and a duration.
5. Select **Add**.
6. Copy the value of the client secret.

## Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

[+ New client secret](#)

DESCRIPTION	EXPIRES	VALUE	Copy to clipboard
test secret	8/16/2020	*sZ1jSe2By:D*-ZRov4NLfdAcY7:+0vA	

You now have a client secret that BlueXP can use it to authenticate with Microsoft Entra ID.

### Step 5: Install the Connector

Download and install the Connector software on an existing Linux host on-premises.

#### Before you begin

You should have the following:

- Root privileges to install the Connector.
- Details about a proxy server, if a proxy is required for internet access from the Connector.

You have the option to configure a proxy server after installation but doing so requires restarting the Connector.

- A CA-signed certificate, if the proxy server uses HTTPS or if the proxy is an intercepting proxy.



You cannot set a certificate for a transparent proxy server when manually installing the Connector. If you need to set a certificate for a transparent proxy server, you must use the Maintenance Console after installation. Learn more about the [Connector Maintenance Console](#).

#### About this task

The installer that is available on the NetApp Support Site might be an earlier version. After installation, the Connector automatically updates itself if a new version is available.

#### Steps

1. If the `http_proxy` or `https_proxy` system variables are set on the host, remove them:

```
unset http_proxy
unset https_proxy
```

If you don't remove these system variables, the installation will fail.

2. Download the Connector software from the [NetApp Support Site](#), and then copy it to the Linux host.

You should download the "online" Connector installer that's meant for use in your network or in the cloud. A separate "offline" installer is available for the Connector, but it's only supported with private mode deployments.

3. Assign permissions to run the script.

```
chmod +x BlueXP-Connector-Cloud-<version>
```

Where <version> is the version of the Connector that you downloaded.

#### 4. Run the installation script.

```
./BlueXP-Connector-Cloud-<version> --proxy <HTTP or HTTPS proxy server>  
--cacert <path and file name of a CA-signed certificate>
```

You'll need to add proxy information if your network requires a proxy for internet access. You can add either a transparent or explicit proxy. The `--proxy` and `--cacert` parameters are optional and you won't be prompted to add them. If you have a proxy server, you will need to enter the parameters as shown.

Here is an example configuring an explicit proxy server with a CA-signed certificate:

```
./BlueXP-Connector-Cloud-v3.9.40--proxy  
https://user:password@10.0.0.30:8080/ --cacert  
/tmp/cacert/certificate.cer
```

`--proxy` configures the Connector to use an HTTP or HTTPS proxy server using one of the following formats:

- `http://address:port`
- `http://user-name:password@address:port`
- `http://domain-name%92user-name:password@address:port`
- `https://address:port`
- `https://user-name:password@address:port`
- `https://domain-name%92user-name:password@address:port`

Note the following:

- The user can be a local user or domain user.
- For a domain user, you must use the ASCII code for a \ as shown above.
- BlueXP doesn't support user names or passwords that include the @ character.
- If the password includes any of the following special characters, you must escape that special character by prepending it with a backslash: & or !

For example:

```
http://bxpproxyuser:netapp1!\@address:3128
```

`--cacert` specifies a CA-signed certificate to use for HTTPS access between the Connector and the proxy server. This parameter is required for HTTPS proxy servers, intercepting proxy servers, and transparent proxy servers.

Here is an example configuring a transparent proxy server. When you configure a transparent proxy, you don't need to define the proxy server. You only add a CA-signed certificate to your Connector host:

```
./BlueXP-Connector-Cloud-v3.9.40 --cacert /tmp/cacert/certificate.cer
```

5. If you used Podman, you'll need to adjust the aardvark-dns port.
  - a. SSH to the BlueXP Connector virtual machine.
  - b. Open podman `/usr/share/containers/containers.conf` file and modify the chosen port for Aardvark DNS service. For example, change it to 54.

```
vi /usr/share/containers/containers.conf
...
# Port to use for dns forwarding daemon with netavark in rootful
bridge
# mode and dns enabled.
# Using an alternate port might be useful if other DNS services
should
# run on the machine.
#
dns_bind_port = 54
...
Esc:wq
```

- c. Reboot the Connector virtual machine.

## Result

At the end of the installation, the Connector service (occm) restarts twice if you specified a proxy server.

## Step 6: Register the Connector with BlueXP

Log into BlueXP and associate the Connector with your organization. How you log in depends on the mode in which you are using BlueXP. If you are using BlueXP in standard mode, you log in through the SaaS website. If you are using BlueXP in restricted or private mode, you log in locally from the Connector host.

## Steps

1. Open a web browser and enter the following URL:

`https://ipaddress`

*ipaddress* can be localhost, a private IP address, or a public IP address, depending on the configuration of the host. For example, if the Connector is in the public cloud without a public IP address, you must enter a private IP address from a host that has a connection to the Connector host.

2. Sign up or log in.
3. After you log in, set up BlueXP:
  - a. Specify the BlueXP organization to associate with the Connector.

- b. Enter a name for the system.
- c. Under **Are you running in a secured environment?** keep restricted mode disabled.

Keep restricted mode disabled because these steps use BlueXP in standard mode. (In addition, restricted mode isn't supported when the Connector is installed on-premises.)

- d. Select **Let's start**.

#### **Step 7: Provide permissions to BlueXP**

After you install and set up the Connector, add your cloud credentials so that BlueXP has the required permissions to perform actions in AWS or Azure.

## AWS

### Before you begin

If you just created these AWS credentials, they may take a few minutes to become available. Wait a few minutes before you add the credentials to BlueXP.

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



2. Select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location:** Select **Amazon Web Services > Connector**.
  - b. **Define Credentials:** Enter an AWS access key and secret key.
  - c. **Marketplace Subscription:** Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
  - d. **Review:** Confirm the details about the new credentials and select **Add**.

### Result

BlueXP now has the permissions that it needs to perform actions in AWS on your behalf.

You can now go to the [BlueXP console](#) to start using the Connector with BlueXP.

## Azure

### Before you begin

If you just created these Azure credentials, they may take a few minutes to become available. Wait a few minutes before you add the credentials to BlueXP.

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



2. Select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location:** Select **Microsoft Azure > Connector**.
  - b. **Define Credentials:** Enter information about the Microsoft Entra service principal that grants the required permissions:
    - Application (client) ID
    - Directory (tenant) ID
    - Client Secret
  - c. **Marketplace Subscription:** Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
  - d. **Review:** Confirm the details about the new credentials and select **Add**.

## Result

BlueXP now has the permissions that it needs to perform actions in Azure on your behalf. You can now go to the [BlueXP console](#) to start using the Connector with BlueXP.

## Subscribe to NetApp Intelligent Services (standard mode)

Subscribe to NetApp Intelligent Services from your cloud provider's marketplace to pay for data services at an hourly rate (PAYGO) or through an annual contract. If you purchased a license from NetApp (BYOL), you also need to subscribe to the marketplace offering. Your license is always charged first, but you'll be charged at the hourly rate if you exceed your licensed capacity or if the license's term expires.

A marketplace subscription enables charging for the following NetApp data services:

- Backup and recovery
- Cloud Volumes ONTAP
- Tiering
- Ransomware protection
- Disaster recovery

Classification is enabled through your subscription, but there is no charge for using classification.

### Before you begin

Subscribing to data services involves associating a marketplace subscription with the cloud credentials that are associated with a Connector. If you followed the "Get started with standard mode" workflow, then you should already have a Connector. To learn more, view the [Quick start for BlueXP in standard mode](#).



## AWS


The following video shows the steps to subscribe to NetApp Intelligent Services from the AWS Marketplace:

### [Subscribe to NetApp Intelligent Services from the AWS Marketplace](#)

#### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.
2. Select the action menu for a set of credentials and then select **Configure Subscription**.

You must select credentials that are associated with a Connector. You can't associate a marketplace subscription with credentials that are associated with BlueXP.



3. To associate the credentials with an existing subscription, select the subscription from the down-down list and select **Configure**.
4. To associate the credentials with a new subscription, select **Add Subscription > Continue** and follow the steps in the AWS Marketplace:
  - a. Select **View purchase options**.
  - b. Select **Subscribe**.
  - c. Select **Set up your account**.

You'll be redirected to the BlueXP website.

- d. From the **Subscription Assignment** page:
  - Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
  - In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

- Select **Save**.

## Azure Steps

1. In the upper right of the console, select the Settings icon, and select **Credentials**.
2. Select the action menu for a set of credentials and then select **Configure Subscription**.

You must select credentials that are associated with a Connector. You can't associate a marketplace subscription with credentials that are associated with BlueXP.

3. To associate the credentials with an existing subscription, select the subscription from the down-down list and select **Configure**.
4. To associate the credentials with a new subscription, select **Add Subscription > Continue** and follow the steps in the Azure Marketplace:
  - a. If prompted, log in to your Azure account.
  - b. Select **Subscribe**.
  - c. Fill out the form and select **Subscribe**.
  - d. After the subscription process is complete, select **Configure account now**.

You'll be redirected to BlueXP.

- e. From the **Subscription Assignment** page:
  - Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
  - In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

- Select **Save**.

The following video shows the steps to subscribe from the Azure Marketplace:


[Subscribe to NetApp Intelligent Services from the Azure Marketplace](#)

## Google Cloud

### Steps

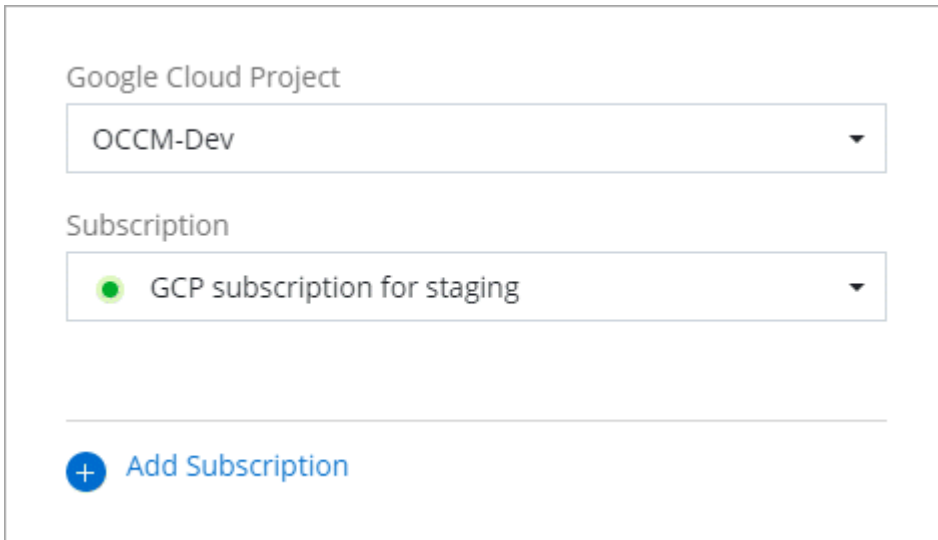
1. In the upper right of the console, select the Settings icon, and select **Credentials**.
2. Select the action menu for a set of credentials and then select **Configure Subscription**.

+new screenshot needed (TS)



3. To configure an existing subscription with the selected credentials, select a Google Cloud project and

subscription from the drop-down list, and then select **Configure**.



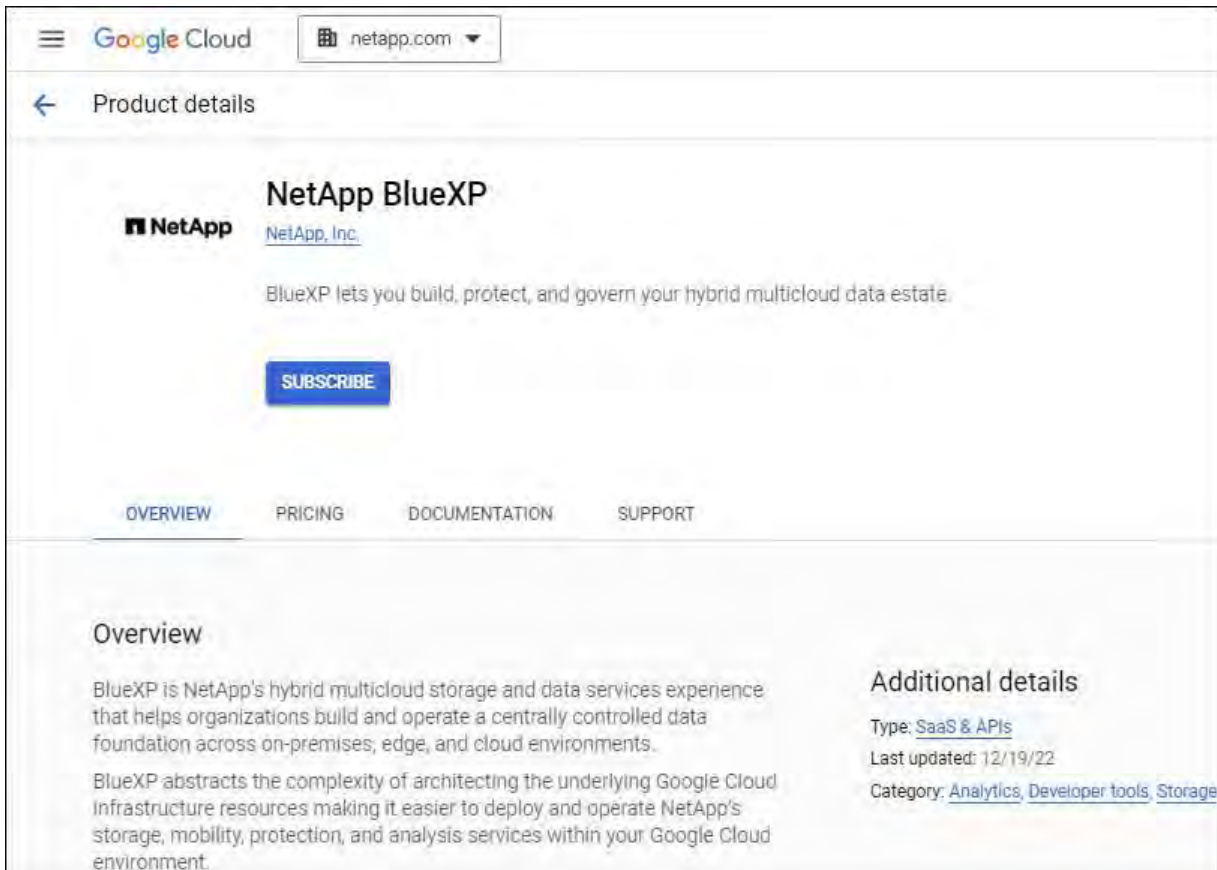
The screenshot shows a configuration interface with two dropdown menus. The first menu is labeled "Google Cloud Project" and has "OCCM-Dev" selected. The second menu is labeled "Subscription" and has "GCP subscription for staging" selected. Below the menus is a blue button with a plus sign and the text "Add Subscription".

4. If you don't already have a subscription, select **Add Subscription > Continue** and follow the steps in the Google Cloud Marketplace.



Before you complete the following steps, ensure that you have both Billing Admin privileges in your Google Cloud account as well as a BlueXP login.

- a. After you're redirected to the [NetApp Intelligent Services page on the Google Cloud Marketplace](#), ensure that the correct project is selected at the top navigation menu.



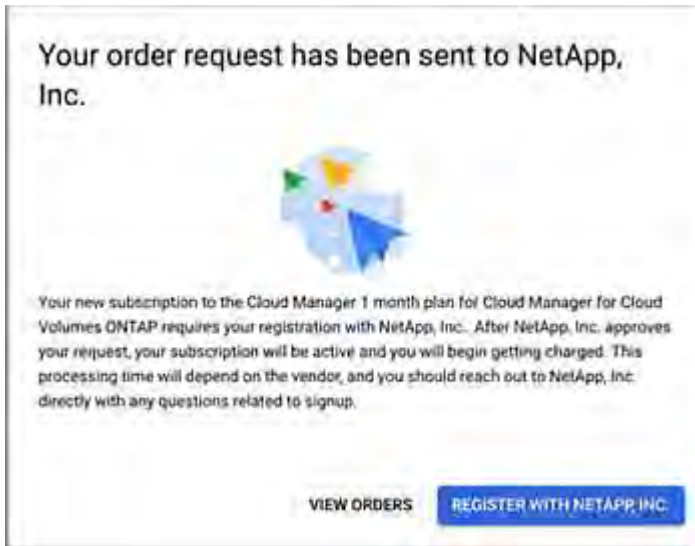
The screenshot shows the "Product details" page for NetApp BlueXP. At the top, there is a navigation bar with the Google Cloud logo and a dropdown menu showing "netapp.com". Below the navigation bar, there is a back arrow and the text "Product details". The main content area features the NetApp logo, the product name "NetApp BlueXP", and the company name "NetApp, Inc.". A description states: "BlueXP lets you build, protect, and govern your hybrid multicloud data estate." A prominent blue "SUBSCRIBE" button is highlighted with a yellow border. Below the description, there are four tabs: "OVERVIEW", "PRICING", "DOCUMENTATION", and "SUPPORT". The "OVERVIEW" tab is selected. The "Overview" section contains two paragraphs of text. The "Additional details" section on the right lists the product type as "SaaS & APIs", the last updated date as "12/19/22", and the category as "Analytics, Developer tools, Storage".

- b. Select **Subscribe**.
- c. Select the appropriate billing account and agree to the terms and conditions.
- d. Select **Subscribe**.

This step sends your transfer request to NetApp.

- e. On the pop-up dialog box, select **Register with NetApp, Inc.**

This step must be completed to link the Google Cloud subscription with your BlueXP organization or account. The process of linking a subscription isn't complete until you are redirected from this page and then sign in to BlueXP.



- f. Complete the steps on the **Subscription Assignment** page:



If someone from your organization has already subscribed to the NetApp BlueXP subscription from your billing account, then you will be redirected to [the Cloud Volumes ONTAP page on the BlueXP website](#) instead. If this is unexpected, contact your NetApp sales team. Google enables only one subscription per Google billing account.

- Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
- In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

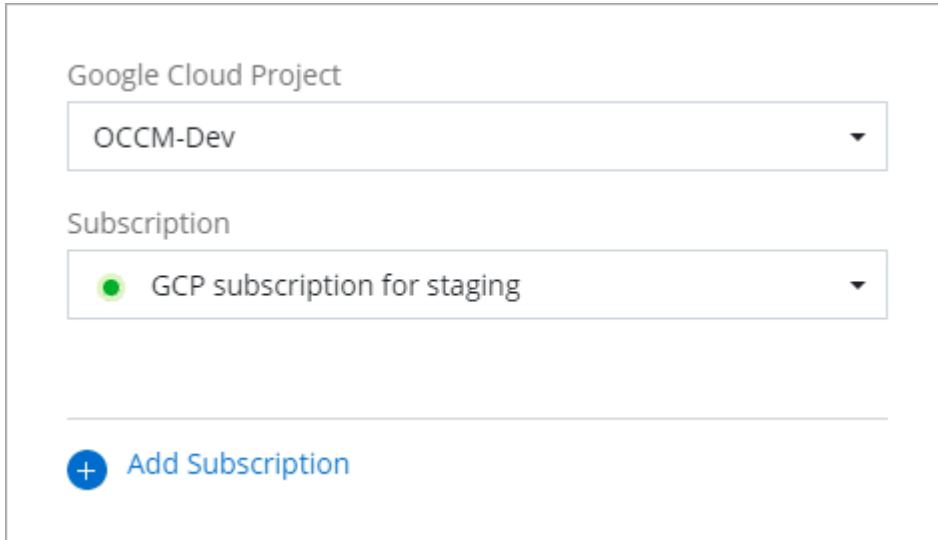
For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

- Select **Save**.

The following video shows the steps to subscribe from the Google Cloud Marketplace:

[Subscribe to BlueXP from the Google Cloud Marketplace](#)

- g. Once this process is complete, navigate back to the Credentials page in BlueXP and select this new subscription.



The screenshot shows a user interface for selecting a Google Cloud Project and Subscription. It features two dropdown menus. The first dropdown is labeled "Google Cloud Project" and has "OCCM-Dev" selected. The second dropdown is labeled "Subscription" and has "GCP subscription for staging" selected, which is preceded by a green circular indicator. Below the dropdowns is a horizontal line, and at the bottom left is a blue button with a plus sign and the text "Add Subscription".

#### Related information

- [Manage BYOL capacity-based licenses for Cloud Volumes ONTAP](#)
- [Manage BYOL licenses for data services](#)
- [Manage AWS credentials and subscriptions](#)
- [Manage Azure credentials and subscriptions](#)
- [Manage Google Cloud credentials and subscriptions](#)

### What you can do next (standard mode)

Now that you've logged in and set up BlueXP in standard mode, users can create and discover working environments and use BlueXP data services.



If you installed a Connector in AWS, Microsoft Azure, or Google Cloud, then BlueXP automatically discovers information about the Amazon S3 buckets, Azure Blob storage, or Google Cloud Storage buckets in the location where the Connector is installed. A working environment is automatically added to the BlueXP canvas.

For help, go to the [home page for the BlueXP documentation](#) to view the docs for all BlueXP services.

#### Related information

[BlueXP deployment modes](#)

## Get started with restricted mode

### Getting started workflow (restricted mode)

Get started with BlueXP in restricted mode by preparing your environment and deploying the Connector.

Restricted mode is typically used by state and local governments and regulated companies, including deployments in AWS GovCloud and Azure Government regions. Before getting started, ensure you have an understanding of [Connectors](#) and [deployment modes](#).

1

### Prepare for deployment

- a. Prepare a dedicated Linux host that meets requirements for CPU, RAM, disk space, container orchestration tool, and more.
- b. Set up networking that provides access to the target networks, outbound internet access for manual installations, and outbound internet for day-to-day access.
- c. Set up permissions in your cloud provider so that you can associate those permissions with the Connector instance after you deploy it.

2

### Deploy the Connector

- a. Install the Connector from your cloud provider's marketplace or by manually installing the software on your own Linux host.
- b. Set up BlueXP by opening a web browser and entering the Linux host's IP address.
- c. Provide BlueXP with the permissions that you previously set up.

3

### Subscribe to NetApp Intelligent Services (optional)

Optional: Subscribe to NetApp Intelligent Services from your cloud provider's marketplace to pay for data services at an hourly rate (PAYGO) or through an annual contract. NetApp Intelligent Services include Backup and recovery, Cloud Volumes ONTAP, Tiering, Ransomware protection, and Disaster recovery. Classification is included with your subscription at no additional cost.

## Prepare for deployment in restricted mode

Prepare your environment before you deploy BlueXP in restricted mode. For example, you need to review host requirements, prepare networking, set up permissions, and more.

### Step 1: Understand how restricted mode works

Before you get started, you should have an understanding of how BlueXP works in restricted mode.

For example, you should understand that you need to use the browser-based interface that is available locally from the BlueXP Connector that you need to install. You can't access BlueXP from the web-based console that's provided through the SaaS layer.

In addition, not all BlueXP services are available.

[Learn how restricted mode works.](#)

### Step 2: Review installation options

In restricted mode, you can only install the Connector in the cloud. The following installation options are available:

- From the AWS Marketplace
- From the Azure Marketplace
- Manually installing the Connector on your own Linux host that's running in AWS, Azure, or Google Cloud

### Step 3: Review host requirements

The Connector software must run on a host that meets specific operating system requirements, RAM requirements, port requirements, and so on.

When you deploy the Connector from the AWS or Azure Marketplace, the image includes the required OS and software components. You simply need to choose an instance type that meets CPU and RAM requirements.

#### Dedicated host

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

Host can be of any architecture that meets the following size requirements:

- CPU: 8 cores or 8 vCPUs
- RAM: 32 GB

#### Hypervisor

A bare metal or hosted hypervisor that is certified to run a supported operating system is required.

#### Operating system and container requirements

BlueXP supports the Connector with the following operating systems when using BlueXP in standard mode or restricted mode. A container orchestration tool is required before you install the Connector.

Operating system	Supported OS versions	Supported Connector versions	Required container tool	SELinux
Red Hat Enterprise Linux	9.1 to 9.4 8.6 to 8.10 7.9	3.9.40 or later with BlueXP in standard mode or restricted mode	Podman version 4.6.1 or 4.9.4 <a href="#">View Podman configuration requirements.</a>	Supported in enforcing mode or permissive mode <sup>1</sup>
Ubuntu	24.04 LTS	3.9.45 or later with BlueXP in standard mode or restricted mode	Docker Engine 23.06 to 28.0.0.	Not supported
	22.04 LTS	3.9.29 or later	Docker Engine 23.0.6 to 28.0.0.	Not supported

Notes:

1. Management of Cloud Volumes ONTAP systems is not supported by Connectors that have SELinux enabled on the operating system.

2. The Connector is supported on English-language versions of these operating systems.
3. For RHEL, the host must be registered with Red Hat Subscription Management. If it's not registered, the host can't access repositories to update required 3rd-party software during Connector installation.

### **AWS EC2 instance type**

An instance type that meets the CPU and RAM requirements above. We recommend t3.2xlarge.

### **Azure VM size**

An instance type that meets the CPU and RAM requirements above. We recommend Standard\_D8s\_v3.

### **Google Cloud machine type**

An instance type that meets the CPU and RAM requirements above. We recommend n2-standard-8.

The Connector is supported in Google Cloud on a VM instance with an OS that supports [Shielded VM features](#)

### **Disk space in /opt**

100 GiB of space must be available

BlueXP uses `/opt` to install the `/opt/application/netapp` directory and its contents.

### **Disk space in /var**

20 GiB of space must be available

BlueXP requires this space in `/var` because Docker or Podman are architected to create the containers within this directory. Specifically, they will create containers in the `/var/lib/containers/storage` directory. External mounts or symlinks do not work for this space.

## **Step 4: Install Podman or Docker Engine**

If you're planning to manually install the Connector software, you need to prepare the host by installing Podman or Docker Engine.

Depending on your operating system, either Podman or Docker Engine is required before you install the Connector.

- Podman is required for Red Hat Enterprise Linux 8 and 9.

[View the Podman versions that BlueXP supports.](#)

- Docker Engine is required for Ubuntu.

[View the Docker Engine versions that BlueXP supports.](#)



## Example 5. Steps

### Podman

Follow these steps to install Podman and configure it to meet the following requirements:

- Enable and start the podman.socket service
- Install python3
- Install the podman-compose package version 1.0.6
- Add podman-compose to the PATH environment variable



When using Podman, adjust the aardvark-dns service port (default: 53) after installing the Connector to avoid conflicts with the DNS port on the host. Follow the instructions to configure the port.

### Steps

1. Remove the podman-docker package if it's installed on the host.

```
dnf remove podman-docker
rm /var/run/docker.sock
```

2. Install Podman.

Podman is available from official Red Hat Enterprise Linux repositories.

For Red Hat Enterprise Linux 9:

```
sudo dnf install podman-2:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

For Red Hat Enterprise Linux 8:

```
sudo dnf install podman-3:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

3. Enable and start the podman.socket service.

```
sudo systemctl enable --now podman.socket
```

4. Install python3.

```
sudo dnf install python3
```

5. Install the EPEL repository package if it's not already available on your system.

This step is required because podman-compose is available from the Extra Packages for Enterprise Linux (EPEL) repository.

For Red Hat Enterprise Linux 9:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-9.noarch.rpm
```

For Red Hat Enterprise Linux 8:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

6. Install podman-compose package 1.0.6.

```
sudo dnf install podman-compose-1.0.6
```



Using the `dnf install` command meets the requirement for adding podman-compose to the PATH environment variable. The installation command adds podman-compose to `/usr/bin`, which is already included in the `secure_path` option on the host.

### Docker Engine

Follow the documentation from Docker to install Docker Engine.

#### Steps

1. [View installation instructions from Docker](#)

Be sure to follow the steps to install a specific version of Docker Engine. Installing the latest version will install a version of Docker that BlueXP doesn't support.

2. Verify that Docker is enabled and running.

```
sudo systemctl enable docker && sudo systemctl start docker
```

## Step 5: Prepare networking

Set up your networking so the Connector can manage resources and processes within your public cloud environment. Other than having a virtual network and subnet for the Connector, you'll need to ensure that the

following requirements are met.

### Connections to target networks

The Connector must have a network connection to the location where you plan to manage storage. For example, the VPC or VNet where you plan to deploy Cloud Volumes ONTAP, or the data center where your on-premises ONTAP clusters reside.

### Prepare networking for user access to BlueXP console

In restricted mode, the BlueXP user interface is accessible from the Connector. As you use the BlueXP user interface, it contacts a few endpoints to complete data management tasks. These endpoints are contacted from a user's computer when completing specific actions from the BlueXP console.

Endpoints	Purpose
<a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a>	The BlueXP web-based console contacts this endpoint to interact with the BlueXP API for actions related to authorization, licensing, subscriptions, credentials, notifications, and more.
<a href="https://signin.b2c.netapp.com">https://signin.b2c.netapp.com</a>	Required to update NetApp Support Site (NSS) credentials or to add new NSS credentials to BlueXP.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://cdn.auth0.com">https://cdn.auth0.com</a> <a href="https://services.cloud.netapp.com">https://services.cloud.netapp.com</a>	Your web browser connects to these endpoints for centralized user authentication through BlueXP.
<a href="https://widget.intercom.io">https://widget.intercom.io</a>	For in-product chat that enables you to talk to NetApp cloud experts.

### Endpoints contacted during manual installation

When you manually install the Connector on your own Linux host, the installer for the Connector requires access to several URLs during the installation process.

- The following endpoints are always contacted no matter where you install the Connector:
  - <https://mysupport.netapp.com>
  - <https://signin.b2c.netapp.com> (this endpoint is the CNAME URL for <https://mysupport.netapp.com>)
  - <https://cloudmanager.cloud.netapp.com/tenancy>
  - <https://stream.cloudmanager.cloud.netapp.com>
  - <https://production-artifacts.cloudmanager.cloud.netapp.com>
- If you install the Connector in an AWS Government region, the installer also needs access to these endpoints:
  - [https://\\*.blob.core.windows.net](https://*.blob.core.windows.net)
  - <https://cloudmanagerinfraproduct.azurecr.io>
- If you install the Connector in an Azure Government region, the installer also needs access to these endpoints:
  - [https://\\*.blob.core.windows.net](https://*.blob.core.windows.net)
  - <https://occmclientinfragov.azurecr.us>

- If you install the Connector in a commercial region or sovereign region, you can choose between two sets of endpoints:
  - Option 1 (recommended):
    - <https://bluexpinfraprod.eastus2.data.azurecr.io>
    - <https://bluexpinfraprod.azurecr.io>
  - Option 2:
    - [https://\\*.blob.core.windows.net](https://*.blob.core.windows.net)
    - <https://cloudmanagerinfraprod.azurecr.io>

The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

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

### Outbound internet access for day-to-day operations

The network location where you deploy the Connector must have an outbound internet connection. The Connector requires outbound internet access to contact the following endpoints in order to manage resources and processes within your public cloud environment.

Endpoints	Purpose
AWS services (amazonaws.com): <ul style="list-style-type: none"> <li>• CloudFormation</li> <li>• Elastic Compute Cloud (EC2)</li> <li>• Identity and Access Management (IAM)</li> <li>• Key Management Service (KMS)</li> <li>• Security Token Service (STS)</li> <li>• Simple Storage Service (S3)</li> </ul>	To manage resources in AWS. The exact endpoint depends on the AWS region that you're using. <a href="#">Refer to AWS documentation for details</a>
<a href="https://management.azure.com">https://management.azure.com</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a> <a href="https://blob.core.windows.net">https://blob.core.windows.net</a> <a href="https://core.windows.net">https://core.windows.net</a>	To manage resources in Azure public regions.

Endpoints	Purpose
<a href="https://management.usgovcloudapi.net">https://management.usgovcloudapi.net</a> <a href="https://login.microsoftonline.us">https://login.microsoftonline.us</a> <a href="https://blob.core.usgovcloudapi.net">https://blob.core.usgovcloudapi.net</a> <a href="https://core.usgovcloudapi.net">https://core.usgovcloudapi.net</a>	To manage resources in Azure Government regions.
<a href="https://management.chinacloudapi.cn">https://management.chinacloudapi.cn</a> <a href="https://login.chinacloudapi.cn">https://login.chinacloudapi.cn</a> <a href="https://blob.core.chinacloudapi.cn">https://blob.core.chinacloudapi.cn</a> <a href="https://core.chinacloudapi.cn">https://core.chinacloudapi.cn</a>	To manage resources in Azure China regions.
<a href="https://www.googleapis.com/compute/v1/">https://www.googleapis.com/compute/v1/</a> <a href="https://compute.googleapis.com/compute/v1">https://compute.googleapis.com/compute/v1</a> <a href="https://cloudresourcemanager.googleapis.com/v1/projects">https://cloudresourcemanager.googleapis.com/v1/projects</a> <a href="https://www.googleapis.com/compute/beta">https://www.googleapis.com/compute/beta</a> <a href="https://storage.googleapis.com/storage/v1">https://storage.googleapis.com/storage/v1</a> <a href="https://www.googleapis.com/storage/v1">https://www.googleapis.com/storage/v1</a> <a href="https://iam.googleapis.com/v1">https://iam.googleapis.com/v1</a> <a href="https://cloudkms.googleapis.com/v1">https://cloudkms.googleapis.com/v1</a> <a href="https://www.googleapis.com/deploymentmanager/v2/projects">https://www.googleapis.com/deploymentmanager/v2/projects</a>	To manage resources in Google Cloud.
<a href="https://support.netapp.com">https://support.netapp.com</a> <a href="https://mysupport.netapp.com">https://mysupport.netapp.com</a>	To obtain licensing information and to send AutoSupport messages to NetApp support.
<a href="https://*.api.bluexp.netapp.com">https://*.api.bluexp.netapp.com</a> <a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a> <a href="https://*.cloudmanager.cloud.netapp.com">https://*.cloudmanager.cloud.netapp.com</a> <a href="https://cloudmanager.cloud.netapp.com">https://cloudmanager.cloud.netapp.com</a> <a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	To provide SaaS features and services within BlueXP.
If the Connector is in an AWS Government region: <a href="https://*.blob.core.windows.net">https://*.blob.core.windows.net</a> <a href="https://cloudmanagerinfragov.azurecr.io">https://cloudmanagerinfragov.azurecr.io</a>	To obtain images for Connector upgrades when the Connector is installed in an AWS Government region.
If the Connector is in an Azure Government region: <a href="https://*.blob.core.windows.net">https://*.blob.core.windows.net</a> <a href="https://occmclientinfragov.azurecr.us">https://occmclientinfragov.azurecr.us</a>	To obtain images for Connector upgrades when the Connector is installed in an Azure Government region.

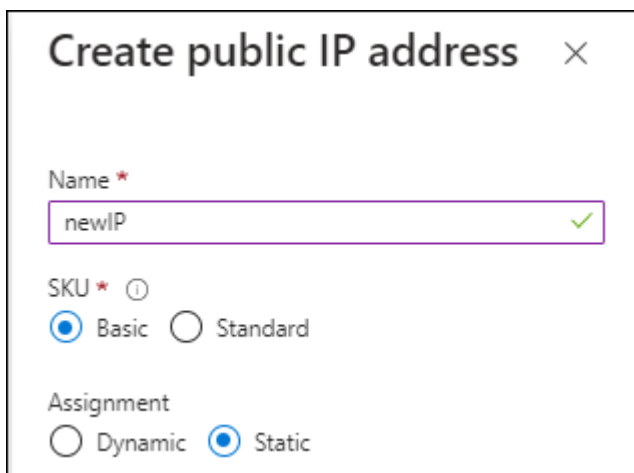
Endpoints	Purpose
<p>If the Connector is in a commercial region or sovereign region, you can choose between two sets of endpoints:</p> <ul style="list-style-type: none"> <li>Option 1 (recommended) <sup>1</sup></li> </ul> <p>https://bluexpinfraprod.eastus2.data.azurecr.io https://bluexpinfraprod.azurecr.io</p> <ul style="list-style-type: none"> <li>Option 2</li> </ul> <p>https://*.blob.core.windows.net https://cloudmanagerinfraprod.azurecr.io</p>	<p>To obtain images for Connector upgrades when the Connector is installed in a commercial region or sovereign region.</p>

<sup>1</sup> The endpoints listed in option 1 are recommended because they are more secure. We recommend that you set up your firewall to allow the endpoints listed in option 1, while disallowing the endpoints listed in option 2. Note the following about these endpoints:

- The endpoints listed in option 1 are supported starting with the 3.9.47 release of the Connector. There is no backwards compatibility with previous releases of the Connector.
- The Connector contacts the endpoints listed in option 2 first. If those endpoints aren't accessible, the Connector automatically contacts the endpoints listed in option 1.
- The endpoints in option 1 are not supported if you use the Connector with BlueXP backup and recovery or BlueXP ransomware protection. In this case, you can disallow the endpoints listed in option 1, while allowing the endpoints listed in option 2.

### Public IP address in Azure

If you want to use a public IP address with the Connector VM in Azure, the IP address must use a Basic SKU to ensure that BlueXP uses this public IP address.



**Create public IP address** ×

Name \*  
newIP ✓

SKU \* ⓘ  
 Basic    Standard

Assignment  
 Dynamic    Static

If you use a Standard SKU IP address instead, then BlueXP uses the *private* IP address of the Connector, instead of the public IP. If the machine that you're using to access the BlueXP Console doesn't have access to that private IP address, then actions from the BlueXP Console will fail.

[Azure documentation: Public IP SKU](#)

## Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

## Ports

There's no incoming traffic to the Connector, unless you initiate it or if the Connector is used as a proxy to send AutoSupport messages from Cloud Volumes ONTAP to NetApp Support.

- HTTP (80) and HTTPS (443) provide access to the local UI, which you'll use in rare circumstances.
- SSH (22) is only needed if you need to connect to the host for troubleshooting.
- Inbound connections over port 3128 are required if you deploy Cloud Volumes ONTAP systems in a subnet where an outbound internet connection isn't available.

If Cloud Volumes ONTAP systems don't have an outbound internet connection to send AutoSupport messages, BlueXP automatically configures those systems to use a proxy server that's included with the Connector. The only requirement is to ensure that the Connector's security group allows inbound connections over port 3128. You'll need to open this port after you deploy the Connector.

## Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

If you're planning to create the Connector from your cloud provider's marketplace, then you'll need to implement this networking requirement after you create the Connector.

## Step 6: Prepare cloud permissions

BlueXP requires permissions from your cloud provider to deploy Cloud Volumes ONTAP in a virtual network and to use BlueXP data services. You need to set up permissions in your cloud provider and then associate those permissions with the Connector.

To view the required steps, select the authentication option that you'd like to use for your cloud provider.

## AWS IAM role

Use an IAM role to provide the Connector with permissions.

If you're creating the Connector from the AWS Marketplace, you'll be prompted to select that IAM role when you launch the EC2 instance.

If you're manually installing the Connector on your own Linux host, you'll need to attach the role to the EC2 instance.

### Steps

1. Log in to the AWS console and navigate to the IAM service.
2. Create a policy:
  - a. Select **Policies > Create policy**.
  - b. Select **JSON** and copy and paste the contents of the [IAM policy for the Connector](#).
  - c. Finish the remaining steps to create the policy.
3. Create an IAM role:
  - a. Select **Roles > Create role**.
  - b. Select **AWS service > EC2**.
  - c. Add permissions by attaching the policy that you just created.
  - d. Finish the remaining steps to create the role.

### Result

You now have an IAM role for the Connector EC2 instance.

## AWS access key

Set up permissions and an access key for an IAM user. You'll need to provide BlueXP with the AWS access key after you install the Connector and set up BlueXP.

### Steps

1. Log in to the AWS console and navigate to the IAM service.
2. Create a policy:
  - a. Select **Policies > Create policy**.
  - b. Select **JSON** and copy and paste the contents of the [IAM policy for the Connector](#).
  - c. Finish the remaining steps to create the policy.

Depending on the BlueXP services that you're planning to use, you might need to create a second policy.

For standard regions, the permissions are spread across two policies. Two policies are required due to a maximum character size limit for managed policies in AWS. [Learn more about IAM policies for the Connector](#).

3. Attach the policies to an IAM user.
  - [AWS Documentation: Creating IAM Roles](#)
  - [AWS Documentation: Adding and Removing IAM Policies](#)



4. Ensure that the user has an access key that you can add to BlueXP after you install the Connector.

### Result

The account now has the required permissions.

### Azure role

Create an Azure custom role with the required permissions. You'll assign this role to the Connector VM.

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would prefer to use a different method, refer to [Azure documentation](#)

### Steps

1. If you're planning to manually install the software on your own host, enable a system-assigned managed identity on the VM so that you can provide the required Azure permissions through a custom role.

[Microsoft Azure documentation: Configure managed identities for Azure resources on a VM using the Azure portal](#)

2. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
3. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription that you want to use with BlueXP.

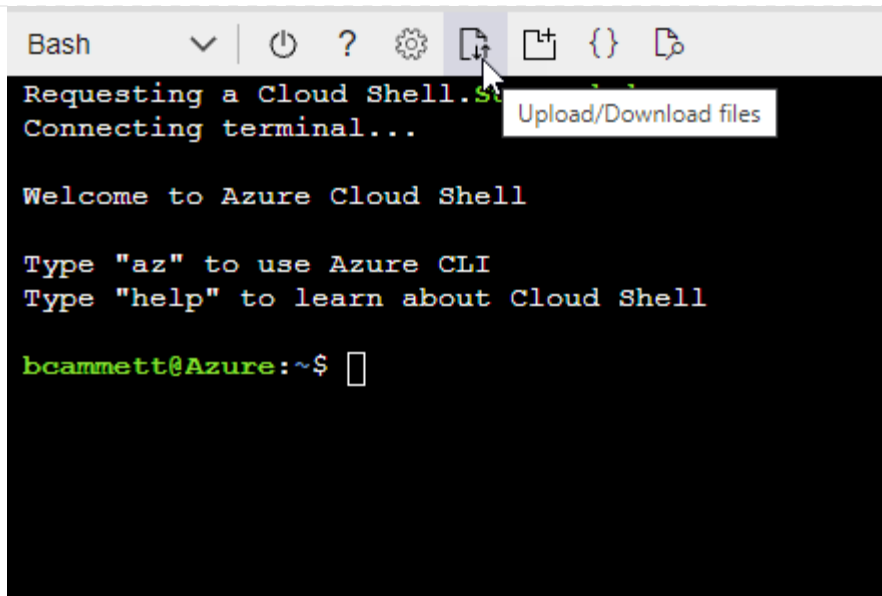
### Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- a. Start [Azure Cloud Shell](#) and choose the Bash environment.
- b. Upload the JSON file.



c. Use the Azure CLI to create the custom role:

```
az role definition create --role-definition Connector_Policy.json
```

### Result

You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.

### Azure service principal

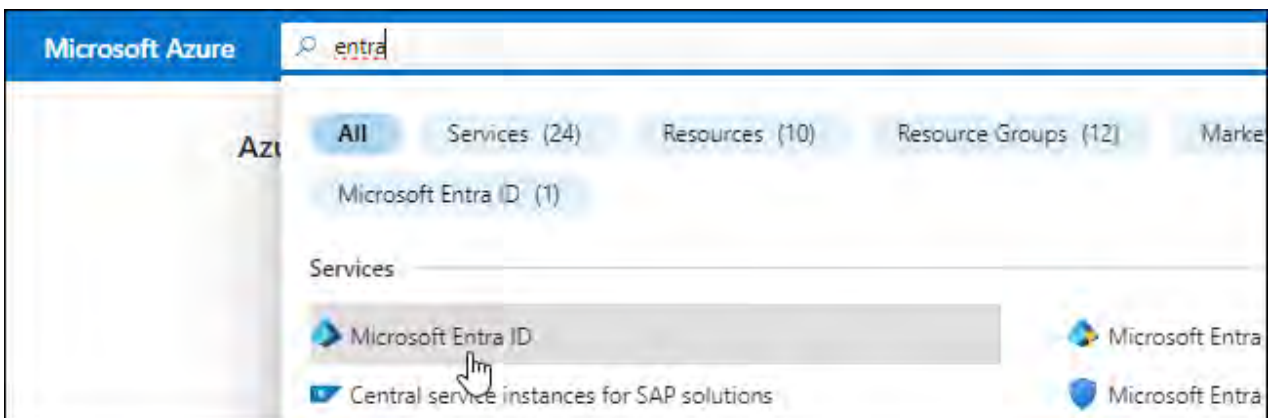
Create and set up a service principal in Microsoft Entra ID and obtain the Azure credentials that BlueXP needs. You'll need to provide BlueXP with these credentials after you install the Connector and set up BlueXP.

### Create a Microsoft Entra application for role-based access control

1. Ensure that you have permissions in Azure to create an Active Directory application and to assign the application to a role.

For details, refer to [Microsoft Azure Documentation: Required permissions](#)

2. From the Azure portal, open the **Microsoft Entra ID** service.



3. In the menu, select **App registrations**.
4. Select **New registration**.
5. Specify details about the application:
  - **Name**: Enter a name for the application.
  - **Account type**: Select an account type (any will work with BlueXP).
  - **Redirect URI**: You can leave this field blank.
6. Select **Register**.

You've created the AD application and service principal.

### Assign the application to a role

1. Create a custom role:

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would prefer to use a different method, refer to [Azure documentation](#)

- a. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
- b. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription from which users will create Cloud Volumes ONTAP systems.

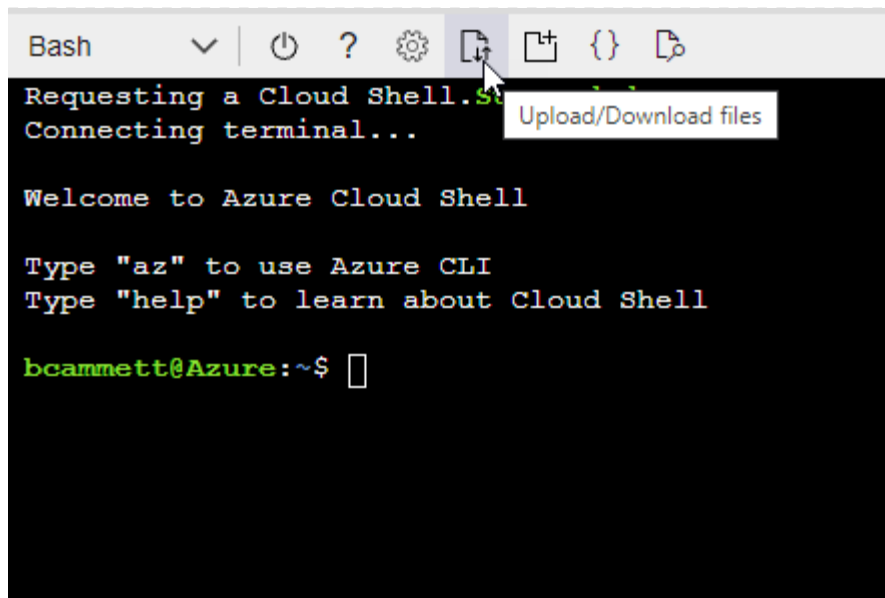
#### Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- Start [Azure Cloud Shell](#) and choose the Bash environment.
- Upload the JSON file.



```
Bash
Requesting a Cloud Shell.
Connecting terminal...
Welcome to Azure Cloud Shell

Type "az" to use Azure CLI
Type "help" to learn about Cloud Shell


bcammett@Azure:~$
```

- Use the Azure CLI to create the custom role:

```
az role definition create --role-definition
Connector_Policy.json
```

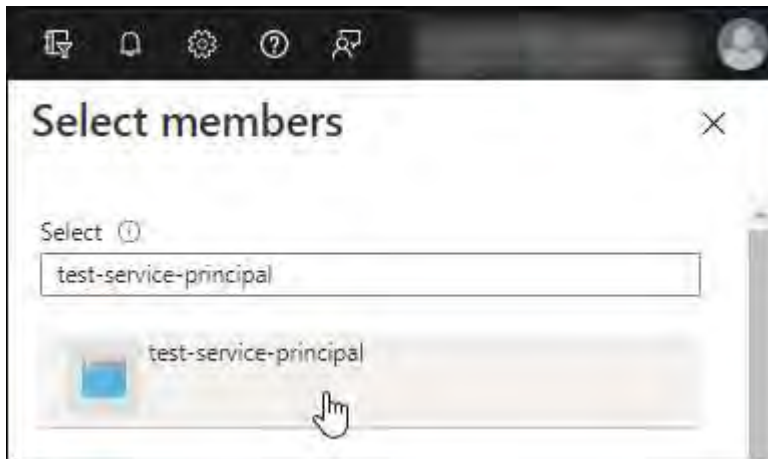
You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.

2. Assign the application to the role:
  - a. From the Azure portal, open the **Subscriptions** service.
  - b. Select the subscription.
  - c. Select **Access control (IAM) > Add > Add role assignment**.
  - d. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.
  - e. In the **Members** tab, complete the following steps:
    - Keep **User, group, or service principal** selected.
    - Select **Select members**.



- Search for the name of the application.

Here's an example:



- Select the application and select **Select**.
  - Select **Next**.
- f. Select **Review + assign**.

The service principal now has the required Azure permissions to deploy the Connector.

If you want to deploy Cloud Volumes ONTAP from multiple Azure subscriptions, then you must bind the service principal to each of those subscriptions. BlueXP enables you to select the subscription that you want to use when deploying Cloud Volumes ONTAP.

#### Add Windows Azure Service Management API permissions














1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Select **API permissions > Add a permission**.
3. Under **Microsoft APIs**, select **Azure Service Management**.

## Request API permissions

Select an API

Microsoft APIs APIs my organization uses My APIs

### Commonly used Microsoft APIs

<b>Microsoft Graph</b> Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10. Access Azure AD, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint. 		
 <b>Azure Batch</b> Schedule large-scale parallel and HPC applications in the cloud	 <b>Azure Data Catalog</b> Programmatic access to Data Catalog resources to register, annotate and search data assets	 <b>Azure Data Explorer</b> Perform ad-hoc queries on terabytes of data to build near real-time and complex analytics solutions
 <b>Azure Data Lake</b> Access to storage and compute for big data analytic scenarios	 <b>Azure DevOps</b> Integrate with Azure DevOps and Azure DevOps server	 <b>Azure Import/Export</b> Programmatic control of import/export jobs
 <b>Azure Key Vault</b> Manage your key vaults as well as the keys, secrets, and certificates within your Key Vaults	 <b>Azure Rights Management Services</b> Allow validated users to read and write protected content	 <b>Azure Service Management</b> Programmatic access to much of the functionality available through the Azure portal
 <b>Azure Storage</b> Secure, massively scalable object and data lake storage for unstructured and semi-structured data	 <b>Customer Insights</b> Create profile and interaction models for your products	 <b>Data Export Service for Microsoft Dynamics 365</b> Export data from Microsoft Dynamics CRM organization to an external destination

4. Select **Access Azure Service Management as organization users** and then select **Add permissions**.

## Request API permissions

< All APIs



Azure Service Management

<https://management.azure.com/> [Docs](#)

What type of permissions does your application require?

### Delegated permissions

Your application needs to access the API as the signed-in user.

### Application permissions

Your application runs as a background service or daemon without a signed-in user.

Select permissions

[expand all](#)

Type to search

PERMISSION

ADMIN CONSENT REQUIRED




user\_impersonation

Access Azure Service Management as organization users (preview)

## Get the application ID and directory ID for the application

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Copy the **Application (client) ID** and the **Directory (tenant) ID**.



When you add the Azure account to BlueXP, you need to provide the application (client) ID and the directory (tenant) ID for the application. BlueXP uses the IDs to programmatically sign in.

## Create a client secret

1. Open the **Microsoft Entra ID** service.
2. Select **App registrations** and select your application.
3. Select **Certificates & secrets > New client secret**.
4. Provide a description of the secret and a duration.
5. Select **Add**.
6. Copy the value of the client secret.

## Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

[+ New client secret](#)

DESCRIPTION	EXPIRES	VALUE	Copy to clipboard
test secret	8/16/2020	*sZ1jSe2By:D*-ZRov4NLfdAcY7:+0vA	

You now have a client secret that BlueXP can use it to authenticate with Microsoft Entra ID.

### Result

Your service principal is now setup and you should have copied the application (client) ID, the directory (tenant) ID, and the value of the client secret. You need to enter this information in BlueXP when you add an Azure account.

### Google Cloud service account

Create a role and apply it to a service account that you'll use for the Connector VM instance.

### Steps

1. Create a custom role in Google Cloud:
  - a. Create a YAML file that includes the permissions defined in the [Connector policy for Google Cloud](#).
  - b. From Google Cloud, activate cloud shell.
  - c. Upload the YAML file that includes the required permissions for the Connector.
  - d. Create a custom role by using the `gcloud iam roles create` command.

The following example creates a role named "connector" at the project level:

```
gcloud iam roles create connector --project=myproject
--file=connector.yaml
```

[Google Cloud docs: Creating and managing custom roles](#)

2. Create a service account in Google Cloud:
  - a. From the IAM & Admin service, select **Service Accounts > Create Service Account**.
  - b. Enter service account details and select **Create and Continue**.
  - c. Select the role that you just created.
  - d. Finish the remaining steps to create the role.

[Google Cloud docs: Creating a service account](#)

### Result

You now have a service account that you can assign to the Connector VM instance.



## Step 7: Enable Google Cloud APIs

Several APIs are required to deploy Cloud Volumes ONTAP in Google Cloud.

### Step

#### 1. [Enable the following Google Cloud APIs in your project](#)

- Cloud Deployment Manager V2 API
- Cloud Logging API
- Cloud Resource Manager API
- Compute Engine API
- Identity and Access Management (IAM) API
- Cloud Key Management Service (KMS) API

(Required only if you are planning to use BlueXP backup and recovery with customer-managed encryption keys (CMEK))

## Deploy the Connector in restricted mode

Deploy the Connector in restricted mode so that you can use BlueXP with limited outbound connectivity. To get started, install the Connector, set up BlueXP by accessing the user interface that's running on the Connector, and then provide the cloud permissions that you previously set up.

### Step 1: Install the Connector

Install the Connector from your cloud provider's marketplace or by manually installing the software on your own Linux host.

## AWS Commercial Marketplace

### Before you begin

You should have the following:

- A VPC and subnet that meets networking requirements.

[Learn about networking requirements](#)

- An IAM role with an attached policy that includes the required permissions for the Connector.

[Learn how to set up AWS permissions](#)

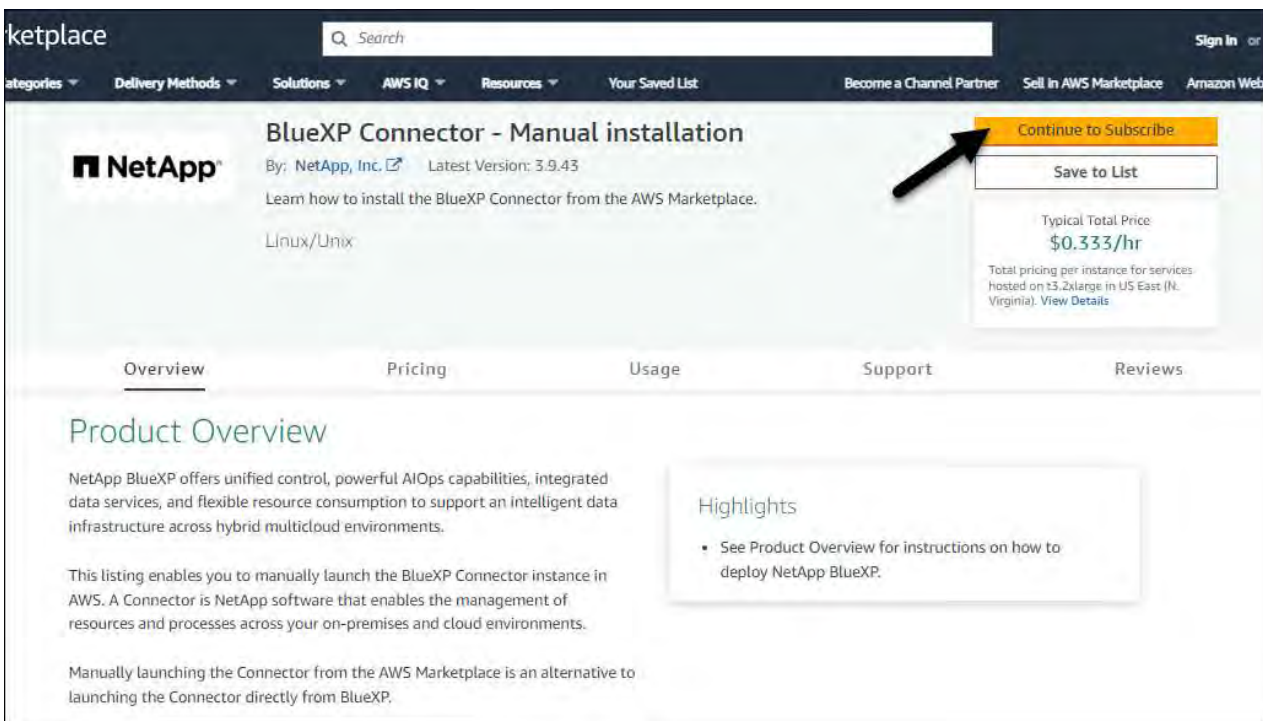
- Permissions to subscribe and unsubscribe from the AWS Marketplace for your IAM user.
- An understanding of CPU and RAM requirements for the instance.

[Review instance requirements.](#)

- A key pair for the EC2 instance.

### Steps

1. Go to the [BlueXP Connector listing on the AWS Marketplace](#)
2. On the Marketplace page, select **Continue to Subscribe**.



The screenshot shows the AWS Marketplace listing for NetApp BlueXP Connector - Manual installation. The page features a dark navigation bar with a search bar and various menu items. The main content area displays the product name, version (3.9.43), and a 'Continue to Subscribe' button highlighted with a black arrow. Below the button is a 'Save to List' button and pricing information: Typical Total Price \$0.333/hr. The page also includes a 'Product Overview' section and a 'Highlights' section.

3. To subscribe to the software, select **Accept Terms**.

The subscription process can take a few minutes.

4. After the subscription process is complete, select **Continue to Configuration**.

ketplace  Hello,

Categories ▾ Delivery Methods ▾ Solutions ▾ AWS IQ ▾ Resources ▾ Your Saved List Become a Channel Partner Sell in AWS Marketplace Amazon Web Se

**NetApp** BlueXP Connector - Manual installation [Continue to Configuration](#)

[< Product Detail](#) [Subscribe](#)

## Subscribe to this software

You're subscribed to this software. Please see the terms and pricing details below or click the button above to configure your software.

### Terms and Conditions

NetApp, Inc. Offer

You have subscribed to this software and agreed that your use of this software is subject to the pricing terms and the seller's End User License Agreement (EULA) [↗](#). You agreed that AWS may share information about this transaction (including your payment terms) with the respective seller, reseller or underlying provider, as applicable, in accordance with the [AWS Privacy Notice](#) [↗](#). AWS will issue invoices and collect payments from you on behalf of the seller through your AWS account. Your use of AWS services remains subject to the [AWS Customer Agreement](#) [↗](#) or other agreement with AWS governing your use of such services.

Product	Effective date	Expiration date	Action
BlueXP Connector - Manual installation	N/A	N/A	<a href="#">Show Details</a>

5. On the **Configure this software** page, ensure that you've selected the correct region and then select **Continue to Launch**.
6. On the **Launch this software** page, under **Choose Action**, select **Launch through EC2** and then select **Launch**.

These steps describe how to launch the instance from the EC2 Console because the console enables you to attach an IAM role to the Connector instance. This isn't possible using the **Launch from Website** action.

7. Follow the prompts to configure and deploy the instance:
  - **Name and tags:** Enter a name and tags for the instance.
  - **Application and OS Images:** Skip this section. The Connector AMI is already selected.
  - **Instance type:** Depending on region availability, choose an instance type that meets RAM and CPU requirements (t3.2xlarge is preselected and recommended).
  - **Key pair (login):** Select the key pair that you want to use to securely connect to the instance.
  - **Network settings:** Edit the network settings as needed:
    - Choose the desired VPC and subnet.
    - Specify whether the instance should have a public IP address.
    - Specify security group settings that enable the required connection methods for the Connector instance: SSH, HTTP, and HTTPS.

[View security group rules for AWS.](#)
  - **Configure storage:** Keep the default size and disk type for the root volume.

If you want to enable Amazon EBS encryption on the root volume, select **Advanced**, expand **Volume 1**, select **Encrypted**, and then choose a KMS key.

- **Advanced details:** Under **IAM instance profile**, choose the IAM role that includes the required permissions for the Connector.
- **Summary:** Review the summary and select **Launch instance**.

## Result

AWS launches the software with the specified settings. The Connector instance and software should be running in approximately five minutes.

## What's next?

Set up BlueXP.

## AWS Gov Marketplace

### Before you begin

You should have the following:

- A VPC and subnet that meets networking requirements.

[Learn about networking requirements](#)

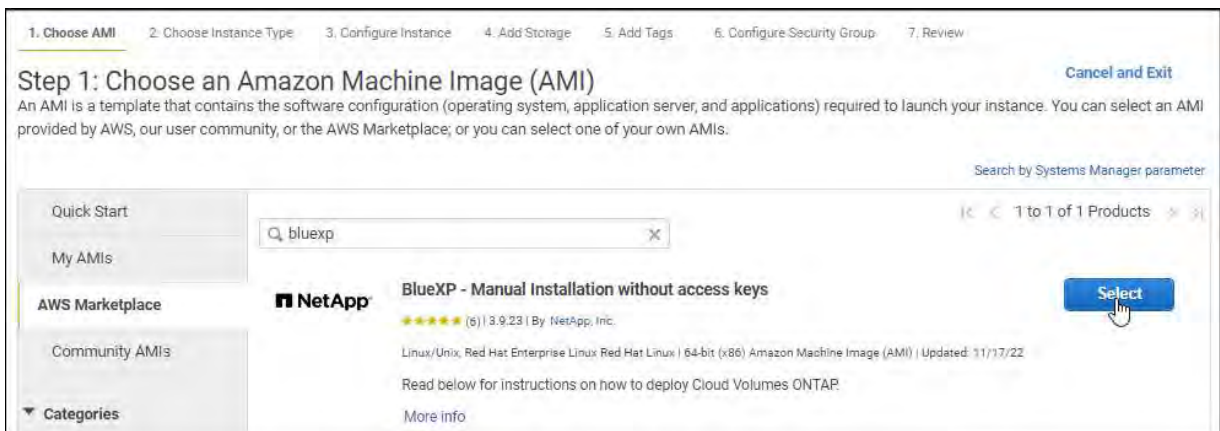
- An IAM role with an attached policy that includes the required permissions for the Connector.

[Learn how to set up AWS permissions](#)

- Permissions to subscribe and unsubscribe from the AWS Marketplace for your IAM user.
- A key pair for the EC2 instance.

## Steps

1. Go to the BlueXP offering in the AWS Marketplace.
  - a. Open the EC2 service and select **Launch instance**.
  - b. Select **AWS Marketplace**.
  - c. Search for BlueXP and select the offering.



- d. Select **Continue**.
2. Follow the prompts to configure and deploy the instance:
    - **Choose an Instance Type:** Depending on region availability, choose one of the supported instance types (t3.2xlarge is recommended).

## Review the instance requirements.

- **Configure Instance Details:** Select a VPC and subnet, choose the IAM role that you created in step 1, enable termination protection (recommended), and choose any other configuration options that meet your requirements.

Number of instances	<input type="text" value="1"/>	<a href="#">Launch into Auto Scaling Group</a>
Purchasing option	<input type="checkbox"/> Request Spot instances	
Network	<input type="text" value="vpc-a76d91c2   VPC4QA (default)"/>	<a href="#">Create new VPC</a>
Subnet	<input type="text" value="subnet-39536c13   QASubnet1   us-east-1b"/> 155 IP Addresses available	<a href="#">Create new subnet</a>
Auto-assign Public IP	<input type="text" value="Enable"/>	
Placement group	<input type="checkbox"/> Add instance to placement group	
Capacity Reservation	<input type="text" value="Open"/>	<a href="#">Create new Capacity Reservation</a>
IAM role	<input type="text" value="Cloud_Manager"/>	<a href="#">Create new IAM role</a>
CPU options	<input type="checkbox"/> Specify CPU options	
Shutdown behavior	<input type="text" value="Stop"/>	
Enable termination protection	<input checked="" type="checkbox"/> Protect against accidental termination	
Monitoring	<input type="checkbox"/> Enable CloudWatch detailed monitoring <a href="#">Additional charges apply.</a>	

- **Add Storage:** Keep the default storage options.
- **Add Tags:** Enter tags for the instance, if desired.
- **Configure Security Group:** Specify the required connection methods for the Connector instance: SSH, HTTP, and HTTPS.
- **Review:** Review your selections and select **Launch**.

## Result

AWS launches the software with the specified settings. The Connector instance and software should be running in approximately five minutes.

## What's next?

Set up BlueXP.

## Azure Marketplace

### Before you begin

You should have the following:

- A VNet and subnet that meets networking requirements.

[Learn about networking requirements](#)

- An Azure custom role that includes the required permissions for the Connector.

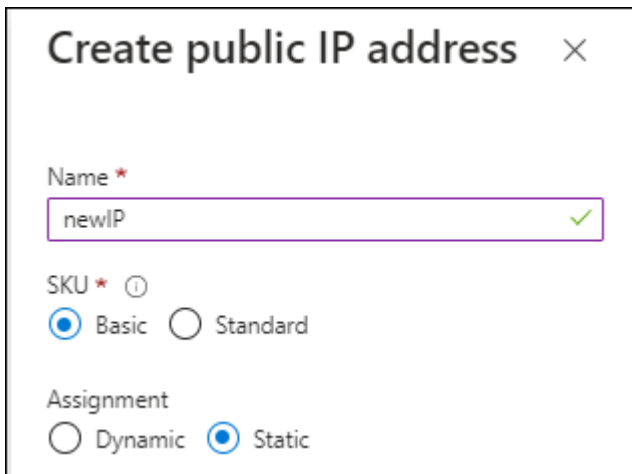
[Learn how to set up Azure permissions](#)

## Steps

1. Go to the NetApp Connector VM page in the Azure Marketplace.
  - [Azure Marketplace page for commercial regions](#)
  - [Azure Marketplace page for Azure Government regions](#)
2. Select **Get it now** and then select **Continue**.
3. From the Azure portal, select **Create** and follow the steps to configure the virtual machine.

Note the following as you configure the VM:

- **VM size:** Choose a VM size that meets CPU and RAM requirements. We recommend Standard\_D8s\_v3.
- **Disks:** The Connector can perform optimally with either HDD or SSD disks.
- **Public IP:** If you want to use a public IP address with the Connector VM, the IP address must use a Basic SKU to ensure that BlueXP uses this public IP address.



The screenshot shows a dialog box titled "Create public IP address" with a close button (X) in the top right corner. The dialog contains the following fields and options:

- Name \***: A text input field containing "newIP" with a green checkmark on the right.
- SKU \***: Radio buttons for "Basic" (selected) and "Standard".
- Assignment**: Radio buttons for "Dynamic" and "Static" (selected).

If you use a Standard SKU IP address instead, then BlueXP uses the *private* IP address of the Connector, instead of the public IP. If the machine that you're using to access the BlueXP Console doesn't have access to that private IP address, then actions from the BlueXP Console will fail.

[Azure documentation: Public IP SKU](#)

- **Network security group:** The Connector requires inbound connections using SSH, HTTP, and HTTPS.

[View security group rules for Azure.](#)

- **Identity:** Under **Management**, select **Enable system assigned managed identity**.

This setting is important because a managed identity allows the Connector virtual machine to identify itself to Microsoft Entra ID without providing any credentials. [Learn more about managed identities for Azure resources.](#)

4. On the **Review + create** page, review your selections and select **Create** to start the deployment.

## Result

Azure deploys the virtual machine with the specified settings. The virtual machine and Connector software should be running in approximately five minutes.

## What's next?

Set up BlueXP.

### Manual install

#### Before you begin

You should have the following:

- Root privileges to install the Connector.
- Details about a proxy server, if a proxy is required for internet access from the Connector.

You have the option to configure a proxy server after installation but doing so requires restarting the Connector.

- A CA-signed certificate, if the proxy server uses HTTPS or if the proxy is an intercepting proxy.



You cannot set a certificate for a transparent proxy server when manually installing the Connector. If you need to set a certificate for a transparent proxy server, you must use the Maintenance Console after installation. Learn more about the [Connector Maintenance Console](#).

- Depending on your operating system, either Podman or Docker Engine is required before you install the Connector.

#### About this task

The installer that is available on the NetApp Support Site might be an earlier version. After installation, the Connector automatically updates itself if a new version is available.

#### Steps

1. If the *http\_proxy* or *https\_proxy* system variables are set on the host, remove them:

```
unset http_proxy
unset https_proxy
```

If you don't remove these system variables, the installation will fail.

2. Download the Connector software from the [NetApp Support Site](#), and then copy it to the Linux host.

You should download the "online" Connector installer that's meant for use in your network or in the cloud. A separate "offline" installer is available for the Connector, but it's only supported with private mode deployments.

3. Assign permissions to run the script.

```
chmod +x BlueXP-Connector-Cloud-<version>
```

Where <version> is the version of the Connector that you downloaded.

4. Run the installation script.

```
./BlueXP-Connector-Cloud-<version> --proxy <HTTP or HTTPS proxy server> --cacert <path and file name of a CA-signed certificate>
```

You'll need to add proxy information if your network requires a proxy for internet access. You can add either a transparent or explicit proxy. The `--proxy` and `--cacert` parameters are optional and you won't be prompted to add them. If you have a proxy server, you will need to enter the parameters as shown.

Here is an example configuring an explicit proxy server with a CA-signed certificate:

```
./BlueXP-Connector-Cloud-v3.9.40--proxy https://user:password@10.0.0.30:8080/ --cacert /tmp/cacert/certificate.cer
```

`--proxy` configures the Connector to use an HTTP or HTTPS proxy server using one of the following formats:

- `http://address:port`
- `http://user-name:password@address:port`
- `http://domain-name%92user-name:password@address:port`
- `https://address:port`
- `https://user-name:password@address:port`
- `https://domain-name%92user-name:password@address:port`

Note the following:

- The user can be a local user or domain user.
- For a domain user, you must use the ASCII code for a \ as shown above.
- BlueXP doesn't support user names or passwords that include the @ character.
- If the password includes any of the following special characters, you must escape that special character by prepending it with a backslash: & or !

For example:

```
http://bxpproxyuser:netapp1!\@address:3128
```

`--cacert` specifies a CA-signed certificate to use for HTTPS access between the Connector and the proxy server. This parameter is required for HTTPS proxy servers, intercepting proxy servers, and transparent proxy servers.

Here is an example configuring a transparent proxy server. When you configure a transparent proxy, you don't need to define the proxy server. You only add a CA-signed certificate to your Connector host:

```
./BlueXP-Connector-Cloud-v3.9.40 --cacert /tmp/cacert/certificate.cer
```



5. If you used Podman, you'll need to adjust the aardvark-dns port.
  - a. SSH to the BlueXP Connector virtual machine.
  - b. Open podman `/usr/share/containers/containers.conf` file and modify the chosen port for Aardvark DNS service. For example, change it to 54.

```
vi /usr/share/containers/containers.conf
...
# Port to use for dns forwarding daemon with netavark in rootful
bridge
# mode and dns enabled.
# Using an alternate port might be useful if other DNS services
should
# run on the machine.
#
dns_bind_port = 54
...
Esc:wq
```

- c. Reboot the Connector virtual machine.

### Result

The Connector is now installed. At the end of the installation, the Connector service (occm) restarts twice if you specified a proxy server.

### What's next?

Set up BlueXP.

## Step 2: Set up BlueXP

When you access the BlueXP console for the first time, you'll be prompted to choose an account to associate the Connector with and you'll need to enable restricted mode.

### Before you begin

The person who sets up the BlueXP Connector must log in to BlueXP using a login that doesn't belong to a BlueXP account or organization.

If your BlueXP login is associated with another account or organization, you'll need to sign up with a new BlueXP login. Otherwise, you won't see the option to enable restricted mode on the setup screen.

### Steps

1. Open a web browser from a host that has a connection to the Connector instance and enter the following URL:

```
https://ipaddress
```

2. Sign up or log in to BlueXP.
3. After you're logged in, set up BlueXP:

- a. Enter a name for the Connector.
- b. Enter a name for a new BlueXP account.
- c. Select **Are you running in a secured environment?**
- d. Select **Enable restricted mode on this account.**

Note that you can't change this setting after BlueXP creates the account. You can't enable restricted mode later and you can't disable it later.

If you deployed the Connector in a Government region, the checkbox is already enabled and can't be changed. This is because restricted mode is the only mode supported in Government regions.

Hi Tami,  
Welcome to NetApp BlueXP

Let's get started by creating an account for your organization.

If your organization already has an existing account, it's best to ask the account admin to add you to it. [Learn how to add user](#)

Connector name: BlueXP1      Account name: MyCompany

Are you running in a secured environment?  Enable restricted mode on this account

[Learn more about BlueXP deployment modes](#)

**Let's start**

- e. Select **Let's start.**

## Result

The Connector is now installed and set up with your BlueXP account. All users need to access BlueXP using the IP address of the Connector instance.

## What's next?

Provide BlueXP with the permissions that you previously set up.

## Step 3: Provide permissions to BlueXP

If you deployed the Connector from the Azure Marketplace or if you manually installed the Connector software, you need to provide the permissions that you previously set up so that you can use BlueXP services.

These steps don't apply if you deployed the Connector from the AWS Marketplace because you chose the required IAM role during deployment.

[Learn how to prepare cloud permissions.](#)

## AWS IAM role

Attach the IAM role that you previously created to the EC2 instance where you installed the Connector.

These steps apply only if you manually installed the Connector in AWS. For AWS Marketplace deployments, you already associated the Connector instance with an IAM role that includes the required permissions.

### Steps

1. Go to the Amazon EC2 console.
2. Select **Instances**.
3. Select the Connector instance.
4. Select **Actions > Security > Modify IAM role**.
5. Select the IAM role and select **Update IAM role**.

### Result

BlueXP now has the permissions that it needs to perform actions in AWS on your behalf.

## AWS access key

Provide BlueXP with the AWS access key for an IAM user that has the required permissions.

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



2. Select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location**: Select **Amazon Web Services > Connector**.
  - b. **Define Credentials**: Enter an AWS access key and secret key.
  - c. **Marketplace Subscription**: Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
  - d. **Review**: Confirm the details about the new credentials and select **Add**.

### Result

BlueXP now has the permissions that it needs to perform actions in AWS on your behalf.

## Azure role

Go to the Azure portal and assign the Azure custom role to the Connector virtual machine for one or more subscriptions.

### Steps

1. From the Azure Portal, open the **Subscriptions** service and select your subscription.

It's important to assign the role from the **Subscriptions** service because this specifies the scope of the role assignment at the subscription level. The *scope* defines the set of resources that the access applies to. If you specify a scope at a different level (for example, at the virtual machine level), your ability to complete actions from within BlueXP will be affected.

2. Select **Access control (IAM) > Add > Add role assignment**.
3. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.



BlueXP Operator is the default name provided in the BlueXP policy. If you chose a different name for the role, then select that name instead.

4. In the **Members** tab, complete the following steps:
  - a. Assign access to a **Managed identity**.
  - b. Select **Select members**, select the subscription in which the Connector virtual machine was created, under **Managed identity**, choose **Virtual machine**, and then select the Connector virtual machine.
  - c. Select **Select**.
  - d. Select **Next**.
  - e. Select **Review + assign**.
  - f. If you want to manage resources in additional Azure subscriptions, switch to that subscription and then repeat these steps.

### Result

BlueXP now has the permissions that it needs to perform actions in Azure on your behalf.

### Azure service principal

Provide BlueXP with the credentials for the Azure service principal that you previously setup.

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



2. Select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location:** Select **Microsoft Azure > Connector**.
  - b. **Define Credentials:** Enter information about the Microsoft Entra service principal that grants the required permissions:
    - Application (client) ID
    - Directory (tenant) ID
    - Client Secret
  - c. **Marketplace Subscription:** Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
  - d. **Review:** Confirm the details about the new credentials and select **Add**.

### Result

BlueXP now has the permissions that it needs to perform actions in Azure on your behalf.

### Google Cloud service account

Associate the service account with the Connector VM.

### Steps

1. Go to the Google Cloud portal and assign the service account to the Connector VM instance.

[Google Cloud documentation: Changing the service account and access scopes for an instance](#)

2. If you want to manage resources in other projects, grant access by adding the service account with the BlueXP role to that project. You'll need to repeat this step for each project.

### Result

BlueXP now has the permissions that it needs to perform actions in Google Cloud on your behalf.

## Subscribe to NetApp Intelligent Services (restricted mode)

Subscribe to NetApp Intelligent Services from your cloud provider's marketplace to pay for data services at an hourly rate (PAYGO) or through an annual contract. If you purchased a license from NetApp (BYOL), you also need to subscribe to the marketplace offering. Your license is always charged first, but you'll be charged at the hourly rate if you exceed your licensed capacity or if the license's term expires.

A marketplace subscription enables charging for the following data services with restricted mode:

- Backup and recovery
- Cloud Volumes ONTAP
- Tiering
- Ransomware protection
- Disaster recovery

Classification is enabled through your subscription, but there is no charge for using classification.

### Before you begin

Subscribing to data services involves associating a marketplace subscription with the cloud credentials that are associated with a Connector. If you followed the "Get started with restricted mode" workflow, then you should already have a Connector. To learn more, view the [Quick start for BlueXP in restricted mode](#).

## AWS

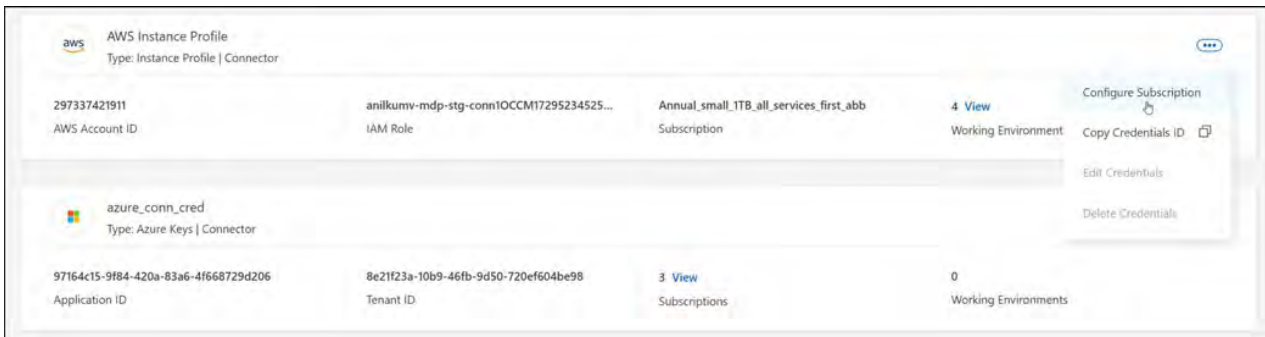
The following video shows the steps to subscribe to NetApp Intelligent Services from the AWS Marketplace:

### [Subscribe to NetApp Intelligent Services from the AWS Marketplace](#)

#### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.
2. Select the action menu for a set of credentials and then select **Configure Subscription**.

You must select credentials that are associated with a Connector. You can't associate a marketplace subscription with credentials that are associated with BlueXP.



3. To associate the credentials with an existing subscription, select the subscription from the down-down list and select **Configure**.
4. To associate the credentials with a new subscription, select **Add Subscription > Continue** and follow the steps in the AWS Marketplace:
  - a. Select **View purchase options**.
  - b. Select **Subscribe**.
  - c. Select **Set up your account**.

You'll be redirected to the BlueXP website.

- d. From the **Subscription Assignment** page:
  - Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
  - In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

- Select **Save**.

## Azure Steps

1. In the upper right of the console, select the Settings icon, and select **Credentials**.
2. Select the action menu for a set of credentials and then select **Configure Subscription**.

You must select credentials that are associated with a Connector. You can't associate a marketplace subscription with credentials that are associated with BlueXP.

3. To associate the credentials with an existing subscription, select the subscription from the down-down list and select **Configure**.
4. To associate the credentials with a new subscription, select **Add Subscription > Continue** and follow the steps in the Azure Marketplace:
  - a. If prompted, log in to your Azure account.
  - b. Select **Subscribe**.
  - c. Fill out the form and select **Subscribe**.
  - d. After the subscription process is complete, select **Configure account now**.

You'll be redirected to BlueXP.

- e. From the **Subscription Assignment** page:

- Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
- In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

- Select **Save**.

The following video shows the steps to subscribe from the Azure Marketplace:

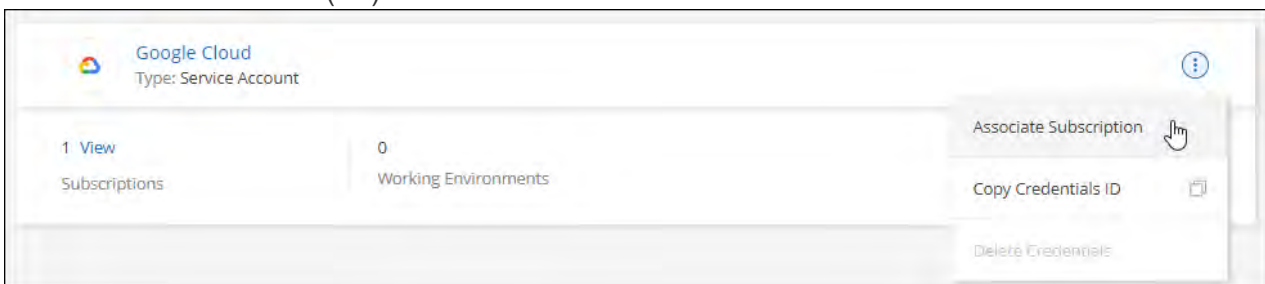
[Subscribe to NetApp Intelligent Services from the Azure Marketplace](#)

## Google Cloud

### Steps


1. In the upper right of the console, select the Settings icon, and select **Credentials**.
2. Select the action menu for a set of credentials and then select **Configure Subscription**.

+new screenshot needed (TS)



3. To configure an existing subscription with the selected credentials, select a Google Cloud project and

subscription from the drop-down list, and then select **Configure**.




4. If you don't already have a subscription, select **Add Subscription > Continue** and follow the steps in the Google Cloud Marketplace.



Before you complete the following steps, ensure that you have both Billing Admin privileges in your Google Cloud account as well as a BlueXP login.

- a. After you're redirected to the [NetApp Intelligent Services page on the Google Cloud Marketplace](#), ensure that the correct project is selected at the top navigation menu.






- b. Select **Subscribe**.
- c. Select the appropriate billing account and agree to the terms and conditions.
- d. Select **Subscribe**.

This step sends your transfer request to NetApp.

- e. On the pop-up dialog box, select **Register with NetApp, Inc.**

This step must be completed to link the Google Cloud subscription with your BlueXP organization or account. The process of linking a subscription isn't complete until you are redirected from this page and then sign in to BlueXP.



- f. Complete the steps on the **Subscription Assignment** page:



If someone from your organization has already subscribed to the NetApp BlueXP subscription from your billing account, then you will be redirected to [the Cloud Volumes ONTAP page on the BlueXP website](#) instead. If this is unexpected, contact your NetApp sales team. Google enables only one subscription per Google billing account.

- Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
- In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

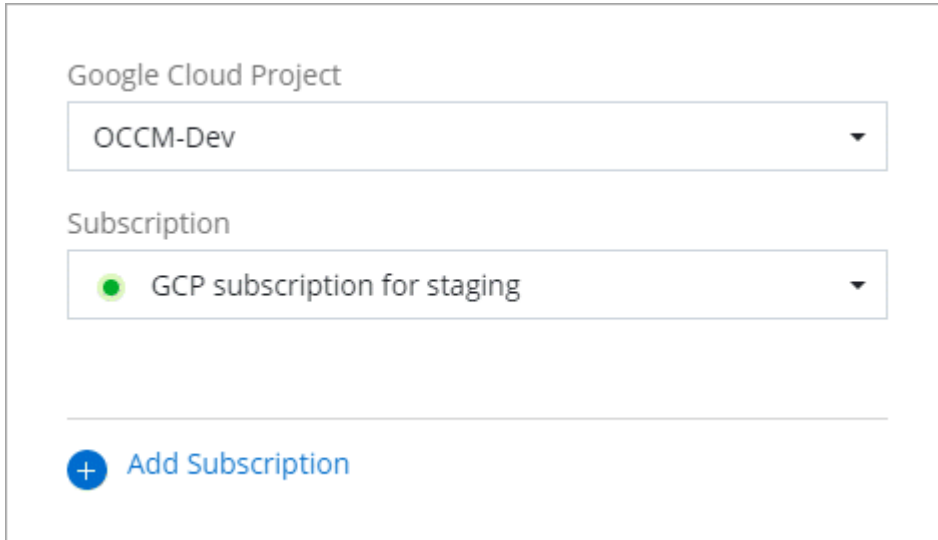
For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

- Select **Save**.

The following video shows the steps to subscribe from the Google Cloud Marketplace:

[Subscribe to BlueXP from the Google Cloud Marketplace](#)

- g. Once this process is complete, navigate back to the Credentials page in BlueXP and select this new subscription.



The screenshot shows a user interface for selecting a Google Cloud Project and a Subscription. The 'Google Cloud Project' dropdown menu is set to 'OCCM-Dev'. The 'Subscription' dropdown menu is set to 'GCP subscription for staging', which is marked with a green dot. Below the dropdowns is a horizontal line, and at the bottom left is a blue button with a plus sign and the text 'Add Subscription'.

#### Related information

- [Manage BYOL capacity-based licenses for Cloud Volumes ONTAP](#)
- [Manage BYOL licenses for data services](#)
- [Manage AWS credentials and subscriptions](#)
- [Manage Azure credentials and subscriptions](#)
- [Manage Google Cloud credentials and subscriptions](#)

#### What you can do next (restricted mode)

After you get up and running with BlueXP in restricted mode, you can start using the BlueXP services that are supported with restricted mode.

For help, refer to the documentation for these services:

- [Azure NetApp Files docs](#)
- [Backup and recovery docs](#)
- [Classification docs](#)
- [Cloud Volumes ONTAP docs](#)
- [Digital wallet docs](#)
- [On-premises ONTAP cluster docs](#)
- [Replication docs](#)

#### Related information

[BlueXP deployment modes](#)

# Get started with private mode

## Getting started workflow (private mode)

Get started with BlueXP in private mode by preparing your environment and deploying the Connector.

Private mode is typically used with on-premises environments that have no internet connection and with secure cloud regions, which includes [AWS Secret Cloud](#), [AWS Top Secret Cloud](#), and [Azure IL6](#)

Before you get started, you should have an understanding of [Connectors](#) and [deployment modes](#).

1

### Prepare for deployment

- a. Prepare a dedicated Linux host that meets requirements for CPU, RAM, disk space, container orchestration tool, and more.
- b. Set up networking that provides access to the target networks.
- c. For cloud deployments, set up permissions in your cloud provider so that you can associate those permissions with the Connector after you install the software.

2

### Deploy the Connector

- a. Install the Connector software on your own Linux host.
- b. Set up BlueXP by opening a web browser and entering the Linux host's IP address.
- c. For cloud deployments, provide BlueXP with the permissions that you previously set up.

## Prepare for deployment in private mode

Prepare your environment before you deploy BlueXP in private mode. For example, you need to review host requirements, prepare networking, set up permissions, and more.



To use BlueXP in the [AWS Secret Cloud](#) or the [AWS Top Secret Cloud](#), follow the specific instructions for those environments. [Learn how to get started with Cloud Volumes ONTAP in the AWS Secret Cloud or Top Secret Cloud](#)

### Step 1: Understand how private mode works

Before you get started, you should understand private mode.

For example, you need to use the browser-based interface that is available locally from the Connector that you install. You can't access BlueXP from the web-based console that's provided through the SaaS layer.

In addition, not all features and services are available.

[Learn how private mode works.](#)

## Step 2: Review installation options

In private mode, you can install the Connector on-premises or in the cloud by manually installing the Connector on your own Linux host.

Where you install the Connector determines which BlueXP services and features are available when using private mode. For example, the Connector must be installed in the cloud if you want to deploy and manage Cloud Volumes ONTAP. [Learn more about private mode.](#)

## Step 3: Review host requirements

The host must meet specific operating system requirements, RAM requirements, port requirements, and so on to run the Connector software.

### Dedicated host

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

Host can be of any architecture that meets the following size requirements:

- CPU: 8 cores or 8 vCPUs
- RAM: 32 GB

### Operating system and container requirements

BlueXP supports the Connector with the following operating systems when using BlueXP in private mode. A container orchestration tool is required before you install the Connector.

Operating system	Supported OS versions	Supported Connector versions	Required container tool	SELinux
Red Hat Enterprise Linux	9.1 to 9.4 8.6 to 8.10	3.9.42 or later with BlueXP in private mode	Podman version 4.6.1 or 4.9.4  <a href="#">View Podman configuration requirements.</a>	Supported in enforcing mode or permissive mode <sup>1</sup>
Ubuntu	22.04 LTS	3.9.29 or later	Docker Engine 23.0.6 to 26.0.0  26.0.0 is supported with <i>new</i> Connector 3.9.44 or later installations	Not supported

Notes:

1. Management of Cloud Volumes ONTAP systems is not supported by Connectors that have SELinux enabled on the operating system.
2. The Connector is supported on English-language versions of these operating systems.
3. For RHEL, the host must be registered with Red Hat Subscription Management. If it's not registered, the

host can't access repositories to update required 3rd-party software during Connector installation.

## Hypervisor

A bare metal or hosted hypervisor that is certified to run a supported operating system is required.

## CPU

8 cores or 8 vCPUs

## RAM

32 GB

## AWS EC2 instance type

An instance type that meets the CPU and RAM requirements above. We recommend t3.2xlarge.

## Azure VM size

An instance type that meets the CPU and RAM requirements above. We recommend Standard\_D8s\_v3.

## Google Cloud machine type

An instance type that meets the CPU and RAM requirements above. We recommend n2-standard-8.

The Connector is supported in Google Cloud on a VM instance with an OS that supports [Shielded VM features](#)

## Disk space in /opt

100 GiB of space must be available

BlueXP uses `/opt` to install the `/opt/application/netapp` directory and its contents.

## Disk space in /var

20 GiB of space must be available

BlueXP requires this space in `/var` because Docker or Podman are architected to create the containers within this directory. Specifically, they will create containers in the `/var/lib/containers/storage` directory. External mounts or symlinks do not work for this space.

## Step 4: Install Podman or Docker Engine

You need to prepare the host for the Connector by installing Podman or Docker Engine.

Depending on your operating system, either Podman or Docker Engine is required before you install the Connector.

- Podman is required for Red Hat Enterprise Linux 8 and 9.

[View the Podman versions that BlueXP supports.](#)

- Docker Engine is required for Ubuntu.

[View the Docker Engine versions that BlueXP supports.](#)

## Example 6. Steps

### Podman

Follow these steps to install Podman and configure it to meet the following requirements:

- Enable and start the podman.socket service
- Install python3
- Install the podman-compose package version 1.0.6
- Add podman-compose to the PATH environment variable



When using Podman, adjust the aardvark-dns service port (default: 53) after installing the Connector to avoid conflicts with the DNS port on the host. Follow the instructions to configure the port.

### Steps

1. Remove the podman-docker package if it's installed on the host.

```
dnf remove podman-docker
rm /var/run/docker.sock
```

2. Install Podman.

Podman is available from official Red Hat Enterprise Linux repositories.

For Red Hat Enterprise Linux 9:

```
sudo dnf install podman-2:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

For Red Hat Enterprise Linux 8:

```
sudo dnf install podman-3:<version>
```

Where <version> is the supported version of Podman that you're installing. [View the Podman versions that BlueXP supports.](#)

3. Enable and start the podman.socket service.

```
sudo systemctl enable --now podman.socket
```

4. Install python3.

```
sudo dnf install python3
```

5. Install the EPEL repository package if it's not already available on your system.

This step is required because podman-compose is available from the Extra Packages for Enterprise Linux (EPEL) repository.

For Red Hat Enterprise Linux 9:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-9.noarch.rpm
```

For Red Hat Enterprise Linux 8:

```
sudo dnf install https://dl.fedoraproject.org/pub/epel/epel-release-latest-8.noarch.rpm
```

6. Install podman-compose package 1.0.6.

```
sudo dnf install podman-compose-1.0.6
```



Using the `dnf install` command meets the requirement for adding podman-compose to the PATH environment variable. The installation command adds podman-compose to `/usr/bin`, which is already included in the `secure_path` option on the host.

### Docker Engine

Follow the documentation from Docker to install Docker Engine.

#### Steps

1. [View installation instructions from Docker](#)

Be sure to follow the steps to install a specific version of Docker Engine. Installing the latest version will install a version of Docker that BlueXP doesn't support.

2. Verify that Docker is enabled and running.

```
sudo systemctl enable docker && sudo systemctl start docker
```

## Step 5: Prepare networking

Set up networking for the Connector to manage resources in your public cloud. Other than having a virtual network and subnet for the Connector, ensure that the following requirements are met.

Connections to target networks::

The Connector must have a network connection to the location where you plan to manage storage. For example, the VPC or VNet where you plan to deploy Cloud Volumes ONTAP, or the data center where your on-premises ONTAP clusters reside.

### Endpoints for day-to-day operations

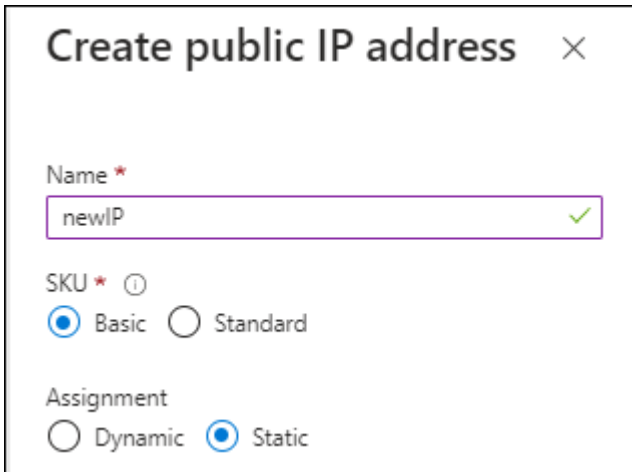
If you are planning to create Cloud Volumes ONTAP systems, the Connector needs connectivity to endpoints in your cloud provider's publicly available resources.

Endpoints	Purpose
<p>AWS services (amazonaws.com):</p> <ul style="list-style-type: none"> <li>• CloudFormation</li> <li>• Elastic Compute Cloud (EC2)</li> <li>• Identity and Access Management (IAM)</li> <li>• Key Management Service (KMS)</li> <li>• Security Token Service (STS)</li> <li>• Simple Storage Service (S3)</li> </ul>	<p>To manage resources in AWS. The exact endpoint depends on the AWS region that you're using. <a href="#">Refer to AWS documentation for details</a></p>
<p>https://management.azure.com            https://login.microsoftonline.com            https://blob.core.windows.net            https://core.windows.net</p>	<p>To manage resources in Azure public regions.</p>
<p>https://management.azure.microsoft.scloud            https://login.microsoftonline.microsoft.scloud            https://blob.core.microsoft.scloud            https://core.microsoft.scloud</p>	<p>To manage resources in the Azure IL6 region.</p>
<p>https://management.chinacloudapi.cn            https://login.chinacloudapi.cn            https://blob.core.chinacloudapi.cn            https://core.chinacloudapi.cn</p>	<p>To manage resources in Azure China regions.</p>
<p>https://www.googleapis.com/compute/v1/            https://compute.googleapis.com/compute/v1            https://cloudresourcemanager.googleapis.com/v1/projects            https://www.googleapis.com/compute/beta            https://storage.googleapis.com/storage/v1            https://www.googleapis.com/storage/v1            https://iam.googleapis.com/v1            https://cloudkms.googleapis.com/v1            https://www.googleapis.com/deploymentmanager/v2/projects</p>	<p>To manage resources in Google Cloud.</p>



## Public IP address in Azure

If you want to use a public IP address with the Connector VM in Azure, the IP address must use a Basic SKU to ensure that BlueXP uses this public IP address.



The screenshot shows a 'Create public IP address' dialog box. It has a title bar with a close button (X). The 'Name' field is labeled with a red asterisk and contains the text 'newIP' with a green checkmark to its right. Below the name field, the 'SKU' section is labeled with a red asterisk and a help icon (i). It has two radio buttons: 'Basic' (selected) and 'Standard'. Below that, the 'Assignment' section has two radio buttons: 'Dynamic' and 'Static' (selected).

If you use a Standard SKU IP address instead, then BlueXP uses the *private* IP address of the Connector, instead of the public IP. If the machine that you're using to access the BlueXP Console doesn't have access to that private IP address, then actions from the BlueXP Console will fail.

[Azure documentation: Public IP SKU](#)

## Proxy server

NetApp supports both explicit and transparent proxy configurations. If you are using a transparent proxy, you only need to provide the certificate for the proxy server. If you are using an explicit proxy, you'll also need the IP address and credentials.

- IP address
- Credentials
- HTTPS certificate

With private mode, the only time that BlueXP sends outbound traffic is to your cloud provider in order to create a Cloud Volumes ONTAP system.

## Ports

There's no incoming traffic to the Connector, unless you initiate it.

HTTP (80) and HTTPS (443) provide access to the BlueXP console. SSH (22) is only needed if you need to connect to the host for troubleshooting.

## Enable NTP

If you're planning to use BlueXP classification to scan your corporate data sources, you should enable a Network Time Protocol (NTP) service on both the BlueXP Connector system and the BlueXP classification system so that the time is synchronized between the systems. [Learn more about BlueXP classification](#)

## Step 6: Prepare cloud permissions

If the Connector is installed in the cloud and you plan to create Cloud Volumes ONTAP systems, BlueXP requires cloud provider permissions. You need to set up permissions in your cloud provider and then associate those permission with the Connector instance after you install it.

To view the required steps, select the authentication option that you'd like to use for your cloud provider.

## AWS IAM role

Use an IAM role to provide the Connector with permissions. You'll need to manually attach the role to the EC2 instance for the Connector.

### Steps

1. Log in to the AWS console and navigate to the IAM service.
2. Create a policy:
  - a. Select **Policies > Create policy**.
  - b. Select **JSON** and copy and paste the contents of the [IAM policy for the Connector](#).
  - c. Finish the remaining steps to create the policy.
3. Create an IAM role:
  - a. Select **Roles > Create role**.
  - b. Select **AWS service > EC2**.
  - c. Add permissions by attaching the policy that you just created.
  - d. Finish the remaining steps to create the role.

### Result

You now have an IAM role for the Connector EC2 instance.

## AWS access key

Set up permissions and an access key for an IAM user. Provide BlueXP with the AWS access key after you install the Connector and set up BlueXP.

### Steps

1. Log in to the AWS console and navigate to the IAM service.
2. Create a policy:
  - a. Select **Policies > Create policy**.
  - b. Select **JSON** and copy and paste the contents of the [IAM policy for the Connector](#).
  - c. Finish the remaining steps to create the policy.

Depending on the BlueXP services that you're planning to use, you might need to create a second policy.

For standard regions, the permissions are spread across two policies. Two policies are required due to a maximum character size limit for managed policies in AWS. [Learn more about IAM policies for the Connector](#).

3. Attach the policies to an IAM user.
  - [AWS Documentation: Creating IAM Roles](#)
  - [AWS Documentation: Adding and Removing IAM Policies](#)
4. Ensure that the user has an access key that you can add to BlueXP after you install the Connector.

### Result

The account now has the required permissions.

## Azure role

Create an Azure custom role with the required permissions. Assign this role to the Connector VM.

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would prefer to use a different method, refer to [Azure documentation](#)

## Steps

1. Enable a system-assigned managed identity on the VM where you plan to install the Connector so that you can provide the required Azure permissions through a custom role.

[Microsoft Azure documentation: Configure managed identities for Azure resources on a VM using the Azure portal](#)

2. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
3. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription that you want to use with BlueXP.


## Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- a. Start [Azure Cloud Shell](#) and choose the Bash environment.
- b. Upload the JSON file.



c. Use the Azure CLI to create the custom role:

```
az role definition create --role-definition Connector_Policy.json
```

## Result

You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.

## Azure service principal

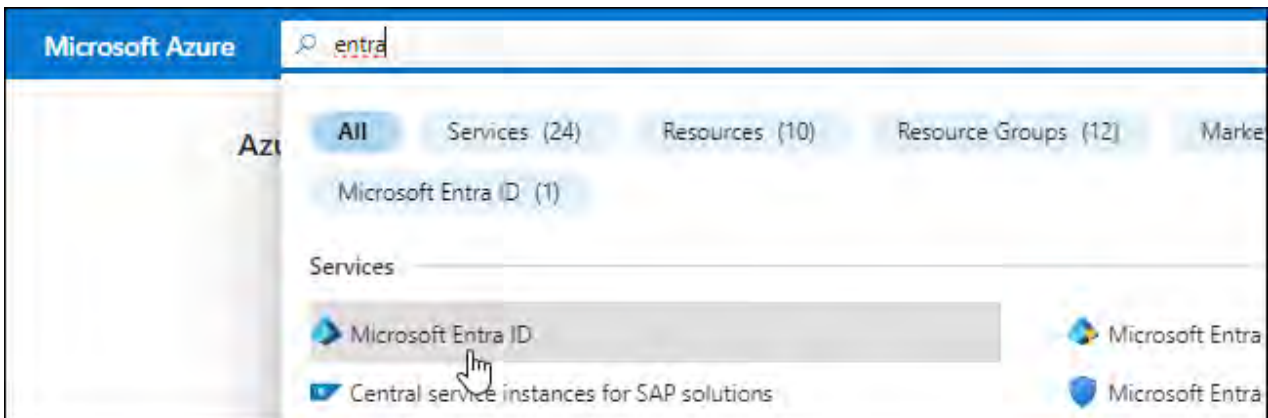
Create and set up a service principal in Microsoft Entra ID and obtain the Azure credentials that BlueXP needs. You'll need to provide BlueXP with these credentials after you install the Connector and set up BlueXP.

### Create a Microsoft Entra application for role-based access control

1. Ensure that you have permissions in Azure to create an Active Directory application and to assign the application to a role.

For details, refer to [Microsoft Azure Documentation: Required permissions](#)

2. From the Azure portal, open the **Microsoft Entra ID** service.



3. In the menu, select **App registrations**.
4. Select **New registration**.
5. Specify details about the application:
  - **Name**: Enter a name for the application.
  - **Account type**: Select an account type (any will work with BlueXP).
  - **Redirect URI**: You can leave this field blank.
6. Select **Register**.

You've created the AD application and service principal.

### Assign the application to a role

1. Create a custom role:

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI,

or REST API. The following steps show how to create the role using the Azure CLI. If you would prefer to use a different method, refer to [Azure documentation](#)

- a. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
- b. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription from which users will create Cloud Volumes ONTAP systems.

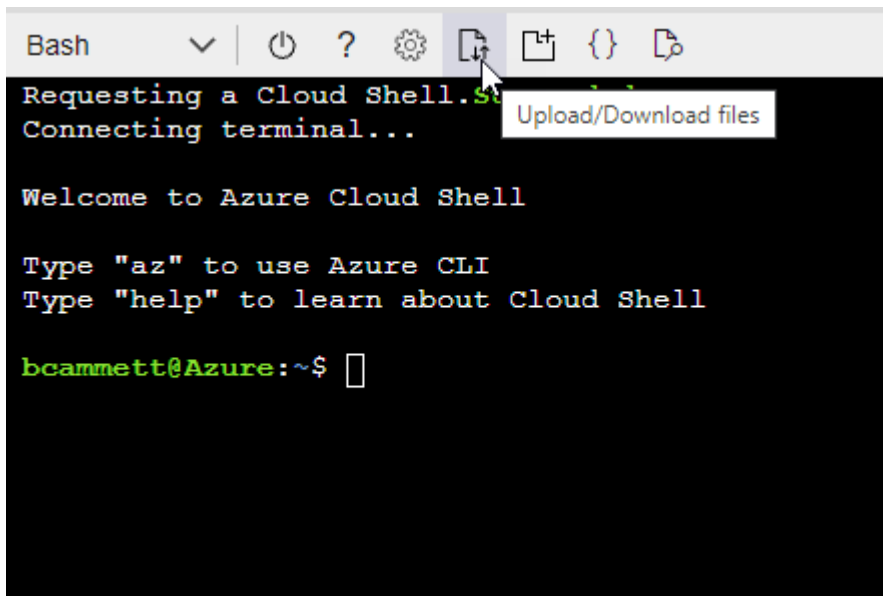
### Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- Start [Azure Cloud Shell](#) and choose the Bash environment.
- Upload the JSON file.

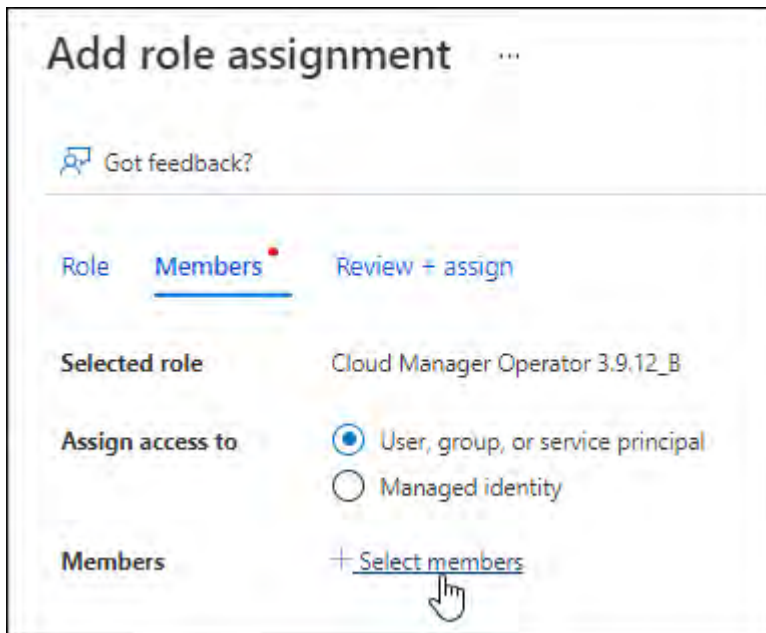


- Use the Azure CLI to create the custom role:

```
az role definition create --role-definition  
Connector_Policy.json
```

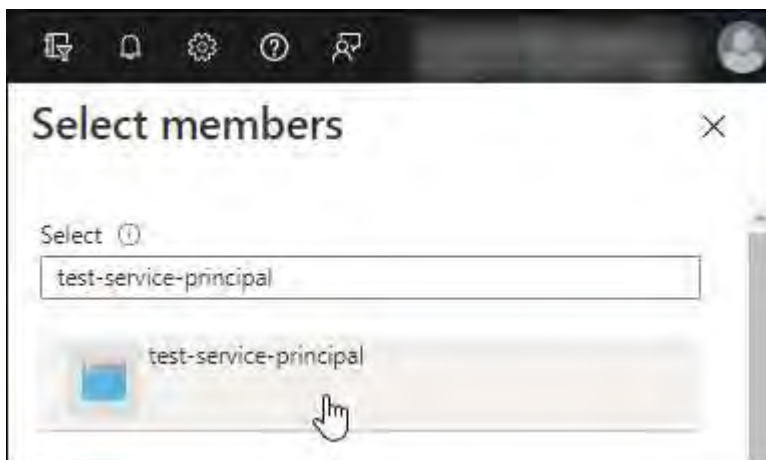
You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.

2. Assign the application to the role:
  - a. From the Azure portal, open the **Subscriptions** service.
  - b. Select the subscription.
  - c. Select **Access control (IAM) > Add > Add role assignment**.
  - d. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.
  - e. In the **Members** tab, complete the following steps:
    - Keep **User, group, or service principal** selected.
    - Select **Select members**.



- Search for the name of the application.

Here's an example:



- Select the application and select **Select**.
  - Select **Next**.
- f. Select **Review + assign**.

The service principal now has the required Azure permissions to deploy the Connector.

If you want to deploy Cloud Volumes ONTAP from multiple Azure subscriptions, then you must bind the service principal to each of those subscriptions. BlueXP enables you to select the subscription that you want to use when deploying Cloud Volumes ONTAP.

### Add Windows Azure Service Management API permissions

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Select **API permissions > Add a permission**.
3. Under **Microsoft APIs**, select **Azure Service Management**.

#### Request API permissions

Select an API

Microsoft APIs   **APIs my organization uses**   My APIs

Commonly used Microsoft APIs

Commonly used Microsoft APIs		
<b>Microsoft Graph</b> Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10, Access Azure AD, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint.		
<b>Azure Batch</b> Schedule large-scale parallel and HPC applications in the cloud	<b>Azure Data Catalog</b> Programmatic access to Data Catalog resources to register, annotate and search data assets	<b>Azure Data Explorer</b> Perform ad-hoc queries on terabytes of data to build near real-time and complex analytics solutions
<b>Azure Data Lake</b> Access to storage and compute for big data analytic scenarios	<b>Azure DevOps</b> Integrate with Azure DevOps and Azure DevOps server	<b>Azure Import/Export</b> Programmatic control of import/export jobs
<b>Azure Key Vault</b> Manage your key vaults as well as the keys, secrets, and certificates within your Key Vaults	<b>Azure Rights Management Services</b> Allow validated users to read and write protected content	<b>Azure Service Management</b> Programmatic access to much of the functionality available through the Azure portal
<b>Azure Storage</b> Secure, massively scalable object and data lake storage for unstructured and semi-structured data	<b>Customer Insights</b> Create profile and interaction models for your products	<b>Data Export Service for Microsoft Dynamics 365</b> Export data from Microsoft Dynamics CRM organization to an external destination

4. Select **Access Azure Service Management as organization users** and then select **Add permissions**.



## Request API permissions

[< All APIs](#)



Azure Service Management

<https://management.azure.com/> [Docs](#)

What type of permissions does your application require?

### Delegated permissions

Your application needs to access the API as the signed-in user.

### Application permissions

Your application runs as a background service or daemon without a signed-in user.

Select permissions

[expand all](#)

Type to search

PERMISSION

ADMIN CONSENT REQUIRED




user\_impersonation

Access Azure Service Management as organization users (preview)

## Get the application ID and directory ID for the application

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Copy the **Application (client) ID** and the **Directory (tenant) ID**.



When you add the Azure account to BlueXP, you need to provide the application (client) ID and the directory (tenant) ID for the application. BlueXP uses the IDs to programmatically sign in.

## Create a client secret

1. Open the **Microsoft Entra ID** service.
2. Select **App registrations** and select your application.
3. Select **Certificates & secrets > New client secret**.
4. Provide a description of the secret and a duration.
5. Select **Add**.
6. Copy the value of the client secret.

## Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

[+ New client secret](#)

DESCRIPTION	EXPIRES	VALUE	Copy to clipboard
test secret	8/16/2020	*sZ1jSe2By:D*-ZRov4NLfdAcY7:+0vA	

You now have a client secret that BlueXP can use it to authenticate with Microsoft Entra ID.

### Result

Your service principal is now setup and you should have copied the application (client) ID, the directory (tenant) ID, and the value of the client secret. Enter this information in BlueXP when you add an Azure account.

### Google Cloud service account

Create a role and apply it to a service account that you'll use for the Connector VM instance.

### Steps

1. Create a custom role in Google Cloud:
  - a. Create a YAML file that includes the permissions defined in the [Connector policy for Google Cloud](#).
  - b. From Google Cloud, activate cloud shell.
  - c. Upload the YAML file that includes the required permissions for the Connector.
  - d. Create a custom role by using the `gcloud iam roles create` command.

The following example creates a role named "connector" at the project level:

```
gcloud iam roles create connector --project=myproject
--file=connector.yaml
```

[Google Cloud docs: Creating and managing custom roles](#)

2. Create a service account in Google Cloud:
  - a. From the IAM & Admin service, select **Service Accounts > Create Service Account**.
  - b. Enter service account details and select **Create and Continue**.
  - c. Select the role that you just created.
  - d. Finish the remaining steps to create the role.

[Google Cloud docs: Creating a service account](#)

### Result

You now have a service account that you can assign to the Connector VM instance.

## Step 7: Enable Google Cloud APIs

You need to enable several APIs to deploy Cloud Volumes ONTAP in Google Cloud.

### Step

#### 1. [Enable the following Google Cloud APIs in your project](#)

- Cloud Deployment Manager V2 API
- Cloud Logging API
- Cloud Resource Manager API
- Compute Engine API
- Identity and Access Management (IAM) API
- Cloud Key Management Service (KMS) API

(Required only if you are planning to use BlueXP backup and recovery with customer-managed encryption keys (CMEK))

## Deploy the Connector in private mode

Deploy the Connector in private mode so that you can use BlueXP with no outbound connectivity to the BlueXP software as a service (SaaS) layer. To get started, install the Connector, set up BlueXP by accessing the user interface that's running on the Connector, and then provide the cloud permissions that you previously set up.

### Step 1: Install the Connector

Download the product installer from the [NetApp Support Site](#) and then manually install the Connector on your own Linux host.

If you want to use BlueXP in the [AWS Secret Cloud](#) or the [AWS Top Secret Cloud](#), then you should follow separate instructions to get started in those environments. [Learn how to get started with Cloud Volumes ONTAP in the AWS Secret Cloud or Top Secret Cloud](#)

### Before you begin

- Root privileges are required to install the Connector.
- Depending on your operating system, either Podman or Docker Engine is required before you install the Connector.

### Steps

#### 1. Download the Connector software from the [NetApp Support Site](#)

Be sure to download the offline installer for private networks without internet access.

2. Copy the installer to the Linux host.
3. Assign permissions to run the script.

```
chmod +x /path/BlueXP-Connector-offline-<version>
```

Where <version> is the version of the Connector that you downloaded.

4. Run the installation script:

```
sudo /path/BlueXP-Connector-offline-<version>
```

Where <version> is the version of the Connector that you downloaded.

### Result

The Connector software is installed. You can now set up BlueXP.

### Step 2: Set up BlueXP

When you access the BlueXP console for the first time, you'll be prompted to set up BlueXP.

### Steps

1. Open a web browser and enter `https://ipaddress` where *ipaddress* is the IP address of the Linux host where you installed the Connector.

You should see the following screen.



2. Select **Set Up New BlueXP Connector** and follow the prompts to set up the system.
  - **System Details:** Enter a name for the Connector and your company name.

1 System Details 2 Create Admin User 3 Review

## System Details

To help us provide better support, enter a name for BlueXP Connector and your company name.

BlueXP Connector Name

Company Name

- **Create an Admin User:** Create the admin user for the system.

This user account runs locally on the system. There's no connection to the auth0 service available through BlueXP.

- **Review:** Review the details, accept the license agreement, and then select **Set Up**.

3. Log in to BlueXP using the admin user that you just created.

### Result

The Connector is now installed and set up.

When new versions of the Connector software are available, they'll be posted to the NetApp Support Site. [Learn how to upgrade the Connector.](#)

### What's next?

Provide BlueXP with the permissions that you previously set up.

### Step 3: Provide permissions to BlueXP

If you want to create Cloud Volumes ONTAP working environments, you'll need to provide BlueXP with the cloud permissions that you previously set up.

[Learn how to prepare cloud permissions.](#)

### AWS IAM role

Attach the IAM role that you previously created to the Connector EC2 instance.

#### Steps

1. Go to the Amazon EC2 console.
2. Select **Instances**.
3. Select the Connector instance.
4. Select **Actions > Security > Modify IAM role**.
5. Select the IAM role and select **Update IAM role**.

#### Result

BlueXP now has the permissions that it needs to perform actions in AWS on your behalf.

### AWS access key

Provide BlueXP with the AWS access key for an IAM user that has the required permissions.

#### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



2. Select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location**: Select **Amazon Web Services > Connector**.
  - b. **Define Credentials**: Enter an AWS access key and secret key.
  - c. **Marketplace Subscription**: Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
  - d. **Review**: Confirm the details about the new credentials and select **Add**.

#### Result

BlueXP now has the permissions that it needs to perform actions in AWS on your behalf.

### Azure role

Go to the Azure portal and assign the Azure custom role to the Connector virtual machine for one or more subscriptions.

#### Steps

1. From the Azure Portal, open the **Subscriptions** service and select your subscription.

It's important to assign the role from the **Subscriptions** service because this specifies the scope of the role assignment at the subscription level. The *scope* defines the set of resources that the access applies to. If you specify a scope at a different level (for example, at the virtual machine level), your ability to complete actions from within BlueXP will be affected.

[Microsoft Azure documentation: Understand scope for Azure RBAC](#)

2. Select **Access control (IAM) > Add > Add role assignment**.

3. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.



BlueXP Operator is the default name provided in the BlueXP policy. If you chose a different name for the role, then select that name instead.

4. In the **Members** tab, complete the following steps:

- a. Assign access to a **Managed identity**.
- b. Select **Select members**, select the subscription in which the Connector virtual machine was created, under **Managed identity**, choose **Virtual machine**, and then select the Connector virtual machine.
- c. Select **Select**.
- d. Select **Next**.
- e. Select **Review + assign**.
- f. If you want to manage resources in additional Azure subscriptions, switch to that subscription and then repeat these steps.

### Result

BlueXP now has the permissions that it needs to perform actions in Azure on your behalf.

### Azure service principal

Provide BlueXP with the credentials for the Azure service principal that you previously setup.

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



2. Select **Add Credentials** and follow the steps in the wizard.

- a. **Credentials Location:** Select **Microsoft Azure > Connector**.
- b. **Define Credentials:** Enter information about the Microsoft Entra service principal that grants the required permissions:
  - Application (client) ID
  - Directory (tenant) ID
  - Client Secret
- c. **Marketplace Subscription:** Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
- d. **Review:** Confirm the details about the new credentials and select **Add**.

### Result

BlueXP now has the permissions that it needs to perform actions in Azure on your behalf.

### Google Cloud service account

Associate the service account with the Connector VM.

### Steps

1. Go to the Google Cloud portal and assign the service account to the Connector VM instance.

[Google Cloud documentation: Changing the service account and access scopes for an instance](#)

2. If you want to manage resources in other projects, grant access by adding the service account with the BlueXP role to that project. You'll need to repeat this step for each project.

**Result**

BlueXP now has the permissions that it needs to perform actions in Google Cloud on your behalf.

**What you can do next (private mode)**

After you get up and running with BlueXP in private mode, you can start using the BlueXP services that are supported with private mode.

For help, refer to the following documentation:

- [Discover on-premises ONTAP clusters](#)
- [Manage software updates](#)
- [Scan on-premises ONTAP volume data using BlueXP classification](#)
- [Monitor license usage with digital wallet](#)
- [View storage health information with digital advisor](#)



# Use BlueXP

## Log in to BlueXP

How you log in to BlueXP depends on the BlueXP deployment mode that you're using for your account.

[Learn about BlueXP deployment modes.](#)

## Standard mode

After you sign up to BlueXP, you can log in from the web-based console to start managing your data and storage infrastructure.

### About this task

You can log in to the BlueXP web-based console using one of the following options:

- Your existing NetApp Support Site (NSS) credentials
- A NetApp cloud login using your email address and a password
- A federated connection

You can use single sign-on to log in using credentials from your corporate directory (federated identity). [Learn how to use identity federation with BlueXP.](#)

### Steps

1. Open a web browser and go to the [BlueXP console](#)
2. On the **Log in** page, enter the email address that's associated with your login.
3. Depending on the authentication method associated with your login, you'll be prompted to enter your credentials:
  - NetApp cloud credentials: Enter your password
  - Federated user: Enter your federated identity credentials
  - NetApp Support Site account: Enter your NetApp Support Site credentials

### Result

You're now logged in and can start using BlueXP to manage your hybrid multi-cloud infrastructure.

## Restricted mode

When you use BlueXP in restricted mode, you need to log in to the BlueXP console from the user interface that runs locally on the Connector.

### About this task

BlueXP supports logging in with one of the following options when your account is set up in restricted mode:

- A NetApp cloud login using your email address and a password
- A federated connection

You can use single sign-on to log in using credentials from your corporate directory (federated identity). [Learn how to use identity federation with BlueXP.](#)

### Steps

1. Open a web browser and enter the following URL:

`https://ipaddress`

*ipaddress* can be localhost, a private IP address, or a public IP address, depending on the configuration of the host where you installed the Connector. For example, you might need to enter a private IP address from a host that has a connection to the Connector host.

2. Enter your user name and password to log in.

### Result

You're now logged in and can start using BlueXP to manage your hybrid multi-cloud infrastructure.

### Private mode

When you use BlueXP in private mode, you need to log in to the BlueXP console from the user interface that runs locally on the Connector.

### About this task

Private mode supports local user management and access. Authentication is not provided through BlueXP's cloud service.

### Steps

1. Open a web browser and enter the following URL:

`https://ipaddress`

*ipaddress* can be localhost, a private IP address, or a public IP address, depending on the configuration of the host where you installed the Connector. For example, you might need to enter a private IP address from a host that has a connection to the Connector host.

2. Enter your user name and password to log in.

### Result

You're now logged in and can start using BlueXP to manage your hybrid multi-cloud infrastructure.

## Manage your BlueXP user settings

You can modify your BlueXP profile including change your password, enable multi-factor authentication (MFA), and see who your BlueXP administrator is.

Within BlueXP, each user has a profile that contains information about the user and their settings. You can view and edit your profile settings.

### Change your display name

You can change your display name. The display name is used to identify you in the BlueXP console and is visible to other users. Your display name is not the same as your username or email address, which cannot be changed.

#### Steps

1. Select the profile icon in the upper right corner of the BlueXP console to view the User settings panel.
2. Select the **Edit** icon next to your name.
3. Enter your new display name in the **Name** field.

### Configure multi-factor authentication

Configure multi-factor authentication (MFA) to enhance account security by requiring a second verification method with your password.

Users using single-sign on with an external identity provider or the NetApp Support Site cannot enable MFA. If either of these are true for you, you won't see the option to enable MFA in your profile settings.

Do not enable MFA if your user account is for BlueXP API access. Multi-factor authentication stops API access when enabled for a user account. Use service accounts for all API access.

### Before you begin

- You must have already downloaded an authentication app, such as Google Authenticator or Microsoft Authenticator, to your device.
- You'll need your password to set up MFA.



If you do not have access to your authentication app or lose your recovery code, contact your BlueXP administrator for help.

### Steps

1. Select the profile icon in the upper right corner of the BlueXP console to view the User settings panel.
2. Select **Configure** next to the **Multi-Factor Authentication** header.
3. Follow the prompts to set up MFA for your account.
4. When you finish, you'll be prompted to save your recovery code. Choose to either copy the code or download a text file containing the code. Keep this code somewhere safe. You need the recovery code if you lose access to your authentication app.

After you set up MFA, you are prompted to enter a one-time code from your authentication app each time you log in to BlueXP.

## Regenerate your MFA recovery code

You can only use recovery codes once. If you use or lose yours, create a new one.

### Steps

1. Select the profile icon in the upper right corner of the BlueXP console to view the User settings panel.
2. Select **...** next to the **Multi-Factor Authentication** header.
3. Select **Regenerate recovery code**.
4. Copy the generated recovery code and save it in a secure location.

## Delete your MFA configuration

To stop using multi-factor authentication (MFA) for your BlueXP account, delete your MFA configuration. This removes the need to enter a one-time code from your authentication app when you log in.



If you are unable to access your authentication app or recovery code, you will need to contact your BlueXP administrator to reset your MFA configuration.

### Steps

1. Select the profile icon in the upper right corner of the BlueXP console to view the User settings panel.
2. Select **...** next to the **Multi-Factor Authentication** header.
3. Select **Delete**.

## Contact your Organization administrator

If you need to contact your organization administrator, you can send an email to them directly from BlueXP. The administrator manages user accounts and permissions within your organization.



You must have a default email application configured for your browser to use the **Contact admins** feature.

### Steps

1. Select the profile icon in the upper right corner of the BlueXP console to view the User settings panel.
2. Select **Contact admins** to send an email to your organization administrator.
3. Select the email application to use.
4. Finish the email and select **Send**.

## Configure dark mode (dark theme)

You can set BlueXP to display in dark mode.

### Steps

1. Select the profile icon in the upper right corner of the BlueXP console to view the User settings panel.
2. Move the **Dark theme** slider to enable it.

# Administer BlueXP

## Identity and access management

### Learn about BlueXP identity and access management

BlueXP identity and access management (IAM) enables you to organize and control access to your NetApp resources. You can organize your resources according to your organization's hierarchy. For example, you can organize resources by geographical location, site, or business unit. You can then assign IAM roles to members at specific parts of the hierarchy, which prevents access to resources in other parts of the hierarchy.

- [Learn about BlueXP deployment modes](#)

### How BlueXP IAM works

BlueXP IAM lets you grant resource access by assigning users access roles to specific parts of the hierarchy. For example, a member can be assigned the Folder or project admin role for a project with five resources.

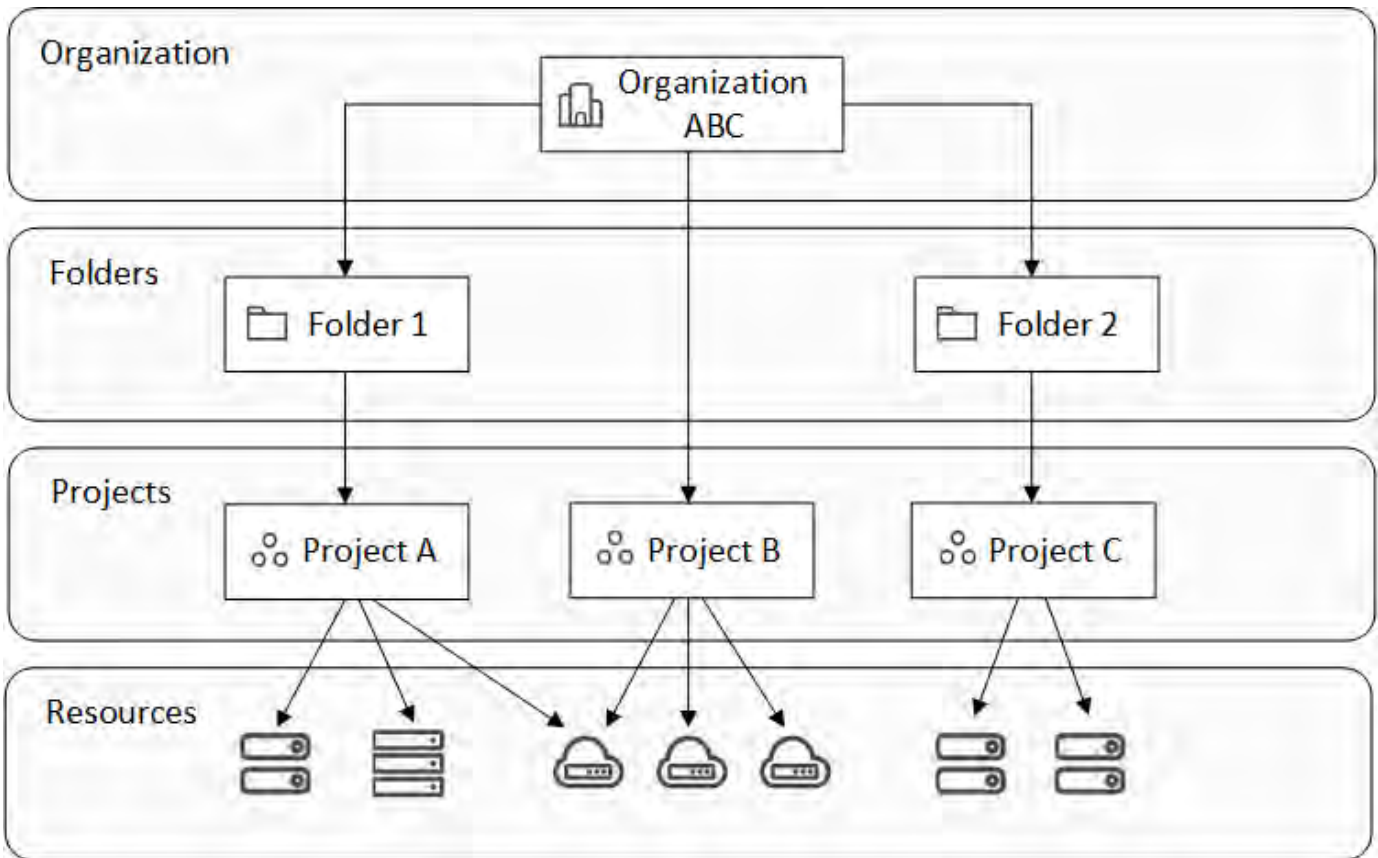
When using BlueXP IAM, you'll manage the following components:

- The organization
- Folders
- Projects
- Resources
- Members
- Roles and permissions
- Connectors

BlueXP resources are organized hierarchically:

- The organization is the top of the hierarchy.
- Folders are children of the organization or of another folder.
- Projects are children of the organization or of a folder.
- Resources are associated with one or more folders or projects.

The following image illustrates this hierarchy at a basic level.



## Organization

An *organization* is the top level of BlueXP's IAM system and typically represents your company. Your organization consists of folders, projects, members, roles, and resources. Connectors are associated with specific projects in the organization.

## Folders

A *folder* enables you to group related projects together and separate them from other projects in your organization. For example, a folder might represent a geographical location (EU or US East), a site (London or Toronto), or a business unit (engineering or marketing).

olders can contain projects, other folders, or both. Creating folders is optional.

## Projects

A *project* represents a workspace in BlueXP that organization members access from the BlueXP canvas in order to manage resources. For example, a project can include a Cloud Volumes ONTAP system, an on-premises ONTAP cluster, or an FSx for ONTAP file system.

An organization can have one or many projects. A project can reside directly underneath the organization or within a folder.

## Resources

A *resource* is a working environment that you created or discovered in BlueXP.

When you create or discover a resource, the resource is associated with the currently selected project. That might be the only project that you want to associate this resource with. But you can choose to associate the

resource with additional projects in your organization.

For example, you might associate a Cloud Volumes ONTAP system with one additional project or with all projects in your organization. How you associate a resource depends on your organization's needs.




You can also associate a Connector with another folder or project in your organization. [Learn more about using Connectors with BlueXP IAM.](#)

### When to associate a resource with a folder


You also have the option to associate a resource with a folder, but this is optional and meets the needs of a specific use case.

An *Organization administrator* might associate a resource with a folder to allow a *Folder or project administrator* to link that resource to the appropriate projects in the folder.

For example, let's say you have a folder that contains two projects:




The *Organization admin* can associate a resource with the folder:



Associating a resource with a folder doesn't make it accessible to all projects; only the *Folder or project admin* can see it. The *Folder or project admin* decides which projects can access it and associates the resource with the appropriate projects.

In this example, the admin associates the resource with Project A:





Members who have permissions for project A can now access the resource.

### Members

Members of your organization are user accounts or service accounts. A service account is typically used by an application to complete specified tasks without human intervention.

Each organization includes at least one user with the *Organization admin* role (BlueXP automatically assigns this role to the user who creates the organization). You can add other members to the organization and assign different permissions across different levels of the resource hierarchy.

### Roles and permissions

In BlueXP IAM, you don't grant permissions directly to organization members. Instead, you grant each member a role. A role contains a set of permissions that enables a member to perform specific actions at a specific level of the resource hierarchy.

Granting permissions at a specific hierarchy level restricts access to the resources a member needs and the services that they can use with those resources.

### Where you can assign roles in the hierarchy

When you associate a member with a role, you need to select the entire organization, a specific folder, or a specific project. The role that you select gives a member permissions to the resources in the selected part of the hierarchy.

### Role inheritance

When you assign a role, the role is inherited down the organization hierarchy:

### Organization

Granting a member an access role at the organization level gives them permissions to all folders, projects, and resources.

### Folders

When you grant an access role at the folder level, all folders, projects, and resources in the folder inherit that role.

For example, if you assign a role at the folder level and that folder has three projects, the member will have permissions to those three projects and any associated resources.

## Projects

When you grant an access role at the project level, all resources associated with that project inherit that role.

## Multiple roles

You can assign each organization member a role at different levels of the organization hierarchy. It can be the same role or a different role. For example, you can assign a member role A for project 1 and project 2. Or you can assign a member role A for project 1 and role B for project 2.

## Access roles

BlueXP supports several predefined roles that you can assign to the members of your organization.

[Learn about access roles.](#)

## Connectors

When an *Organization admin* creates a Connector, BlueXP automatically associates that Connector with the organization and the currently selected project. The *Organization admin* automatically has access to that Connector from anywhere in the organization. But if you have other members in your organization with different roles, those members can only access that Connector from the project in which it was created, unless you associate that Connector with other projects.

You make a Connector available for another project in these cases:

- You want to allow members in your organization to use an existing Connector to create or discover additional working environments in another project
- You associated an existing resource with another project and that resource is managed by a Connector

If a resource that you associate with an additional project is discovered using a BlueXP Connector, then you also need to associate the Connector with the project that the resource is now associated with. Otherwise, the Connector and its associated resource aren't accessible from the BlueXP canvas by members who don't have the *Organization admin* role.

You can create an association from the **Connectors** page in BlueXP IAM:

- Associate a Connector with a project

When you associate a Connector with a project, that Connector is accessible from the BlueXP canvas when viewing the project.

- Associate a Connector with a folder

Associating a Connector with a folder doesn't automatically make that Connector accessible from all projects in the folder. Organization members can't access a Connector from a project until you associate the Connector with that specific project.

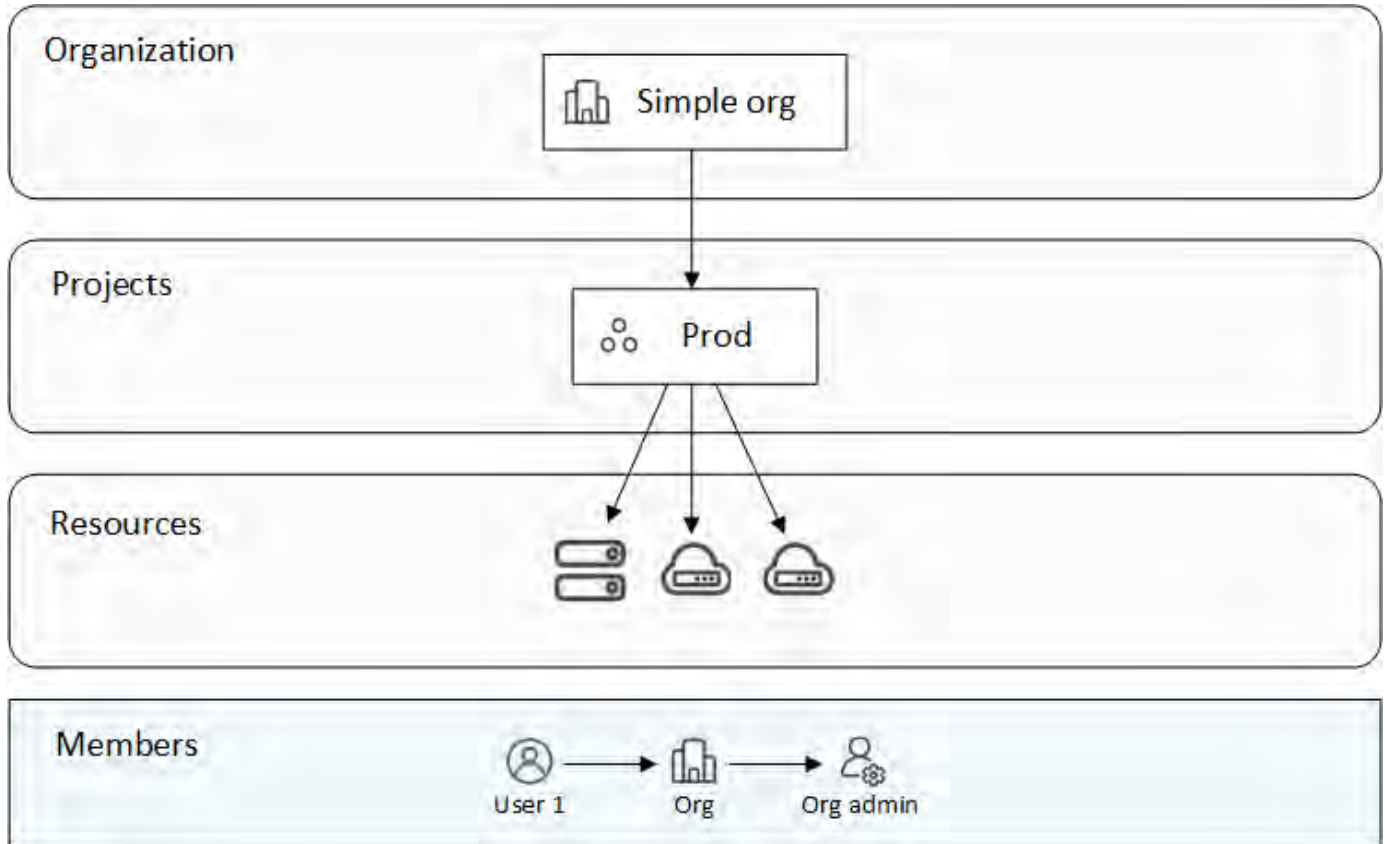
An *Organization admin* might associate a Connector with a folder so that the *Folder or project admin* can make the decision to associate that Connector with the appropriate projects that reside in the folder.

## IAM examples

These examples demonstrate how you might set up your organization.

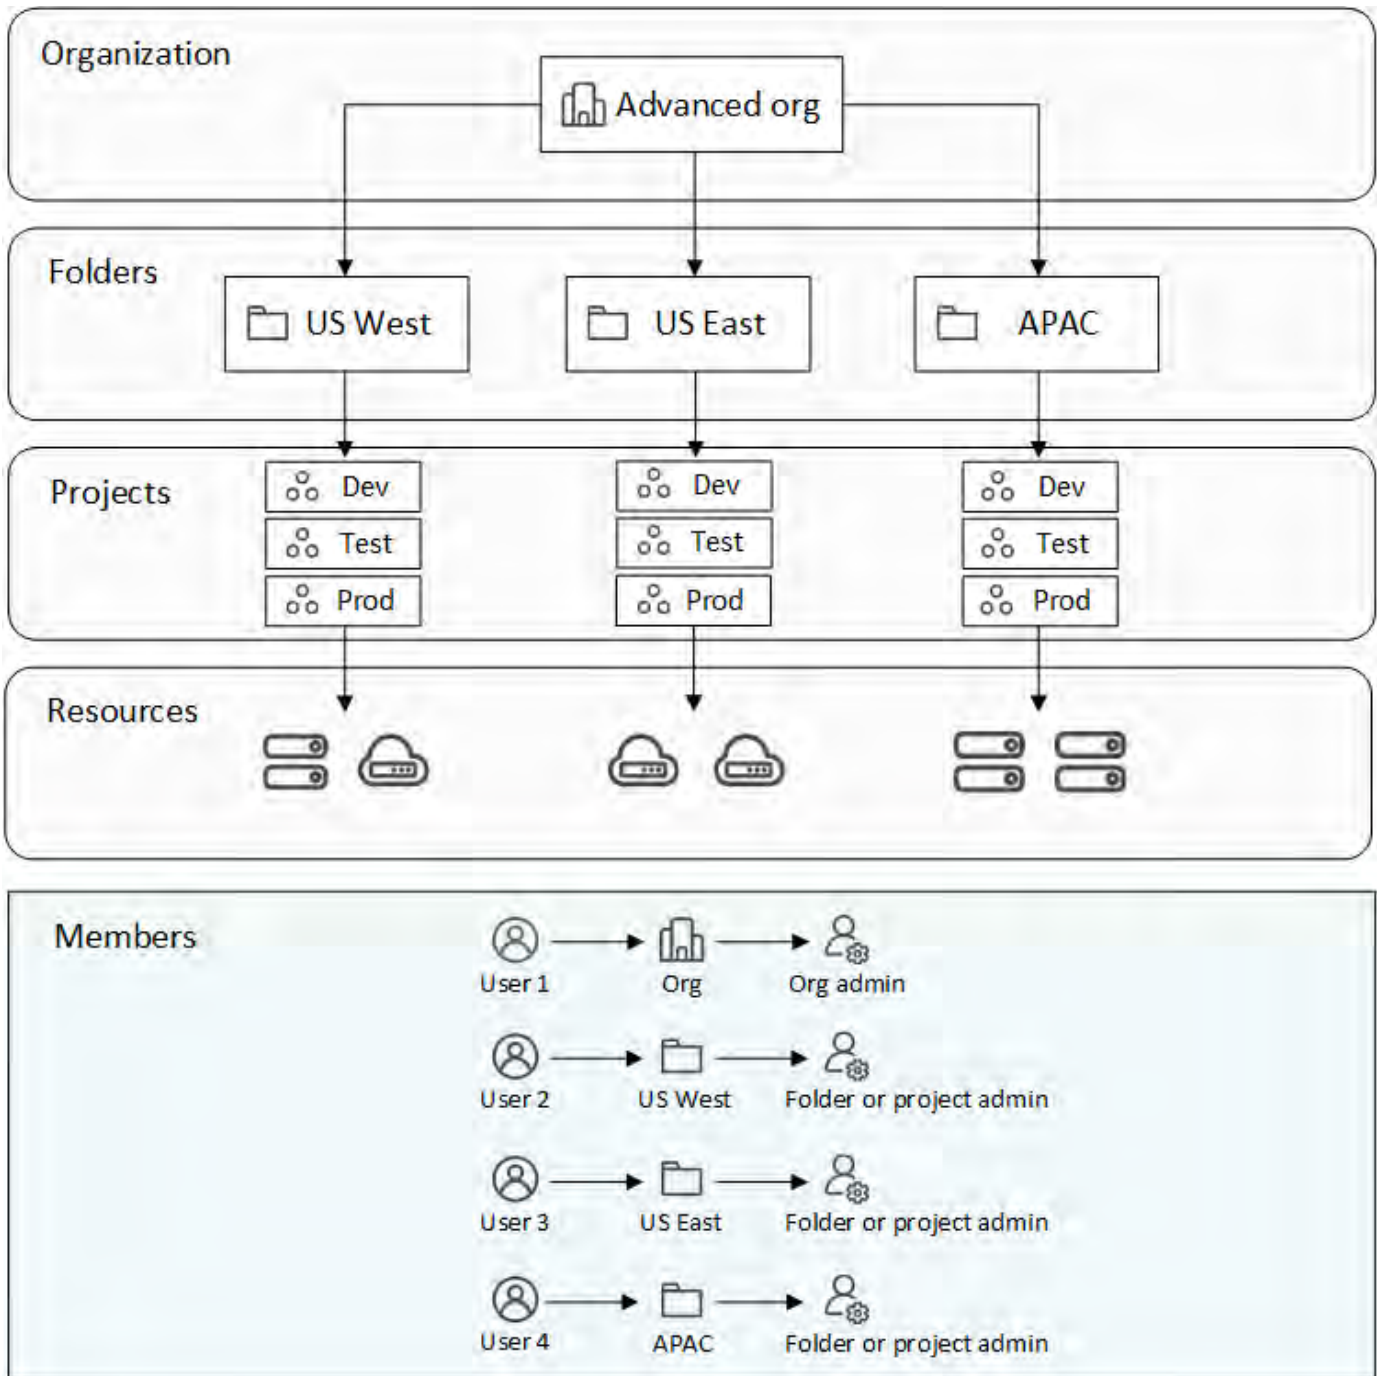
### Simple organization

The following diagram shows a simple example of an organization that uses the default project and no folders. A single member manages the entire organization.



### Advanced organization

The following diagram shows an organization that uses folders to organize the projects for each geographic location in the business. Each project has its own set of associated resources. The members include an organization admin and an admin for each folder in the organization.



### What you can do with BlueXP IAM

The following examples describe how you might use IAM to manage your BlueXP organization:

- Grant specific roles to specific members so that they can only complete the required tasks.
- Modify member permissions because they moved departments or because they have additional responsibilities.
- Remove a user who left the company.
- Add folders or projects to your hierarchy because a new business unit has added NetApp storage.
- Associate a resource with another project because that resource has capacity that another team can utilize.

- View the resources that a member can access.
- View the members and resources associated with a specific project.

## Where to go next

- [Get started with BlueXP IAM](#)
- [Organize your resources in BlueXP with folders and projects](#)
- [Manage BlueXP members and their permissions](#)
- [Manage the resource hierarchy in your BlueXP organization](#)
- [Associate Connectors with folders and projects](#)
- [Switch between BlueXP projects and organizations](#)
- [Rename your BlueXP organization](#)
- [Monitor or audit IAM activity](#)
- [BlueXP access roles](#)
- [Learn about the API for BlueXP IAM](#)

## Get started with BlueXP identity and access management

When you sign up to BlueXP, you're prompted to create a new organization. The organization includes one member (an Organization admin) and one default project. To set up BlueXP identity and access management (IAM) to meet your business needs, you'll need to customize your organization's hierarchy, add additional members, add or discover resources, and associate those resources across your hierarchy.

You must have **Organization admin** permissions to administer the entire organization from BlueXP IAM. If you have **Folder or project admin** permissions, you can only administer the folders and projects for which you have permissions.

Follow these steps to set up a new BlueXP organization. The order in which you complete these steps might be different, depending on your organization's needs.

1

### Edit the default project or add to your organization's hierarchy

Use the default project or create additional projects and folders matching your business hierarchy.

[Learn how to organize your resources with folders and projects.](#)

2

### Associate members with your organization

If multiple people in your business need access to BlueXP, associate their user accounts with your organization and assign the necessary permissions. You also have the option to add service accounts to your organization.

[Learn how to manage members and their permissions.](#)

3

### Add or discover resources

Add or discover resources in BlueXP as *the working environments*. Organization members manage a working environment, which represents a storage system, from within a project.

Learn how to create or discover resources:

- [Amazon FSx for NetApp ONTAP](#)
- [Azure NetApp Files](#)
- [Cloud Volumes ONTAP](#)
- [E-Series systems](#)
- [On-premises ONTAP clusters](#)
- [StorageGRID](#)



#### Associate resources with additional projects

When you create or discover a resource in BlueXP, that resource is automatically associated with the project that was selected when you created or discovered the working environment. If you want to make that resource available to another project in your organization, then you'll need to create an association between them. If a Connector manages the resource, associate the Connector with the respective project.S

- [Learn how to manage your organization's resource hierarchy.](#)
- [Learn how to associate a Connector with a folder or project.](#)

#### Related information

- [Learn about BlueXP identity and access management](#)
- [Learn about the API for BlueXP IAM](#)

### Organize your resources in BlueXP IAM with folders and projects

BlueXP identity and access management (IAM) enables you to organize your NetApp resources using projects and folders. A *project* represents a workspace in BlueXP that organization members access to manage *resources* (for example, a Cloud Volumes ONTAP system). A *folder* groups related projects together. After you organize your resources into folders and projects, you can grant granular access to resources by providing organization members with permissions to specific folders and projects.

#### Add a folder or project


When you create your BlueXP organization, it includes a single project. You can create additional projects to manage your organization's resources. You can optionally create folders to group related projects together.

#### About this task

Your organization's resource hierarchy can have up to 7 levels, with nested folders down to 6 levels and projects at the seventh level.

The following image illustrates the maximum depth of your organization's resource hierarchy:

#### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. From the **Organization** page, select **Add folder or project**.
3. Select **Folder** or **Project**.
4. Provide details about the folder or project:
  - **Name and location:** Enter a name and choose a location in the hierarchy for the folder or project. A folder or project can reside directly underneath the organization or within a folder.
  - **Resources:** Select the resources that you want to associate with this folder or project.

You can select resources associated with the parent folder or project: all resources for an organization parent, or folder-specific resources for a folder parent.

[Learn when you might associate a resource with a folder.](#)

- **Access:** View the members who will have access to the folder or project based on the existing permissions already defined in your resource hierarchy.

If needed, select **Add a member** to specify additional organization members who should have access to the folder or project and then select a role. A role defines the permissions that members have for the folder or project.

[Learn about predefined IAM roles.](#)

5. Select **Add**.

### Obtain the ID for a project

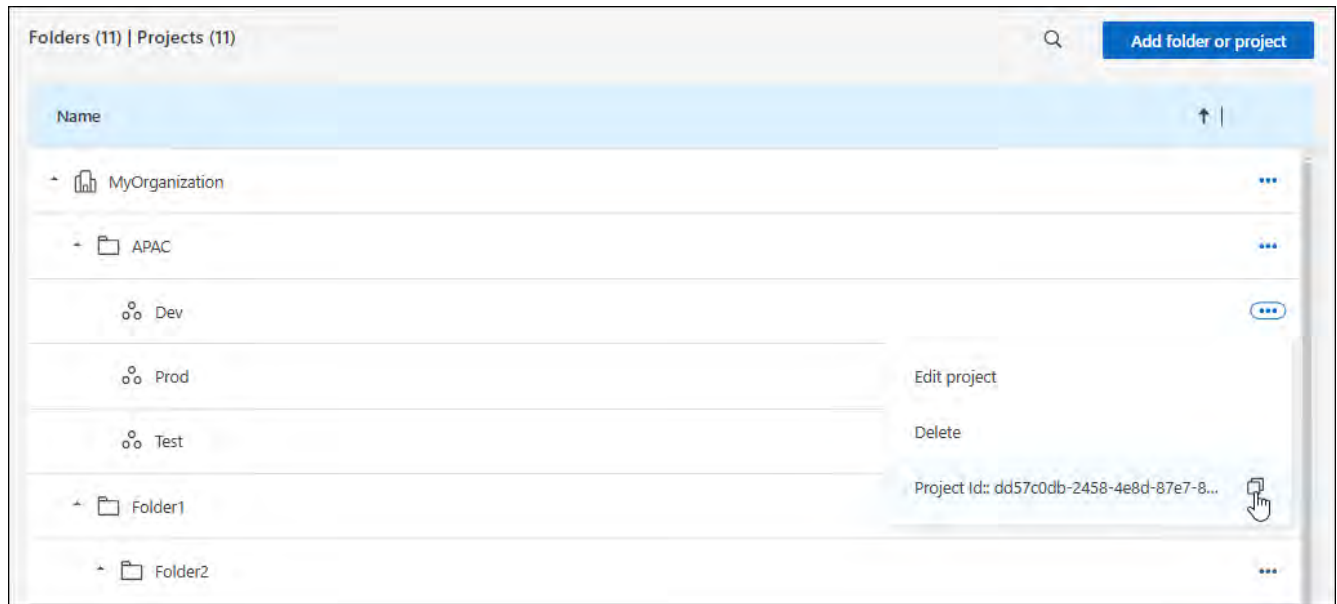
If you're using the BlueXP API, you might need to obtain the ID for a project. For example, when creating a Cloud Volumes ONTAP working environment.

#### Steps

1. From the **Organization** page, navigate to a project in the table and select **...**

The system displays the project ID.

2. To copy the ID, select the copy button.



## Rename a folder or project

If needed, you can change the name of your folders and projects.

### Steps

1. From the **Organization** page, navigate to a project or folder in the table, select **...** and then select **Edit folder** or **Edit project**.
2. On the **Edit** page, enter a new name and select **Apply**.

## Delete a folder or project

You can delete the folders and projects that you no longer need.

### Before you begin

- The folder or project must not have any associated resources. [Learn how to disassociate resources.](#)
- A folder must not contain any subfolders or projects. You need to delete those folders and projects first.

### Steps

1. From the **Organization** page, navigate to a project or folder in the table, select **...** and then select **Delete**.
2. Confirm that you want to delete the folder or project.

## View the resources associated with a folder or project

To verify that your resources are organized appropriately and accessible to the right members in your organization, you can view which resources and members are associated with a folder or project.

### Steps

1. From the **Organization** page, navigate to a project or folder in the table, select **...** and then select **Edit folder** or **Edit project**.





2. On the **Edit** page, you can view details about the selected folder or project by expanding the **Resources** or **Access** sections.

- Select **Resources** to view the associated resources. In the table, the **Status** column identifies the resources that are associated with the folder or project.

Available resources (45)

<input type="checkbox"/>	Platform Type	Resource Type	Resource Name	Status
<input type="checkbox"/>	AWS	Cloud Volumes ONTAP HA	Keystonecvo2	Associated
<input type="checkbox"/>	AWS	Cloud Volumes ONTAP HA	kfuKeystone1vadim	Associated
<input type="checkbox"/>	AWS	Cloud Volumes ONTAP	cvo1Vadim	Associated
<input type="checkbox"/>	AWS	Cloud Volumes ONTAP HA	cvoparts11test	Associated

## Modify the resources associated with a folder or project

Members with permissions for a folder or project can access its associated resources.

### Before you begin

[Learn when you might associate a resource with a folder.](#)

### Steps

1. From the **Organization** page, navigate to a project or folder in the table, select **...** and then select **Edit folder** or **Edit project**.
2. On the **Edit** page, select **Resources**.

In the table, the **Status** column identifies the resources that are associated with the folder or project.

3. Select the resources that you'd like to associate or disassociate.
4. Depending on the resources that you selected, select either **Associate with the project** or **Disassociate from the project**.

Available resources (45) | Selected (3)

Actions: Associate with the project | **Disassociate from the project**

<input type="checkbox"/>	Platform Type	Resource Type	Resource Name	Status
<input checked="" type="checkbox"/>	aws	Cloud Volumes ONTAP HA	Keystonecvo2	Associated
<input checked="" type="checkbox"/>	aws	Cloud Volumes ONTAP HA	kfuKeystone1vadim	Associated
<input checked="" type="checkbox"/>	aws	Cloud Volumes ONTAP	cvo1Vadim	Associated
<input type="checkbox"/>	aws	Cloud Volumes ONTAP HA	cvoparts11test	Associated
<input type="checkbox"/>	aws	Cloud Volumes ONTAP	cvosecondaryparts11	Associated
<input type="checkbox"/>	aws	Cloud Volumes ONTAP HA	keystonetest	Associated
<input type="checkbox"/>	aws	Cloud Volumes ONTAP HA	keystonetesting55	Associated

5. Select **Apply**

### View members associated with a folder or project

- Select **Access** to view the members who have access to the folder or project.

Access

Members (2) Learn more about user roles Add a member

Load users which inherits access

<input type="checkbox"/>	Type	Name	Role
<input type="checkbox"/>		Gabriel	Folder or project admin
<input type="checkbox"/>		Ben	Organization admin

### Modify member access to a folder or project

Modify member access to ensure the right members can access the associated resources.

Member access provided at a higher hierarchy level cannot be changed at lower levels. You need to switch to that part of the hierarchy and update the member's permissions there. Alternatively, you can [manage permissions from the Members page](#).

[Learn more about role inheritance.](#)

## Steps

1. From the **Organization** page, navigate to a project or folder in the table, select **...** and then select **Edit folder** or **Edit project**.
2. On the **Edit** page, select **Access** to view the list of members who have access to the selected folder or project.
3. Modify member access:
  - **Add a member**: Select the member that you'd like to add to the folder or project and assign them a role.
  - **Change a member's role**: For any members with a role other than Organization Admin, select their existing role and then choose a new role.
  - **Remove member access**: For members who have a role defined at the folder or project for which you're viewing, you can remove their access.
4. Select **Apply**.

## Related information

- [Learn about BlueXP identity and access management](#)
- [Get started with BlueXP IAM](#)
- [Learn about the API for BlueXP IAM](#)

## Add BlueXP members and service accounts

BlueXP identity and access management (IAM) enables you to add members to your organization and assign them one or more roles across your resource hierarchy. A *role* contains a set of permissions that enables a member to perform specific actions at a specific level of the resource hierarchy. You can associate new user accounts and service accounts, manage member roles, and more.



Ensure two members have the Organization admin role to avoid losing access to your BlueXP organization.

To manage users and their permissions, you must be assigned one of the following roles:

- Organization admin

Users with this role can manage all members

- Folder or project admin

Users with this role can manage members only of a designated folder or project

```
_Folder or project admin_ can view all members on the *Members* page but manage permissions only for folders and projects they have access to. xref:{relative_path}reference-iam-predefined-roles.html[Learn more about the actions that a _Folder or project admin_ can complete].
```

## Add members to your organization

You can add two types of members to your organization: a user account and a service account. A service account is used by applications to perform tasks via the BlueXP API without human intervention. A user account is typically used by a person to log in to BlueXP and manage resources.

Users must sign up for BlueXP before being added to an organization or assigned a role. However, you can create service accounts directly from BlueXP.

To manage users and their permissions, you must have the **Organization admin** role or the **Folder or project admin** role. Remember that users with the **Folder or project admin** role can only manage members for the folder or projects of which they have admin permissions.


## User account

### Steps

1. Direct the user to visit [NetApp BlueXP website](#) to sign up.

Once users sign up, they complete the **Sign up** page, check their email, and log in. If BlueXP prompts users to create an organization, they close it and notify you of their account creation. You can then add the user to your existing BlueXP organization.

[Learn how to sign up to BlueXP.](#)

2. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
3. Select **Members**.
4. Select **Add a member**.
5. To add the member, complete the steps in the dialog box:
  - **Entity Type**: Keep **User** selected.
  - **User's email**: Enter the user's email address that is associated with the BlueXP login that they created.
  - **Select an organization, folder, or project**: Choose the level of your resource hierarchy that the member should have permissions for.

Note the following:

- You can only select from the folders and projects for which you have admin permissions.
- Selecting an organization or folder grants the member permissions to all its contents.
- **Select a category** and then select a **Role** that provides the member with permissions for the resources that are associated with the organization, folder, or project that you selected.
  - If you selected a folder or project, you can choose from any role other than **Organization admin**.


[Learn about access roles.](#)

- **Add role**: If you want to provide access to additional folders or projects within your organization or grant the user further permissions in the selected area, select **Add role**, specify another folder or project or a different role category and then choose a role.
6. Select **Add**.

NetApp BlueXP sends the user an email with information on how to access BlueXP.

## Service account

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select **Members**.
3. Select **Add a member**.
4. To add the member, complete the steps in the dialog box:
  - **Entity Type**: Select **Service account**.

- **Service account name:** Enter a name for the service account.
- **Select an organization, folder, or project:** Choose the level of your resource hierarchy that the member should have permissions for.

Note the following:

- You can only select from the folders and projects for which you have admin permissions.
- Selecting an organization or folder grants the member permissions to all its contents.
- **Select a category** then select a **Role** that provides the member with permissions for the resources that are associated with the organization, folder, or project that you selected.
  - If you selected a folder or project, you can choose from any role other than **Organization admin**.

[Learn about predefined IAM roles.](#)

- **Add role:** If you want to provide access to additional folders or projects within your organization or grant the user further permissions in the selected area, select **Add role**, specify another folder or project or a different role category and then choose a role.

5. Download or copy the client ID and client secret.

BlueXP displays the client secret only once. Copy or download it and store it securely. Note that you can recreate the client ID and client secret later on as needed.

6. Select **Close**.

## View organization members


You can view a list of all members in your BlueXP organization. To understand which resources and permissions are available to a member, you can view the roles assigned to the member at different levels of your organization's resource hierarchy. [Learn how to use roles to control access to BlueXP resources.](#)

You can view both user accounts and service accounts from the **Members** page.



You can also view all of the members associated with a specific folder or project. [Learn more.](#)

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select **Members**.

The **Members** table lists the members of your organization.

3. From the **Members** page, navigate to a member in the table, select  and then select **View details**.

## Remove a member from your organization

You might need to remove a member from your organization—for example, if they leave your company.

Removing a member removes their permissions but keeps their BlueXP and NetApp Support Site accounts.

### Steps

1. From the **Members** page, navigate to a member in the table, select **⋮** then select **Delete user**.
2. Confirm that you want to remove the member from your organization.


## Recreate the credentials for a service account

Create new credentials if lost or when updating security credentials becomes necessary.

### About this task

When you recreate the credentials, you delete the existing credentials for the service account and create new ones. You cannot use the previous credentials.

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select **Members**.
3. In the **Members** table, navigate to a service account, select **⋮** and then select **Recreate secrets**.
4. Select **Recreate**.
5. Download or copy the client ID and client secret.

BlueXP displays the client secret only once. Copy or download it and store it securely.

## Manage a user's multi-factor authentication (MFA)

If a user has loses access to their MFA device, you can either remove or disable their MFA configuration.

If you remove their MFA configuration, the user needs to set up MFA again when they log in to BlueXP. If the user has only lost access to their MFA device temporarily, they can use the recovery code that they saved when they set up MFA to log in to BlueXP.

If they do not have their recovery code, temporarily disable MFA to allow login. When you disable MFA for a user, it is disabled for only eight hours and then re-enabled automatically. The user is allowed one login during that time without MFA. After the eight hours, the user must use MFA to log in to BlueXP.



You must have an email address in the same domain as the affected user to manage that user's multi-factor authentication.

### Steps

1. In the upper right of the console, select  > **Identity & Access Management**.
2. Select **Members**.

The members of your organization appear in the **Members** table.

3. From the **Members** page, navigate to a member in the table, select **⋮** and then select **Manage multi-factor authentication**.
4. Choose whether to remove or to disable the user's MFA configuration.

### Related information

- [Learn about BlueXP identity and access management](#)

- [Get started with BlueXP IAM](#)
- [Predefined BlueXP IAM roles](#)
- [Learn about the API for BlueXP IAM](#)

## Use roles to manage user access to resources

Within BlueXP, you can assign roles to users based on what they need to do and where.

Users with the **Organization admin** or **Folder or project admin** role have the responsibility of assigning roles to other users. You can assign access roles on a project or folder basis. For example, you can assign a user the Ransomware protection admin role for one project and the SnapCenter admin role for a different project. Alternatively, if a user needs the Classification admin role for all projects within a specific folder, you can give them this role at the folder level.

Use access roles to assign access to storage resources based on the specific tasks that users need to perform. For example, if a user needs to interact with ransomware protection services, they must be given an access role that includes either viewing or administrative permissions for the ransomware protection service for the project for which the access role is granted.

Assign roles to users based on your IAM strategy for enhanced security. IAM roles ensure users have only the access they need.



Remember that you can't directly grant access to resources. Assign resources to projects first. Consider setting up your resource hierarchy before assigning users access. [Learn how to organize your resources in BlueXP IAM with folders and projects.](#)

### View roles(s) assigned to a member

When you add a member to your organization, you are prompted to assign them a role. You can members to verify which roles they are currently assigned.

If you have the *Folder or project admin* role, the page displays all members in the organization. However, you can only view and manage member permissions for the folders and projects for which you have permissions. [Learn more about the actions that a \*Folder or project admin\* can complete.](#)

1. From the **Members** page, navigate to a member in the table, select **...** and then select **View details**.
2. In the table, expand the respective row for organization, folder, or project where you want to view the member's assigned role and select **View** in the **Role** column.

### Add an access role to a member

You typically assign a role when adding a member to your organization, but you can update it at any time by removing or adding roles.

You can assign a user an access role for your organization, folder, or project.

Members can have multiple roles within the same project and in different projects. For example, smaller organizations may assign all available access roles to the same user, while larger organizations may have users do more specialized tasks. Alternatively, you could also assign one user the Ransomware protection admin role for an organization. In that example, the user would be able to perform ransomware protection tasks on all projects within your organization.



Your access role strategy should align with the way you have organized your NetApp resources.



A member who is assigned the Organization admin role can't be assigned any additional roles. They already have permissions across the entire organization. A member with the Folder or project role can't be assigned any other roles within the folder or project where they have that role already. Both of these roles provide access to all services within the scope that they are assigned.

### Steps

1. From the **Members** page, navigate to a member in the table, select **...** and then select **Add a role**.
2. To add a role, complete the steps in the dialog box:

- **Select an organization, folder, or project:** Choose the level of your resource hierarchy that the member should have permissions for.

If you select the organization or a folder, the member will have permissions to everything that resides within the organization or folder.

- **Select a category:** Choose a role category. [Learn about access roles](#).
- **Select a Role:** Choose a role that provides the member with permissions for the resources that are associated with the organization, folder, or project that you selected.

[Learn about access roles](#).

\* **Add role:** If you want to provide access to additional folders or projects within your organization, select **Add role**, specify another folder or project or role category, and then select a role category and a corresponding role.

1. Select **Add new roles**.


### Change a member's assigned role

You can change the assigned roles for a member should you need to adjust the access for a user.



Users must have at least one role assigned to them. You can't remove all roles from a user. If you need to remove all roles, you must delete the user from your organization.

### Steps

1. From the **Members** page, navigate to a member in the table, select **...** and then select **View details**.
2. In the table, expand the respective row for organization, folder, or project where you want to change the member's assigned role and select **View** in the **Role** column to view the roles assigned to this member.
3. You can change an existing role for a member or remove a role.
  - a. To change a member's role, select **Change** next to the role you want to change. You can only change this role to a role within the same role category. For example, you can change from one data service role to another. Confirm the change.
  - b. To unassign a member's role, select  next to the role to unassign the member the respective role. You'll be asked to confirm the removal.

## Manage the resource hierarchy in your BlueXP organization

When you use associate a member with your organization, you provide permissions at

the organization, folder, or project level. To ensure that those members have permissions to access the right resources, you'll need to manage the resource hierarchy of your organization by associating resources with specific projects and folders. A *resource* is a storage resource that BlueXP already manages or is aware of.


### View the resources in your organization

You can view both discovered and undiscovered resources associated with your organization. Undiscovered resources are storage resources identified by digital advisor but not added as working environments.



The IAM resources page excludes Amazon FSx for NetApp ONTAP resources because you cannot associate them with an IAM role. View these resources on their respective canvas or from workloads.

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select **Resources** to view the Resources page.
3. Select **Advanced Search & Filtering**.
4. Use any of the available options to find the resource that you're looking for:
  - **Search by resource name:** Enter a text string and select **Add**.
  - **Platform:** Select one or more platforms, such as Amazon Web Services.
  - **Resources:** Select one or more resources, such as Cloud Volumes ONTAP.
  - **Organization, folder, or project:** Select the entire organization, a specific folder, or a specific project.
5. Select **Search**.


### Associate a resource with folders and projects

Associate a resource to a folder or project to make it available.

### Before you begin

You should understand how resource association works. [Learn about resources, including when to associate a resource with a folder.](#)

### Steps

1. From the **Resources** page, navigate to a resource in the table, select  and then select **Associate to folders or projects**.
2. Select a folder or project and then select **Accept**.
3. To associate an additional folder or project, select **Add folder or project** and then select the folder or project.

Note that you can only select from the folders and projects for which you have admin permissions.

4. Select **Associate resources**.
  - If you associated the resource with projects, members who have permissions for those projects now have the ability to access the resource in BlueXP.
  - If you associated the resource with a folder, a *Folder or project admin* can now access the resource

from within BlueXP IAM. [Learn about associating a resource with a folder.](#)

## After you finish

If you discover a resource using a BlueXP Connector associate the Connector with the project to grant them access. Otherwise, the Connector and its associated resource are not accessible from the BlueXP canvas by members without the *Organization admin* role.

[Learn how to associate a Connector with a folder or project.](#)

## View the folders and projects associated with a resource

To identify where a resource is available in your organization's hierarchy, you can view the folders and projects that are associated with that resource.

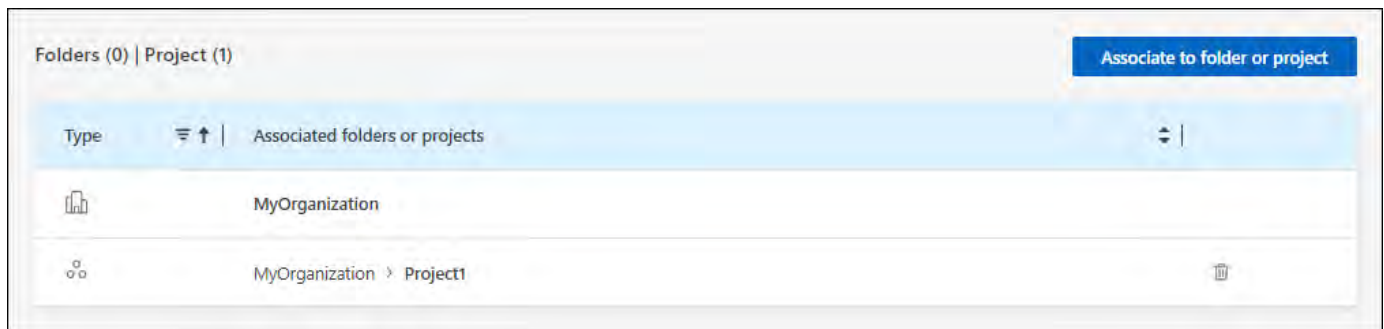


If you need to determine which organization members have access to the resource, you can [view the members who have access to the folders and projects that are associated with the resource.](#)

## Steps

1. From the **Resources** page, navigate to a resource in the table, select **...** and then select **View details**.

The following example shows a resource that is associated with one project.



If you need to determine which organization members have access to the resource, you can [view the members who have access to the folders and projects that are associated with the resource.](#)

## Remove a resource from a folder or project

To remove a resource from a folder or project, you need to remove the association between the folder or project and the resource. Removing the association prevents members from managing the resource in the folder or project.



If you want to remove a discovered resource from the entire organization, you need to remove the working environment from the BlueXP canvas.

## Steps

1. From the **Resources** page, navigate to a resource in the table, select **...** and then select **View details**.
2. For the folder or project for which you want to remove the resource, select
3. Confirm that you want to remove the association by selecting **Delete**.

## Related information

- [Learn about BlueXP identity and access management](#)
- [Get started with BlueXP IAM](#)
- [Learn about the API for BlueXP IAM](#)

## Associate a BlueXP Connector with other folders and projects

When an `_Organization admin_` creates a Connector, it is automatically associated with currently selected project within the organization. Although someone with the `_Organization admin_` can access to that Connector from anywhere in the organization. Other members in your organization can only access that Connector from the project in which it was created, unless you associate that Connector with other projects.


### Before you begin

You should understand how Connector association works. [Learn about using Connectors with BlueXP IAM.](#)

### About this task

- When a *Folder or project admin* views the **Connectors** page, the page displays all Connectors in the organization. However, a member with this role can only view and associate Connectors with the folders and projects for which they have permissions. [Learn more about the actions that a \*Folder or project admin\* can complete.](#)

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select **Connectors**.
3. From the table, find the Connector that you want to associate.

Use the search above the table to find a specific Connector or filter the table by resource hierarchy.

4. To view the folders and projects linked to the Connector, select  and then select **View details**.

BlueXP displays details about the folders and projects that the Connector is associated with.

5. Select **Associate to folder or project**.
6. Select a folder or project and then select **Accept**.
7. To associate the Connector with an additional folder or project, select **Add a folder or project** and then select the folder or project.
8. Select **Associate Connector**.

### After you finish

If you want to associate the resources that the Connector manages with the same folders and projects, you can do so from the Resources page.

[Learn how to associate a resource with folders and projects.](#)

## Related information

- [Learn about BlueXP Connectors](#)
- [Learn about BlueXP identity and access management](#)
- [Get started with BlueXP IAM](#)
- [Learn about the API for BlueXP IAM](#)

## Switch between BlueXP organizations, projects, and Connectors

You might belong to multiple BlueXP organizations or have permissions to access multiple projects or Connectors within a BlueXP organization. When needed, you can easily switch between organizations, projects, and Connectors to access the resources associated with that organization, project, or Connector.




You might belong to multiple organizations if you were invited to join another organization or if you created an additional organization yourself. You can create an additional organization by using the API. [Learn how to create a new organization](#)

### Switch between organizations

If you are a member of multiple organizations, you can switch between them at any time.

#### Steps

1. At the top of BlueXP, select **Organization**.



2. Select another organization and then select **Switch**.

#### Result

BlueXP switches to the selected organization and displays the resources associated with that organization.

### Switch between projects


If your organization includes multiple projects and you have access to those projects, you can switch between them at any time.

#### Before you begin

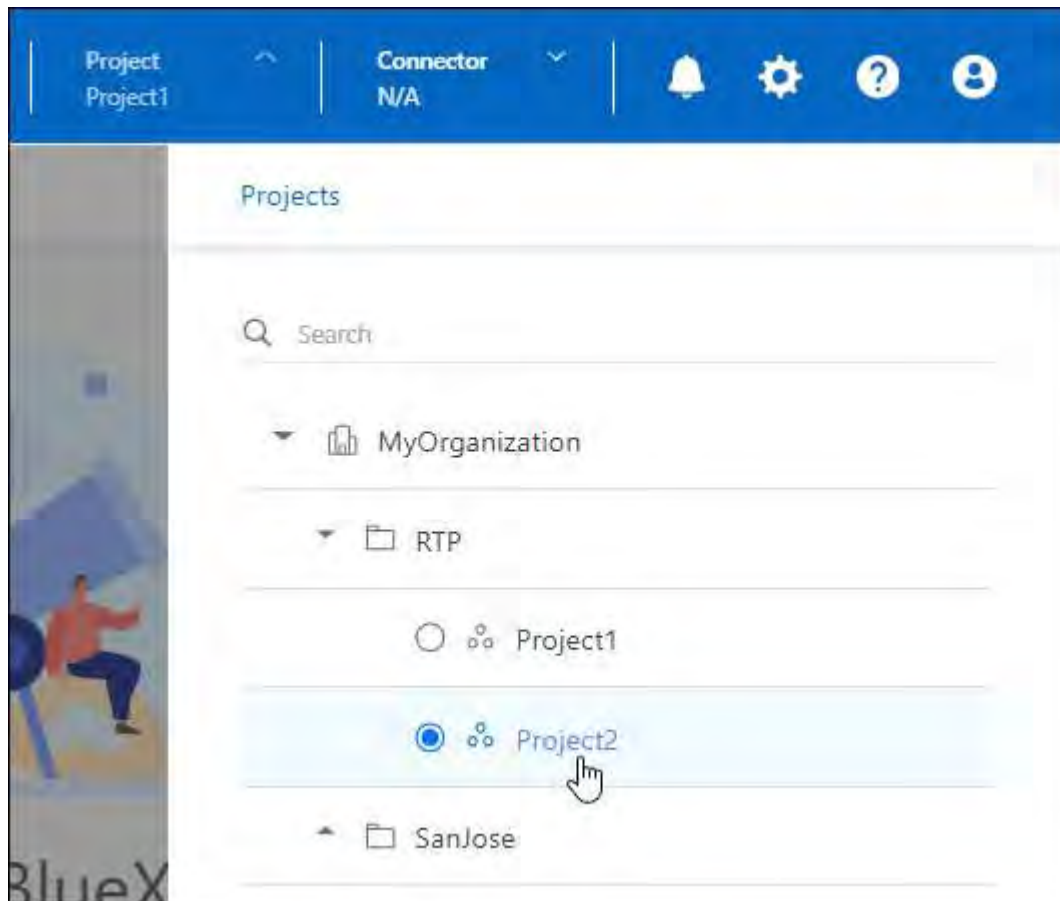
You must be on any page in the BlueXP console other than the BlueXP identity and access management (IAM) pages. You can't switch to another project when viewing any of the IAM pages.

#### Steps

1. At the top of BlueXP, select **Project**.



2. Browse through the folders and projects in your organization, select the project that you want, and then select **Switch**.



### Result

BlueXP switches to the selected project and displays the resources associated with that project.

### Switch between Connectors

If you have multiple Connectors, you can switch between them to see the working environments that are associated with a specific Connector.

### Steps

1. At the top of BlueXP, select **Connector**.
2. Select another Connector and then select **Switch**.

### Result

BlueXP refreshes and shows the working environments associated with the selected Connector.

## Related information

[Associate Connectors with folders and projects.](#)

## Related information

- [Learn about BlueXP identity and access management](#)
- [Get started with BlueXP IAM](#)
- [Learn about the API for BlueXP IAM](#)



## Organization and project IDs

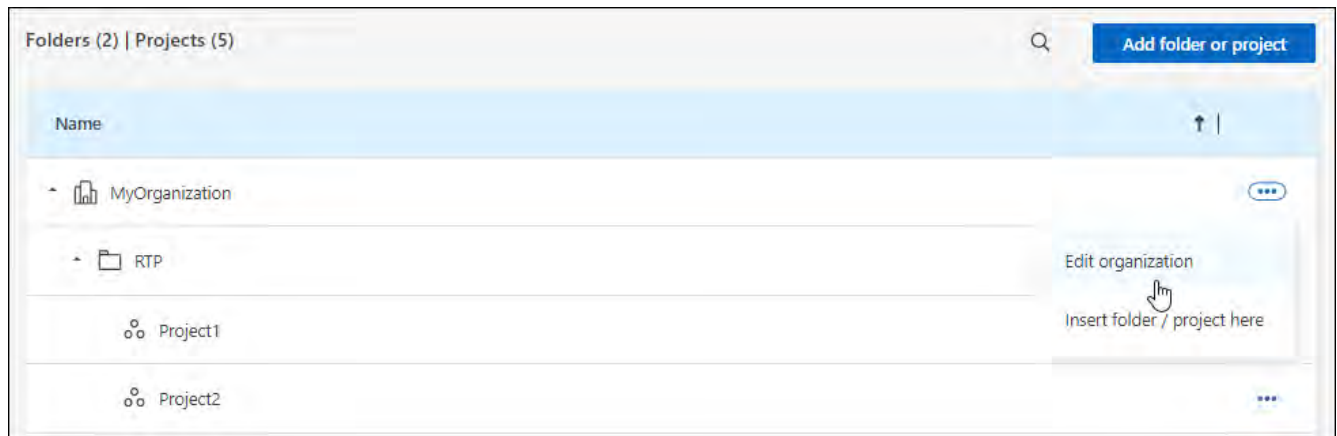
Your BlueXP organization has a name and an ID. You can choose a name for your organization to help identify it in your BlueXP deployment. You may also need to retrieve the organization ID for certain integrations.

### Rename your organization

You can rename your organization within BlueXP. This is helpful if you support more than organization within your BlueXP deployment.

#### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. From the **Organization** page, navigate to the first row in the table, select  and then select **Edit organization**.




3. Enter a new organization name and select **Apply**.

### Get the organization ID

The organization ID is used for certain integrations with BlueXP.

You can view the organization ID from the Organizations page and copy it to the clipboard for your needs.

#### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Organization** tab to view the **Organization** page.

3. On the **Organization** page, look for your organization ID in the summary bar and copy it to the clipboard. You can save this for use later or copy it directly to where you need to use it.

## Obtain the ID for a project

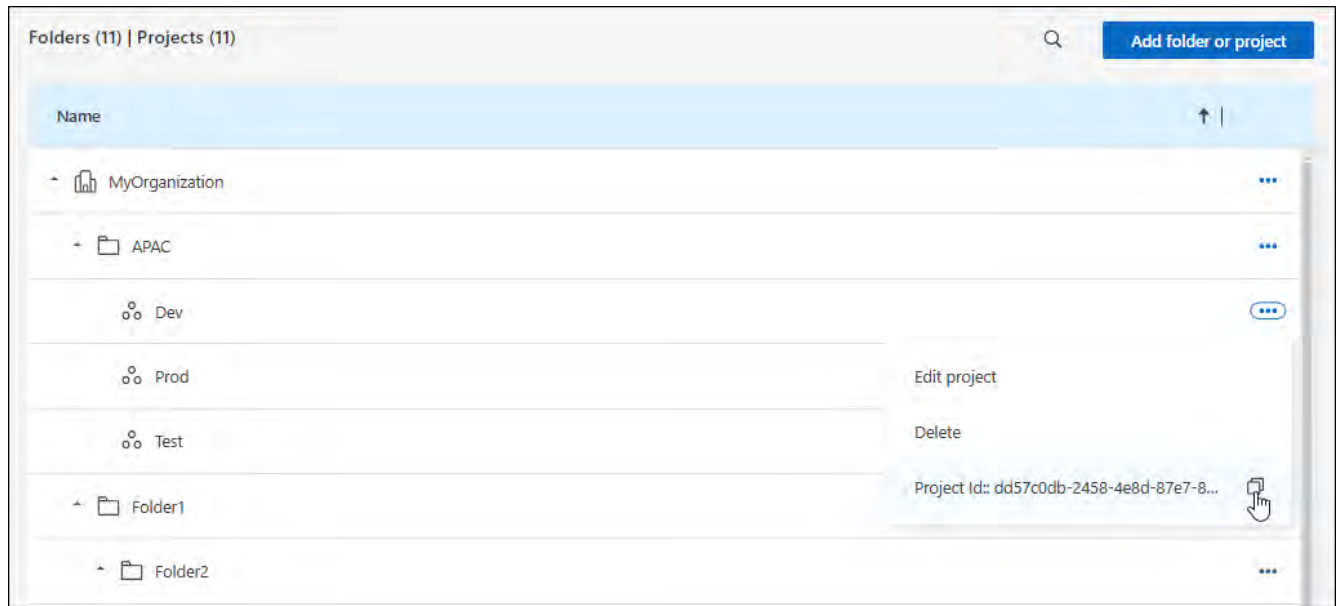
If you're using the BlueXP API, you might need to obtain the ID for a project. For example, when creating a Cloud Volumes ONTAP working environment.

### Steps

1. From the **Organization** page, navigate to a project in the table and select **...**

The project ID displays.

2. To copy the ID, select the copy button.




### Related information

- [Learn about BlueXP identity and access management](#)
- [Get started with BlueXP IAM](#)
- [Learn about the API for BlueXP IAM](#)

## Monitor or audit IAM activity from the BlueXP timeline

If you need to monitor or audit an action that was completed from BlueXP identity and access management (IAM), you can view details from the BlueXP Timeline. For example, you might want to verify who added a member to an organization or that a project was deleted successfully.

### Steps

1. In the upper right of the BlueXP console, select  > **Timeline**.
2. From the filters, select **Service** and then select **Tenancy**.
3. Use any of the other filters to change which actions display in the table.



For example, you can use the **User** filter to show actions related to a specific user account.

## Result

The Timeline updates to show you completed management actions related to BlueXP IAM.

## BlueXP access roles

### Learn about BlueXP access roles

BlueXP identity and access management (IAM) includes predefined roles that you can assign to the members of your organization across different levels of your resource hierarchy. Before you assign these roles, you should understand the permissions that each role includes. Roles fall into the following categories: platform, application, and data service.

### Platform roles

Platform roles grant all BlueXP administration permissions, including assigning roles and adding users. Platform roles provide access to all BlueXP data services and applications. BlueXP IAM includes two platform roles: Organization admin and Folder or project admin. The main difference between the two BlueXP IAM platform roles is scope.

Platform role	Responsibilities
<a href="#">Organization admin</a>	<p>Allows a user unrestricted access to all projects and folders within an organization, add members to any project or folder, as well as perform any BlueXP task and use any data service that does not have an explicit role associated with it.</p> <p>Users with this role organize and manage your BlueXP organization. They create folders and projects, assign roles, add users, and can manage all working environments, if they have the credentials to do so.</p> <p>This is the only access role that can create Connectors.</p>
<a href="#">Folder or project admin</a>	<p>Allows a user unrestricted access to specific projects and folders to which they are assigned. Can add members to folders or projects they manage, as well as perform any BlueXP task and use any data service or application on resources within the folder or project they are assigned.</p> <p>Folder or project admins cannot create Connectors.</p>
<a href="#">Federation admin</a>	<p>Allows a user create and manage federations with BlueXP, which enables single-sign on (SSO).</p>
<a href="#">Federation viewer</a>	<p>Allows a user to view existing federations with BlueXP. Cannot create or manage federations</p>

### Application roles

The following is a list of roles in the application category. Each role grants specific permissions within its designated scope. Users who do not have the required application role or a platform role will be unable to access the application.

Application role	Responsibilities
<a href="#">Google Cloud NetApp Volumes admin</a>	Users with the Google Cloud NetApp Volumes role can discover and manage Google Cloud NetApp Volumes.
<a href="#">Keystone admin</a>	Users with the Keystone admin role can create service requests. Allows users to monitor and view consumption, assets, and administrative information within the Keystone tenant they are accessing.
<a href="#">Keystone viewer</a>	Users with the Keystone viewer role CANNOT create service requests. Allows users to monitor and view consumption, assets, and administrative information within the Keystone tenant they are accessing.
ONTAP Mediator setup role	Service accounts with the ONTAP Mediator setup role can create service requests. This role is required in a service account to configure an instance of the <a href="#">ONTAP Cloud Mediator</a> .
<a href="#">Operation support analyst</a>	Provides access to alerts and monitoring tools and ability to enter and manage support cases.
<a href="#">Storage admin</a>	Administer storage health and governance functions, discover storage resources, as well as modify and delete existing working environments.
<a href="#">Storage viewer</a>	View storage health and governance functions, as well as view previously discovered storage resources. Cannot discover, modify, or delete existing storage working environments.
<a href="#">System health specialist</a>	Administer storage and health and governance functions, all permissions of the Storage admin except cannot modify or delete existing working environments.

#### Data service roles

The following is a list of roles in the data service category. Each role grants specific permissions within its designated scope. Users who do not have the required data service role or a platform role will be unable to access the data service.

Data service role	Responsibilities
<a href="#">Backup and recovery super admin</a>	Perform any actions in the Backup and recovery service.
<a href="#">Backup and recovery admin</a>	Perform backups to local snapshots, replicate to secondary storage, and back up to object storage.
<a href="#">Backup and recovery restore admin</a>	Restore workloads in the Backup and recovery service.
<a href="#">Backup and recovery clone admin</a>	Clone applications and data in the Backup and recovery service.
<a href="#">Backup and recovery viewer</a>	View Backup and recovery information.
<a href="#">Disaster recovery admin</a>	Perform any actions in the Disaster recovery service.
<a href="#">Disaster recovery failover admin</a>	Perform failover and migrations.

Data service role	Responsibilities
<a href="#">Disaster recovery application admin</a>	Create replication plans, modify replication plans, and start test failovers.
<a href="#">Disaster recovery viewer</a>	View information only.
Classification viewer	Provides the ability to view BlueXP classification scan results.  Users with this role can view compliance information and generate reports for resources that they have permission to access. These users can't enable or disable scanning of volumes, buckets, or database schemas. Classification does not have a viewer role.
<a href="#">Ransomware protection admin</a>	Manage actions on the Protect, Alerts, Recover, Settings, and Reports tabs of the Ransomware protection service.
<a href="#">Ransomware protection viewer</a>	View workload data, view alert data, download recovery data, and download reports in the Ransomware protection service.
SnapCenter admin	Provides the ability to back up snapshots from on-premises ONTAP clusters using BlueXP backup and recovery for applications. A member who has this role can complete the following actions in BlueXP:  * Complete any action from Backup and recovery > Applications * Manage all working environments in the projects and folders for which they have permissions * Use all BlueXP services  SnapCenter does not have a viewer role.

#### Related links

- [Learn about BlueXP identity and access management](#)
- [Get started with BlueXP IAM](#)
- [Manage BlueXP members and their permissions](#)
- [Learn about the API for BlueXP IAM](#)

#### BlueXP platform access roles

Assign platform roles to users to grant permissions to perform administration tasks in BlueXP, assign roles, add users, create Connectors, and manage federations.

#### Example for organization roles in BlueXP for a large multi-national organization

XYZ Corporation organizes data storage access by region—North America, Europe, and Asia-Pacific—providing regional control with centralized oversight.

The **Organization admin** in XYZ Corporation's BlueXP creates an initial organization and separate folders for each region. The **Folder or project admin** for each region organizes projects (with associated resources) within the region's folder.

Regional admins with the **Folder or project admin** role actively manage their folders by adding resources and users. These regional admins can also add, remove, or rename folders and projects they manage. The **Organization admin** inherits permissions for any new resources, maintaining visibility of storage usage across

the entire organization.

Within the same organization, one user is assigned the **Federation admin** role to manage the federation of the organization with their corporate IdP. This user can add or remove federated organizations, but cannot manage users or resources within the organization. The **Organization admin** assigns a user the **Federation viewer** role to check federation status and view federated organizations.

The following tables indicate the actions that each BlueXP platform role can perform.

#### Organization administration roles

Task	Organization admin	Folder or project admin
Create Connectors	Yes	No
Create, modify or delete working environments (add or discover new resources using the BlueXP canvas)	Yes	Yes
Create folders and projects, including deleting	Yes	No
Rename existing folders and projects	Yes	Yes
Assign roles and add users	Yes	Yes
Associate resources with folders and projects	Yes	Yes
Associate Connectors with folders and projects	Yes	No
Remove Connectors from a folders and projects	Yes	No
Manage Connectors (edit certificates, settings, and so on)	Yes	No
Manage credentials from Settings > Credentials	Yes	Yes
Create, manage, and view federations	Yes	No
Register for support and submit cases through BlueXP	Yes	Yes
Use data services	Yes	Yes
View the BlueXP timeline and notifications	Yes	Yes

#### Federation roles

Task	Federation admin	Federation viewer
Create a federation	Yes	No
Verify a domain	Yes	No
Add a domain to a federation	Yes	No
Disable and delete federations	Yes	No
Test federations	Yes	No
View federations and their details	Yes	Yes

## Application roles

### Keystone access roles for BlueXP

Keystone roles provide access to the Keystone dashboards and allow users to view and manage their Keystone subscription. There are two Keystone roles: Keystone admin and Keystone viewer. The main difference between the two roles is the actions they can take in Keystone. The Keystone admin role is the only role that is allowed to create service requests or modify subscriptions.

### Example for Keystone roles in BlueXP

XYZ Corporation has four storage engineers from different departments who view Keystone subscription information. Although all of these users need to monitor the Keystone subscription, only the team lead is allowed to make service requests. Three of the team members are given the **Keystone viewer** role, while the team lead is given the **Keystone admin** role so that there is a point of control over service requests for the company.

The following table indicates the actions that each Keystone role can perform.

Feature and action	Keystone admin	Keystone viewer
View the following tabs: Subscription, Assets, Monitor, and Administration	Yes	Yes
<b>Keystone subscription page:</b>		
View subscriptions	Yes	Yes
Amend or renew subscriptions	Yes	No
<b>Keystone assets page:</b>		
View assets	Yes	Yes
Manage assets	Yes	No
<b>Keystone alerts page:</b>		
View alerts	Yes	No
Manage alerts	Yes	No
Create alerts for self	Yes	Yes
<b>Digital wallet:</b>		
Can view digital wallet	Yes	Yes
<b>Keystone reports page:</b>		
Download reports	Yes	Yes

Feature and action	Keystone admin	Keystone viewer
Manage reports	Yes	Yes
Create reports for self	Yes	Yes
<b>Service requests:</b>		
Create service requests	Yes	No
View service requests created by any user within the Organization	Yes	Yes

#### Operational support analyst access role for BlueXP

You can assign the following role to users to provide them access to alerts and monitoring. Users with this role can also open support cases.

#### Operational support analyst

Task	Can perform
Manage own user credentials from Settings > Credentials	Yes
View discovered resources	Yes
Register for support and submit cases through BlueXP	Yes
View the BlueXP timeline and notifications	Yes
View, download, and configure alerts	Yes

#### Storage access roles for BlueXP

You can assign the following roles to users to provide them access to the storage management features in BlueXP that are associated with supported storage resources. You can assign users an administrative role to manage storage or a viewer role for monitoring.



These roles are not available from the BlueXP partnership API.

Administrators can assign storage roles to users for the following storage resources and features:

Storage resources:

- On-premises ONTAP clusters
- StorageGRID
- E-Series

BlueXP services and features:

- Digital advisor
- Software updates
- Economic efficiency
- Sustainability

### Example for storage roles in BlueXP

XYZ Corporation, a multinational company, has a large team of storage engineers and storage administrators. They allow this team to manage storage assets for their regions while limiting access to core BlueXP tasks like user management, Connector creation, and cost tools such as the digital wallet.

Within a team of 12, two users are given the **Storage viewer** role which allows them to monitor the storage resources associated with the BlueXP projects they are assigned to. The remaining nine are given the **Storage admin** role which includes the ability to manage software updates, access ONTAP System Manager through BlueXP, as well as discover storage resources (add working environments). One person on the team is given the **System health specialist** role so they can manage the health of the storage resources in their region, but not modify or delete any working environments. This person can also perform software updates on the storage resources for projects they are assigned.

The organization has two additional users with the **Organization admin** role who can manage all aspects of BlueXP, including user management, Connector creation, and cost management tools like digital wallet, as well as several users with the **Folder or project admin** role who can perform BlueXP administration tasks for the folders and projects they are assigned to.

The following table shows the actions each BlueXP storage role performs.

Feature and action	Storage admin	System health specialist	Storage viewer
<b>Canvas:</b>			
Discover new resources (create new working environment)	Yes	Yes	No
View discovered resources	Yes	Yes	No
Delete working environments	Yes	No	No
Modify working environments	Yes	No	No
<b>Create Connector</b>	No	No	No
<b>Digital advisor</b>			
View all pages and functions	Yes	Yes	Yes
<b>Digital wallet</b>			
View all pages and functions	No	No	No
<b>Software updates</b>			

<b>Feature and action</b>	<b>Storage admin</b>	<b>System health specialist</b>	<b>Storage viewer</b>
View landing page and recommendations	Yes	Yes	Yes
Review potential version recommendations and key benefits	Yes	Yes	Yes
View update details for a cluster	Yes	Yes	Yes
Run pre-update checks and download upgrade plan	Yes	Yes	Yes
Install software updates	Yes	Yes	No
<b>Economic efficiency</b>			
Review capacity planning status	Yes	Yes	Yes
Choose next action (best practice, tier)	Yes	No	No
Tier cold data to cloud storage and free up storage	Yes	Yes	No
Set up reminders	Yes	Yes	Yes
<b>Sustainability</b>			
View dashboard and recommendations	Yes	Yes	Yes
Download report data	Yes	Yes	Yes
Edit carbon mitigation percentage	Yes	Yes	No
Fix recommendations	Yes	Yes	No
Defer recommendations	Yes	Yes	No
<b>System manager access</b>			
May enter credentials	Yes	Yes	No
<b>Credentials</b>			
User credentials	Yes	Yes	No

## Data services roles

### BlueXP backup and recovery roles

You can assign the following roles to users to provide them access to the Backup and



recovery service within BlueXP. Backup and recovery roles give you the flexibility to assign users a role specific to the tasks they need to accomplish within your organization. How you assign roles depends on your own business and storage management practices.

Backup and recovery uses the following roles:

- **Backup and recovery super admin:** Perform any actions.
- **Backup and recovery admin:** Perform backups to local snapshots, replicate to secondary storage, and back up to object storage.
- **Backup and recovery restore admin:** Restore workloads.
- **Backup and recovery clone admin:** Clone applications and data.
- **Backup and recovery viewer:** View backup and recovery information.

The following table indicates the actions that each role can perform.

Feature and action	Backup and recovery super admin	Backup admin	Restore admin	Clone admin	Viewer
Add, edit, or delete hosts	Yes	No	No	No	No
Install plugins	Yes	No	No	No	No
Add credentials (host, instance, vCenter)	Yes	No	No	No	No
View dashboard and all tabs	Yes	Yes	Yes	Yes	Yes
Start free trial	Yes	No	No	No	No
Initiate discovery of workloads	No	Yes	Yes	Yes	No
View license information	Yes	Yes	Yes	Yes	Yes
Activate license	Yes	No	No	No	No
View hosts	Yes	Yes	Yes	Yes	Yes
<b>Schedules:</b>					
Activate schedules	Yes	Yes	Yes	Yes	No
Suspend schedules	Yes	Yes	Yes	Yes	No
<b>Policies and protection:</b>					

Feature and action	Backup and recovery super admin	Backup admin	Restore admin	Clone admin	Viewer
View protection plans	Yes	Yes	Yes	Yes	Yes
Create, modify, or delete protection plans	Yes	Yes	No	No	No
Restore workloads	Yes	No	Yes	No	No
Create, split, or delete clones	Yes	No	No	Yes	No
Create, modify, or delete policy	Yes	Yes	No	No	No
<b>Reports:</b>					
View reports	Yes	Yes	Yes	Yes	Yes
Create reports	Yes	Yes	Yes	Yes	No
Delete reports	Yes	No	No	No	No
<b>Import from SnapCenter and manage host:</b>					
View imported SnapCenter data	Yes	Yes	Yes	Yes	Yes
Import data from SnapCenter	Yes	Yes	No	No	No
Manage (migrate) host	Yes	Yes	No	No	No
<b>Configure settings:</b>					
Configure log directory	Yes	Yes	Yes	No	No
Associate or remove instance credentials	Yes	Yes	Yes	No	No
<b>Buckets:</b>					
View storage buckets	Yes	Yes	Yes	Yes	Yes
Create, edit, or delete storage buckets	Yes	Yes	No	No	No

#### BlueXP disaster recovery roles

You can assign the following roles to users to provide them access to the Disaster

recovery within BlueXP. Disaster recovery roles give you the flexibility to assign users a role specific to the tasks they need to accomplish within your organization. How you assign roles depends on your own business and storage management practices.

Disaster recovery uses the following roles:

- **Disaster recovery admin:** Perform any actions.
- **Disaster recovery failover admin:** Perform failover and migrations.
- **Disaster recovery application admin:** Create replication plans. Modify replication plans. Start test failovers.
- **Disaster recovery viewer:** View information only.

The following table indicates the actions that each role can perform.

Feature and action	Disaster recovery admin	Disaster recovery failover admin	Disaster recovery application admin	Disaster recovery viewer
View dashboard and all tabs	Yes	Yes	Yes	Yes
Start free trial	Yes	No	No	No
Initiate discovery of workloads	Yes	No	No	No
View license information	Yes	Yes	Yes	Yes
Activate license	Yes	No	Yes	No
<b>On the Sites tab:</b>				
View sites	Yes	Yes	Yes	Yes
Add, modify, or delete sites	Yes	No	No	No
<b>On the Replication plans tab:</b>				
View replication plans	Yes	Yes	Yes	Yes
View replication plan details	Yes	Yes	Yes	Yes
Create or modify replication plans	Yes	Yes	Yes	No
Create reports	Yes	No	No	No
View snapshots	Yes	Yes	Yes	Yes
Perform failover tests	Yes	Yes	Yes	No

Feature and action	Disaster recovery admin	Disaster recovery failover admin	Disaster recovery application admin	Disaster recovery viewer
Perform failovers	Yes	Yes	No	No
Perform failbacks	Yes	Yes	No	No
Perform migrations	Yes	Yes	No	No
<b>On the Resource groups tab:</b>				
View resource groups	Yes	Yes	Yes	Yes
Create, modify, or delete resource groups	Yes	No	Yes	No
<b>On the Job Monitoring tab:</b>				
View jobs	Yes	No	Yes	Yes
Cancel jobs	Yes	Yes	Yes	No

#### Ransomware protection access roles for BlueXP

Ransomware roles provide users access to the Ransomware protection service. The two roles are Ransomware protection admin and Ransomware protection viewer. The main difference between the two roles is the actions they can take in Ransomware protection.

The following table shows the actions each BlueXP ransomware protection role can perform.

Feature and action	Ransomware protection admin	Ransomware protection viewer
View dashboard and all tabs	Yes	Yes
Start free trial	Yes	No
Discover workloads	Yes	No
<b>On the Protect tab:</b>		
Add, modify, or delete policies	Yes	No
Protect workloads	Yes	No
Identify sensitive data	Yes	No
Edit workload protection	Yes	No

Feature and action	Ransomware protection admin	Ransomware protection viewer
View workload details	Yes	Yes
Download data	Yes	Yes
<b>On the Alerts tab:</b>		
View alert details	Yes	Yes
Edit incident status	Yes	No
View incident details	Yes	Yes
Get full list of impacted files	Yes	No
Download alerts data	Yes	Yes
<b>On the Recover tab:</b>		
Download impacted files	Yes	No
Restore workload	Yes	No
Download recovery data	Yes	Yes
Download reports	Yes	Yes
<b>On the Settings tab:</b>		
Add or modify backup targets	Yes	No
Add or modify SIEM targets	Yes	No
<b>On the Reports tab:</b>		
Download reports	Yes	Yes

## Identity federation

### Enable single sign-on by using identity federation with BlueXP

Single-sign on (federation) simplifies the login process and enhances security by allowing users to log in to BlueXP using their corporate credentials. You can enable single sign-on (SSO) with your identity provider (IdP) or with the NetApp Support site.

#### Required role

Organization admin, Federation admin, Federation viewer. [Learn more about access roles.](#)

## Identity federation with NetApp Support Site

When you federate with the NetApp Support Site, users can login with the same credentials to access BlueXP as you use for the NetApp Support Site, Active IQ Digital Advisor and other apps associated with your NetApp Support Site account. After you set up federation, any new users who create a NetApp Support Site accounts are also be able to access BlueXP.



If you federate with the NetApp Support Site, you can't also federate with your corporate identity management provider. Choose which one works best for your organization.

### Steps

1. Download and complete the [NetApp Federation Request Form](#).
2. Submit the form to the email address specified in the form.

The NetApp support team reviews and processes your request.

## Set up a federated connection with your identity provider

You can set up a federated connection with your identity provider to enable single sign-on (SSO) for BlueXP. The process involves configuring your identity provider to trust NetApp as a service provider and then creating the connection in BlueXP.



If you previously configured federation using NetApp Cloud Central (an external application to BlueXP), you'll need to import your federation using the BlueXP Federation page to be able to manage it within BlueXP. [Learn how to import your federation.](#)

### Supported identity providers

NetApp supports the following protocols and identity providers for federation:

#### Protocols

- Security Assertion Markup Language (SAML) identity providers
- Active Directory Federation Services (AD FS)

#### Identity providers

- Microsoft Entra ID
- PingFederate

### Federation with BlueXP workflow

NetApp supports service provider-initiated (SP-initiated) SSO only. You need to first configure the identity provider to trust NetApp as a service provider. Then, you can create a connection in BlueXP that uses the identity provider's configuration.

You can federate with your email domain or with a different domain that you own. To federate with a domain different from your email domain, first verify you own the domain.



**Verify your domain (if not using your email domain)**

To federate with a domain different from your email domain, verify that you own it. You can federate your email domain without any extra steps.

2

### Configure your IdP to trust NetApp as a service provider

Configure your identity provider to trust NetApp by creating a new application and providing the necessary information, such as the ACS URL, Entity ID or other credential information. Service provider information varies by identity provider, so refer to the documentation for your specific identity provider for details. You'll need to work with your IdP administrator to complete this step.

3

### Create the federated connection in BlueXP

To create the connection, you need to provide the necessary information from your identity provider, such as the SAML metadata URL or file. This information is used to establish the trust relationship between BlueXP and your identity provider. The information you provide depends on the IdP that you are using. For example, if you're using Microsoft Entra ID, you need to provide the client ID, secret, and domain.

4

### Test your federation in BlueXP

Test your federated connection before enabling it. The Federation page in BlueXP provides a test option that allows you to verify your test user is able to authenticate successfully. If the test is successful, you can enable the connection.

5

### Enable your connection in BlueXP

After you enable the connection, users can log in to BlueXP using their corporate credentials.

Review the topic for your respective protocol or IdP to get started:

- [Set up a federated connection with AD FS](#)
- [Set up a federated connection with Microsoft Entra ID](#)
- [Set up a federated connection with PingFederate](#)
- [Set up a federated connection with a SAML identity provider](#)

## Domain verification

### Verify the email domain for your federated connection


If you want to federate with a domain that is different than your email domain, you must first verify that you own the domain. You can only use verified domains for federation.

#### Required roles

Organization admin or Federation admin. [Learn more about access roles.](#)

Verifying your domain involves adding a TXT record to your domain's DNS settings. This record is used to prove that you own the domain and allows BlueXP to trust the domain for federation. You may need to coordinate with your IT or network administrator to complete this step.

#### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select **Configure new federation**.
4. Select **Verify domain ownership**.
5. Enter the domain that you want to verify and select **Continue**.
6. Copy the TXT record that is provided.
7. Go to your domain's DNS settings and configure the TXT value that was provided as a TXT record for your domain. Work with your IT or network administrator if needed.
8. After the TXT record is added, return to BlueXP and select **Verify**.

## Configure federations

### Federate BlueXP with Active Directory Federation Services (AD FS)

Federate your Active Directory Federation Services (AD FS) with BlueXP to enable single sign-on (SSO) for BlueXP. This allows users to log in to BlueXP using their corporate credentials.

#### Required roles

Organization admin or Federation admin is required to make create and manage federations. Federation viewer can view the Federation page. [Learn more about access roles](#).



You can federate with your corporate IdP or with the NetApp Support Site. NetApp recommends choosing one or the other, but not both.


NetApp supports service provider-initiated (SP-initiated) SSO only. First, configure the identity provider to trust BlueXP as a service provider. Then, create a connection in BlueXP using your identity provider's configuration.

You can set up federation with your AD FS server to enable single sign-on (SSO) for BlueXP. The process involves configuring your AD FS to trust BlueXP as a service provider and then creating the connection in BlueXP.

#### Before you begin

- An IdP account with administrative privileges is required. Coordinate with your IdP administrator to complete the steps.
- Identify the domain you want to use for federation. You can use your email domain or a different domain that you own. If you want to use a domain other than your email domain, you must first verify the domain in BlueXP. You can do this by following the steps in the [Verify your domain in BlueXP](#) topic.

#### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select **Configure new federation**.
4. Enter your domain details:
  - a. Choose whether you want to use a verified domain or your email domain. The email domain is the domain associated with the account you are logged in with.



- b. Enter the name of the federation you are configuring.
  - c. If you choose a verified domain, select the domain from the list.
5. Select **Next**.
  6. For your connection method, choose **Protocol** and then select **Active Directory Federation Services (AD FS)**.
  7. Select **Next**.
  8. Create a Relying Party Trust in your AD FS server. You can use PowerShell or manually configure it on your AD FS server. Consult the AD FS documentation for details on how to create a relying party trust.
    - a. Create the trust using PowerShell by using following script:

```
(new-object Net.WebClient -property @{Encoding = [Text.Encoding]
::UTF8}) .DownloadString("https://raw.githubusercontent.com/auth0/AD FS-
auth0/master/AD FS.ps1") | iex
AddRelyingParty "urn:auth0:netapp-cloud-account" "https://netapp-
cloud-account.auth0.com/login/callback"
```

- b. Alternatively, you can create the trust manually in the AD FS management console. Use the following BlueXP values when creating the trust:
  - When creating the Relying Trust Identifier, use the **YOUR\_TENANT** value: netapp-cloud-account
  - When you select **Enable support for the WS-Federation**, use the **YOUR\_AUTH0\_DOMAIN** value: netapp-cloud-account.auth0.com
- c. After creating the trust, copy the metadata URL from your AD FS server or download the federation metadata file. You'll need this URL or file to complete the connection in BlueXP.

NetApp recommends using the metadata URL to let BlueXP automatically retrieve the latest AD FS configuration. If you download the federation metadata file, you will need to update it manually in BlueXP whenever there are changes to your AD FS configuration.

9. Return to BlueXP, and select **Next** to create the connection.
10. Create the connection with AD FS.
  - a. Enter the **AD FS URL** that you copied from your AD FS server in the previous step or upload the federation metadata file that you downloaded from your AD FS server.
11. Select **Create connection**. Creating the connection might take a few seconds.
12. Select **Next**.
13. Select **Test connection** to test your connection. You are directed to a login page for your IdP server. Log in with your IdP credentials to complete the test and return to BlueXP to enable the connection.
14. Select **Next**.
15. On the **Enable federation** page, review the federation details and then select **Enable federation**.
16. Select **Finish** to complete the process.

After you enable the federation, users can log in to BlueXP using their corporate credentials.

## Federate BlueXP with Microsoft Entra ID

Federate with your Microsoft Entra ID IdP provider to enable single sign-on (SSO) for BlueXP. This allows users to log in using their corporate credentials.

### Required roles

Organization admin or Federation admin is required to make create and manage federations. Federation viewer can view the Federation page. [Learn more about access roles.](#)



You can federate with your corporate IdP or with the NetApp Support Site. NetApp recommends choosing one or the other, but not both.


NetApp supports service provider-initiated (SP-initiated) SSO only. You need to first configure the identity provider to trust NetApp as a service provider. Then, you can create a connection in BlueXP that uses the identity provider's configuration.

You can set up a federated connection with Microsoft Entra ID to enable single sign-on (SSO) for BlueXP. The process involves configuring your Microsoft Entra ID to trust BlueXP as a service provider and then creating the connection in BlueXP.

### Before you begin

- An IdP account with administrative privileges is required. Coordinate with your IdP administrator to complete the steps.
- Identify the domain you want to use for federation. You can use your email domain or a different domain that you own. If you want to use a domain other than your email domain, you must first verify the domain in BlueXP. You can do this by following the steps in the [Verify your domain in BlueXP](#) topic.

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select **Configure new federation**.

### Domain details

4. Enter your domain details:
  - a. Choose whether you want to use a verified domain or your email domain. The email domain is the domain associated with the account you are logged in with.
  - b. Enter the name of the federation you are configuring.
  - c. If you choose a verified domain, select the domain from the list.
5. Select **Next**.

### Connection method

6. For your connection method, choose **Provider** and then select **Microsoft Entra ID**.
7. Select **Next**.

### Configuration instructions

1. Configure your Microsoft Entra ID to trust NetApp as a service provider. You need to do this step on your

Microsoft Entra ID server.

- a. Use the following values when registering your Microsoft Entra ID app to trust BlueXP:
  - For the **Redirect URL**, use <https://services.cloud.netapp.com>
  - For the **Reply URL**, use <https://netapp-cloud-account.auth0.com/login/callback>
- b. Create a client secret for your Microsoft Entra ID app. You'll need to provide the client ID, the client secret, and the Entra ID domain name to complete the federation.

2. Return to BlueXP, and select **Next** to create the connection.

### Create connection

1. Create the connection with Microsoft Entra ID
  - a. Enter the client ID and Client secret that you created in the previous step.
  - b. Enter the Microsoft Entra ID domain name.
2. Select **Create connection**. The system creates the connection in a few seconds.

### Test and enable the connection

1. Select **Next**.
2. Select **Test connection** to test your connection. You are directed to a login page for your IdP server. Log in with your IdP credentials to complete the test and return to BlueXP to enable the connection.
3. Select **Next**.
4. On the **Enable federation** page, review the federation details and then select **Enable federation**.
5. Select **Finish** to complete the process.

After you enable the federation, users can log in to BlueXP using their corporate credentials.

### Federate BlueXP with PingFederate

Federate with your PingFederate IdP provider to enable single sign-on (SSO) for BlueXP. This allows users to log in using their corporate credentials.

#### Required roles

Organization admin or Federation admin is required to make create and manage federations. Federation viewer can view the Federation page. [Learn more about access roles.](#)



You can federate with your corporate IdP or with the NetApp Support Site. NetApp recommends choosing one or the other, but not both.


NetApp supports service provider-initiated (SP-initiated) SSO only. You need to first configure the identity provider to trust NetApp as a service provider. Then, you can create a connection in BlueXP that uses the identity provider's configuration.

You can set up a federated connection with PingFederate to enable single sign-on (SSO) for BlueXP. The process involves configuring your PingFederate server to trust BlueXP as a service provider and then creating the connection in BlueXP.

### Before you begin

- An IdP account with administrative privileges is required. Coordinate with your IdP administrator to complete the steps.
- Identify the domain you want to use for federation. You can use your email domain or a different domain that you own. If you want to use a domain other than your email domain, you must first verify the domain in BlueXP. You can do this by following the steps in the [Verify your domain in BlueXP](#) topic.

## Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select **Configure new federation**.
4. Enter your domain details:
  - a. Choose whether you want to use a verified domain or your email domain. The email domain is the domain associated with the account you are logged in with.
  - b. Enter the name of the federation you are configuring.
  - c. If you choose a verified domain, select the domain from the list.
5. Select **Next**.
6. For your connection method, choose **Provider** and then select **PingFederate**.
7. Select **Next**.
8. Configure your PingFederate server to trust NetApp as a service provider. You need to do this step on your PingFederate server.
  - a. Use the following values when configuring PingFederate to trust BlueXP:
    - For the **Reply URL** or **Assertion Consumer Service (ACS) URL**, use <https://netapp-cloud-account.auth0.com/login/callback>
    - For the **Logout URL**, use <https://netapp-cloud-account.auth0.com/logout>
    - For **Audience/Entity ID**, use `urn:auth0:netapp-cloud-account:<fed-domain-name-saml>` where `<fed-domain-name-pingfederate>` is the domain name for the federation. For example, if your domain is `example.com`, the Audience/Entity ID would be `urn:auth0:netappcloud-account:fed-example-com-pingfederate`.
  - b. Copy the PingFederate server URL. You will need this URL when creating the connection in BlueXP.
  - c. Download the X.509 certificate from your PingFederate server. It needs to be in Base64-encoded PEM format (.pem, .crt, .cer).
9. Return to BlueXP, and select **Next** to create the connection.
10. Create the connection with PingFederate
  - a. Enter the PingFederate server URL that you copied in the previous step.
  - b. Upload the X.509 signing certificate. The certificate must be in PEM, CER, or CRT format.
11. Select **Create connection**. The system creates the connection in a few seconds.
12. Select **Next**.
13. Select **Test connection** to test your connection. You are directed to a login page for your IdP server. Log in with your IdP credentials to complete the test and return to BlueXP to enable the connection.
14. Select **Next**.
15. On the **Enable federation** page, review the federation details and then select **Enable federation**.

16. Select **Finish** to complete the process.

After you enable the federation, users can log in to BlueXP using their corporate credentials.

### Federate with a SAML identity provider

Federate with your SAML 2.0 IdP provider to enable single sign-on (SSO) for BlueXP. This allows users to log in using their corporate credentials.

#### Required role

Organization admin. [Learn more about access roles.](#)



You can federate with your corporate IdP or with the NetApp Support Site. You can't federate with both.


NetApp supports service provider-initiated (SP-initiated) SSO only. You need to first configure the identity provider to trust NetApp as a service provider. Then, you can create a connection in BlueXP that uses the identity provider's configuration.

You can set up a federated connection with your SAML 2.0 provider to enable single sign-on (SSO) for BlueXP. The process involves configuring your provider to trust NetApp as a service provider and then creating the connection in BlueXP.

#### Before you begin

- An IdP account with administrative privileges is required. Coordinate with your IdP administrator to complete the steps.
- Identify the domain you want to use for federation. You can use your email domain or a different domain that you own. If you want to use a domain other than your email domain, you must first verify the domain in BlueXP. You can do this by following the steps in the [Verify your domain in BlueXP](#) topic.

#### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select **Configure new federation**.
4. Enter your domain details:
  - a. Choose whether you want to use a verified domain or your email domain. The email domain is the domain associated with the account you are logged in with.
  - b. Enter the name of the federation you are configuring.
  - c. If you choose a verified domain, select the domain from the list.
5. Select **Next**.
6. For your connection method, choose **Protocol** and then select **SAML Identity Provider**.
7. Select **Next**.
8. Configure your SAML identity provider to trust NetApp as a service provider. You need to do this step on your SAML provider server.
  - a. Ensure that your IdP has the attribute `email` set to the user's email address. This is required for BlueXP to identify users correctly:

```

<saml:AttributeStatement
  xmlns:x500="urn:oasis:names:tc:SAML:2.0:profiles:attribute:X500"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <saml:Attribute Name="email"
    NameFormat="urn:oasis:names:tc:SAML:1.1:nameid-
    format:X509SubjectName">
    <saml:AttributeValue xsi:type="xs:string">
    email@domain.com</saml:AttributeValue>
  </saml:Attribute>
</saml:AttributeStatement>

```

- b. Use the following values when registering your SAML application with BlueXP:
    - For the **Reply URL** or **Assertion Consumer Service (ACS) URL**, use <https://netapp-cloud-account.auth0.com/login/callback>
    - For the **Logout URL**, use <https://netapp-cloud-account.auth0.com/logout>
    - For **Audience/Entity ID**, use `urn:auth0:netapp-cloud-account:<fed-domain-name-saml>` where `<fed-domain-name-saml>` is the domain name you want to use for federation. For example, if your domain is `example.com`, the Audience/Entity ID would be `urn:auth0:netapp-cloud-account:fed-example-com-samlp`.
  - c. After creating the trust, copy the following values from your SAML provider server:
    - Sign In URL
    - Sign Out URL (optional)
  - d. Download the X.509 certificate from your SAML provider server. It needs to be in PEM, CER, or CRT format.
9. Return to BlueXP, and select **Next** to create the connection.
  10. Create the connection with SAML.
    - a. Enter the **Sign In URL** of your SAML server.
    - b. Upload the X.509 certificate that you downloaded from your SAML provider server.
    - c. Optionally, enter the **Sign Out URL** of your SAML server.
  11. Select **Create connection**. The system creates the connection in a few seconds.
  12. Select **Next**.
  13. Select **Test connection** to test your connection. You are directed to a login page for your IdP server. Log in with your IdP credentials to complete the test and return to BlueXP to enable the connection.
  14. Select **Next**.
  15. On the **Enable federation** page, review the federation details and then select **Enable federation**.
  16. Select **Finish** to complete the process.

After you enable the federation, users can log in to BlueXP using their corporate credentials.

## Manage federations in BLueXP

You can manage your federation in BlueXP. You can disable it, update expired credentials, as well as disable it if you no longer need it.



If you previously configured federation using NetApp Cloud Central (an external application to BlueXP), you'll need to import your federation using the BlueXP Federation page to be able to manage it within BlueXP. [Learn how to import your federation](#)

You can also add a verified domain to an existing federation, which allows you to use multiple domains for your federated connection.



Federation management events such as enabling, disabling, and updating federations display in the Timeline. [Learn more about monitoring operations in BlueXP.](#)



### Required roles

Organization admin or Federation admin is required to make create and manage federations. Federation viewer can view the Federation page. [Learn more about access roles.](#)

### Enable a federation

If you have created a federation but it is not enabled, you can enable it through the Federation tab in BlueXP. Enabling a federation allows users associated with the federation to log in to BlueXP using their corporate credentials. You must have already created the federation and tested it successfully before enabling it.

### Steps


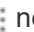
1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select the actions menu  next to the federation that you want to enable and select **Enable**.

### Add a verified domain to an existing federation

You can add a verified domain to an existing federation in BlueXP to use multiple domains with the same identity provider (IdP).

You must have already verified the domain in BlueXP before you can add it to a federation. If you haven't verified the domain yet, you can do so by following the steps in [Verify your domain in BlueXP](#).

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select the actions menu  next to the federation that you want to add a verified domain to and select **Update domains**. The **Update domains** dialog box lists the domain already associated with this federation.
4. Select a verified domain from the list of available domains.
5. Select **Update**. It may take up to 30 seconds for users of the new domain to have federated access to BlueXP.


## Updating an expiring federated connection

You can update the details of a federation in BlueXP. For example, you'll need to update the federation if the credentials such as a certificate or client secret expire. When needed, update the notification date to remind you to update the connection before it expires.



Update BlueXP first before updating your IdP to avoid login issues. Stay logged in to BlueXP during the process.



### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select the actions menu (three vertical dots) next to the federation that you want to update and select **Update federation**.
4. Update the details of the federation as needed.
5. Select **Update**.

## Test an existing federation

If you are having trouble with an existing federation, you can test the connection to see if it is working properly. This can help you identify any issues with the federation and troubleshoot them.

### Steps



1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select the actions menu  next to the federation that you want to add a verified domain to and select **Test connection**.
4. Select **Test**. You're prompted to log in with your corporate credentials. If the connection is successful, you will be redirected to the BlueXP console. If the connection fails, you will see an error message indicating the issue with the federation.
5. Select **Done** to return to the **Federation** tab.

## Disable a federation

If you no longer need a federation, you can disable it. This prevents users associated with the federation from logging in to BlueXP using their corporate credentials. You can re-enable the federation later if needed.

You should disable a federation before deleting it. For example, if you are decommissioning the IdP in favor of another IdP or no longer want to use federation. This allows you to re-enable it later if needed.

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select the actions menu  next to the federation that you want to add a verified domain to and select **Disable**.





## Delete a federation

If you no longer need a federation, you can delete it. This removes the federation from BlueXP and prevents any users associated with the federation from logging in to BlueXP using their corporate credentials. For example, if the IdP is being decommissioned or if the federation is no longer needed. After you delete a federation, you cannot recover it. You must create a new federation.



You must disable a federation before you can delete it. You cannot undelete a federation after you delete it.

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select the actions menu  next to the federation that you want to add a verified domain to and select **Delete**.

## Import your federation to BlueXP

If you have previously setup federation through NetApp Cloud Central (an external application to BlueXP) the Federation page prompts you to import your existing federated connection to BlueXP to manage it in the new interface. This allows you to take advantage of the latest enhancements without having to recreate your federated connections.

Existing customers who have already set up federated connections to BlueXP can import their existing federations to the new interface. This allows you to manage your federated connections in the new Federations page without having to recreate them.




After you import your existing federation, you can manage the federation from the Federations page. [Learn more about managing federations.](#)

### Required role

Organization admin or Federation admin. [Learn more about access roles.](#)

### Steps

1. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
2. Select the **Federation** tab.
3. Select **Import Federation**.

## Connectors

### Maintain the Connector VM and operating system

Maintaining the operating system on the Connector host is your (the customer's) responsibility. For example, you (the customer) should apply security updates to the operating system on the Connector host by following your company's standard procedures for operating system distribution.



If you have an existing Connector, you should be aware of [changes to supported Linux operating systems](#).

## Operating system patches and the Connector

Apply OS security patches without stopping Connector host services.

## VM or instance type

If you create a Connector from BlueXP, it deploys a VM instance in your cloud provider with a default configuration. After you create the Connector, don't switch to a smaller VM instance with less CPU or RAM.

The following table lists the CPU and RAM requirements:

### CPU

8 cores or 8 vCPUs

### RAM

32 GB

[Learn about the default configuration for the Connector.](#)

## Monitor the Connector

BlueXP notifies you when the Connector VM is unhealthy, including disk space, RAM, and CPU issues. Monitor these notifications in the Notifications Center within BlueXP or configure email notifications. Occasional increases in disk space, memory, or CPU usage are normal, but if it happens frequently, you should take steps to resolve.

BlueXP notifies you when a Connector resource (CPU, RAM, or disk space) exceeds 90% of its total capacity for 30 consecutive minutes. Afterwards, if the resource usage drops below that threshold, the notification displays as resolved (green) in the Notifications Center.



Work with NetApp support if you have questions about modifying your Connector VM.

[Learn more.](#)

Notification	Action needed
Disk space is too high	<a href="#">Review the NetApp Knowledge Base article.</a>
CPU usage is too high	Increase the CPU size of the Connector VM in your hyperscaler or on-premises, depending on where you installed it. Alternatively, create additional Connectors and distribute the workload across multiple Connectors. RAM utilization can vary based on your environment, ONTAP workloads, number of Cloud Volumes ONTAP systems, and the data services that you are using.

Notification	Action needed
RAM usage is too high	Increase the RAM of the Connector VM in your hyperscaler or on-premises, depending on where you installed it. Alternatively, create additional Connectors and distribute the workload across multiple Connectors. RAM utilization can vary based on your environment, ONTAP workloads, number of Cloud Volumes ONTAP systems, and the data services that you are using.

## Stopping and starting the Connector VM

If you need to, stop and start the Connector VM using your cloud provider's console or standard on-premises procedures.

Be aware that the Connector must be operational at all times.

## Connect to the Linux VM

If you need to connect to the Linux VM that the Connector runs on, use the connectivity options from your cloud provider.

### AWS

When you create the Connector instance in AWS, provide an AWS access key and secret key. You can use this key pair to SSH to the instance. Use the user name 'ubuntu' for the EC2 Linux instance. For Connectors created prior to May 2023, use the user name 'ec2-user'.

[AWS Docs: Connect to your Linux instance](#)

### Azure

When you create the Connector VM in Azure, you specify a user name and choose to authenticate with a password or SSH public key. Use the authentication method that you chose to connect to the VM.

[Azure Docs: SSH into your VM](#)

### Google Cloud

You can't specify an authentication method when you create a Connector in Google Cloud. However, you can connect to the Linux VM instance using the Google Cloud Console or Google Cloud CLI (gcloud).

[Google Cloud Docs: Connect to Linux VMs](#)

## Change the IP address for a Connector

You can change the internal and public IP addresses of the Connector instance assigned by your cloud provider if needed.

### Steps

1. Follow the instructions from your cloud provider to change the local IP address or public IP address (or both) for the Connector instance.
2. Restart the Connector instance to register a new public IP address with BlueXP.
3. If you changed the private IP address, update the backup location for Cloud Volumes ONTAP configuration files so that the backups are being sent to the new private IP address on the Connector.

Update the backup location for each Cloud Volumes ONTAP system.

- a. From the Cloud Volumes ONTAP CLI, set the privilege level to advanced:

```
set -privilege advanced
```

- b. Run the following command to display the current backup target:

```
system configuration backup settings show
```

- c. Run the following command to update the IP address for the backup target:

```
system configuration backup settings modify -destination <target-  
location>
```

## Edit a Connector's URIs

You can add and remove the Uniform Resource Identifier (URI) for a Connector.

### Steps

1. Select the **Connector** drop-down from the BlueXP header.
2. Select **Manage Connectors**.
3. Expand the **Connector URIs** bar to view connector URIs.
4. Add and remove URIs and then select **Apply**.

## Install a CA-signed certificate for web-based console access

When you use BlueXP in restricted mode or private mode, the user interface is accessible from the Connector virtual machine that's deployed in your cloud region or on-premises. By default, BlueXP uses a self-signed SSL certificate to provide secure HTTPS access to the web-based console running on the Connector. If required by your business, you can install a certificate signed by a certificate authority (CA), which provides better security protection than a self-signed certificate. After you install the certificate, BlueXP uses the CA-signed certificate when users access the web-based console.

### Before you begin

You need to create a Connector before you can change BlueXP settings. [Learn how to create a Connector](#).

### Install an HTTPS certificate

Install a certificate signed by a CA for secure access to the web-based console running on the Connector.

### About this task

You can install the certificate using one of the following options:

- Generate a certificate signing request (CSR) from BlueXP, submit the certificate request to a CA, and then install the CA-signed certificate on the Connector.

The key pair that BlueXP uses to generate the CSR is stored internally on the Connector. BlueXP automatically retrieves the same key pair (private key) when you install the certificate on the Connector.

- Install a CA-signed certificate that you already have.

With this option, the CSR is not generated through BlueXP. You generate the CSR separately and store the private key externally. You provide BlueXP with the private key when you install the certificate.

## Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **HTTPS Setup**.

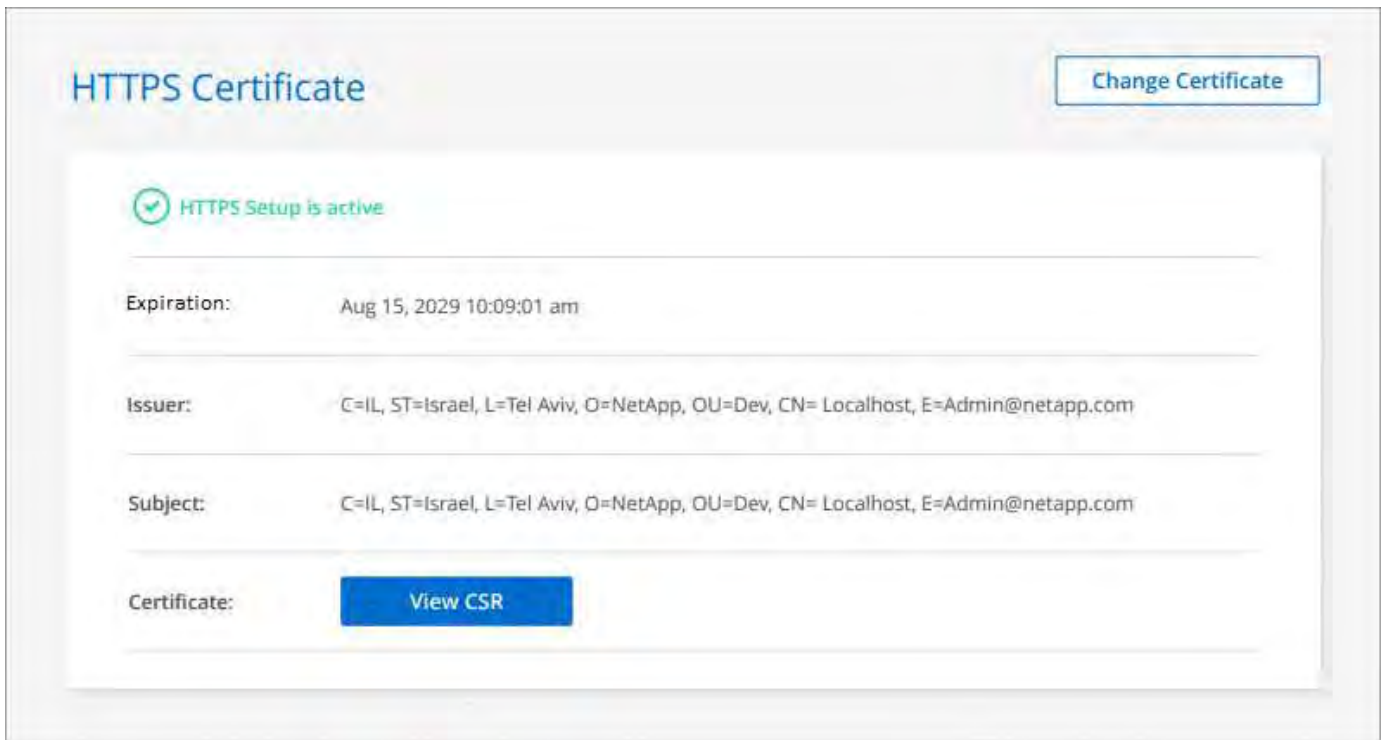


2. In the HTTPS Setup page, install a certificate by generating a certificate signing request (CSR) or by installing your own CA-signed certificate:

Option	Description
Generate a CSR	<ol style="list-style-type: none"> <li>Enter the host name or DNS of the Connector host (its Common Name), and then select <b>Generate CSR</b>.  BlueXP displays a certificate signing request.</li> <li>Use the CSR to submit an SSL certificate request to a CA.  The certificate must use the Privacy Enhanced Mail (PEM) Base-64 encoded X.509 format.</li> <li>Upload the certificate file and then select <b>Install</b>.</li> </ol>
Install your own CA-signed certificate	<ol style="list-style-type: none"> <li>Select <b>Install CA-signed certificate</b>.</li> <li>Load both the certificate file and the private key and then select <b>Install</b>.  The certificate must use the Privacy Enhanced Mail (PEM) Base-64 encoded X.509 format.</li> </ol>

## Result

BlueXP now uses the CA-signed certificate to provide secure HTTPS access. The following image shows a Connector that is configured for secure access:



## Renew the BlueXP HTTPS certificate

You should renew the BlueXP HTTPS certificate before it expires to ensure secure access to the BlueXP console. If you don't renew the certificate before it expires, a warning appears when users access the web console using HTTPS.

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **HTTPS Setup**.

Details about the BlueXP certificate displays, including the expiration date.

2. Select **Change Certificate** and follow the steps to generate a CSR or install your own CA-signed certificate.

### Result

BlueXP uses the new CA-signed certificate to provide secure HTTPS access.

## Configure a Connector to use a proxy server

If your corporate policies require you to use a proxy server for all communication to the internet, then you need to configure your Connectors to use that proxy server. If you didn't configure a Connector to use a proxy server during installation, then you can configure the Connector to use that proxy server at any time.

The Connector's proxy server enables outbound internet access without a public IP or NAT gateway. The proxy server provides outbound connectivity only for the Connector, not for Cloud Volumes ONTAP systems.

If Cloud Volumes ONTAP systems lack outbound internet access, BlueXP configures them to use the Connector's proxy server. You must ensure that the Connector's security group allows inbound connections over port 3128. Open this port after deploying the Connector.

If the Connector itself doesn't have an outbound internet connection, Cloud Volumes ONTAP systems cannot use the configured proxy server.

### Supported configurations

- Transparent proxy servers are supported for Connectors that serve Cloud Volumes ONTAP systems. If you use BlueXP services with Cloud Volumes ONTAP, create a dedicated Connector for Cloud Volumes ONTAP where you can use a transparent proxy server.
- Explicit proxy servers are supported with all Connectors, including those that manage Cloud Volumes ONTAP systems and those that manage BlueXP services.
- HTTP and HTTPS.
- The proxy server can reside in the cloud or in your network.



Once you have configured a proxy, you cannot change the proxy type. If you need to change the proxy type, you remove the Connector and add a new Connector with the new proxy type.

### Enable an explicit proxy on a Connector

When you configure a Connector to use a proxy server, that Connector and the Cloud Volumes ONTAP systems that it manages (including any HA mediators), all use the proxy server.

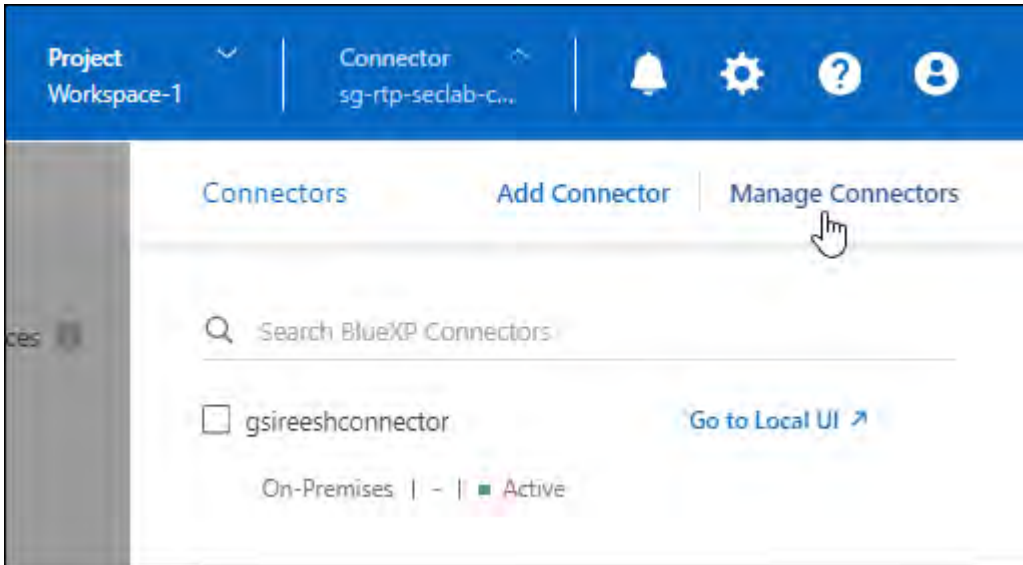
This operation restarts the Connector. Verify the Connector is idle before proceeding.

### Steps

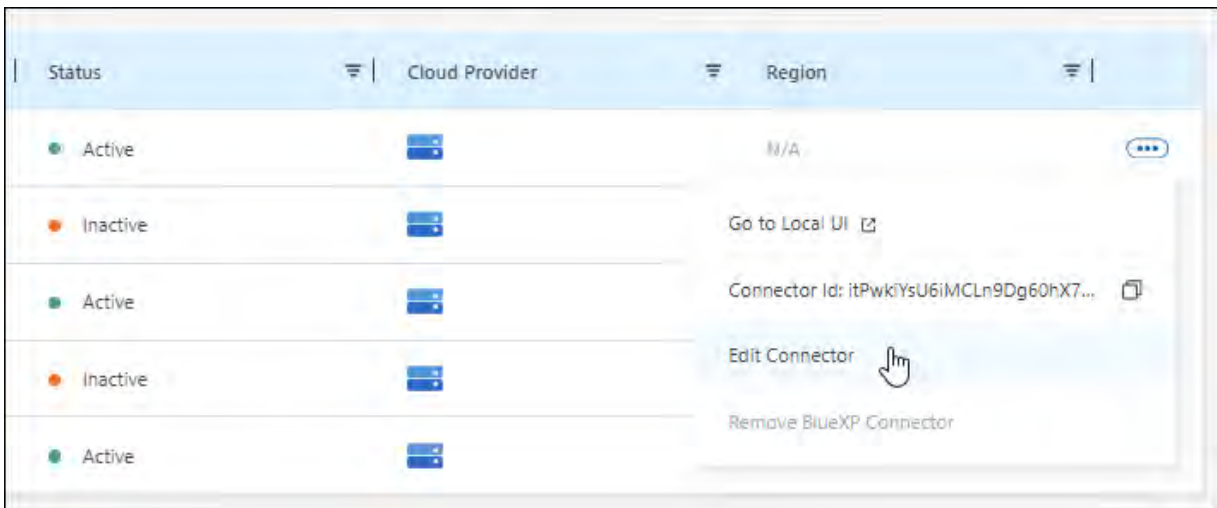
1. Navigate to the **Edit BlueXP Connector** page.

**Standard mode**

- a. Select the **Connector** drop-down from the BlueXP header.
- b. Select **Manage Connectors**.



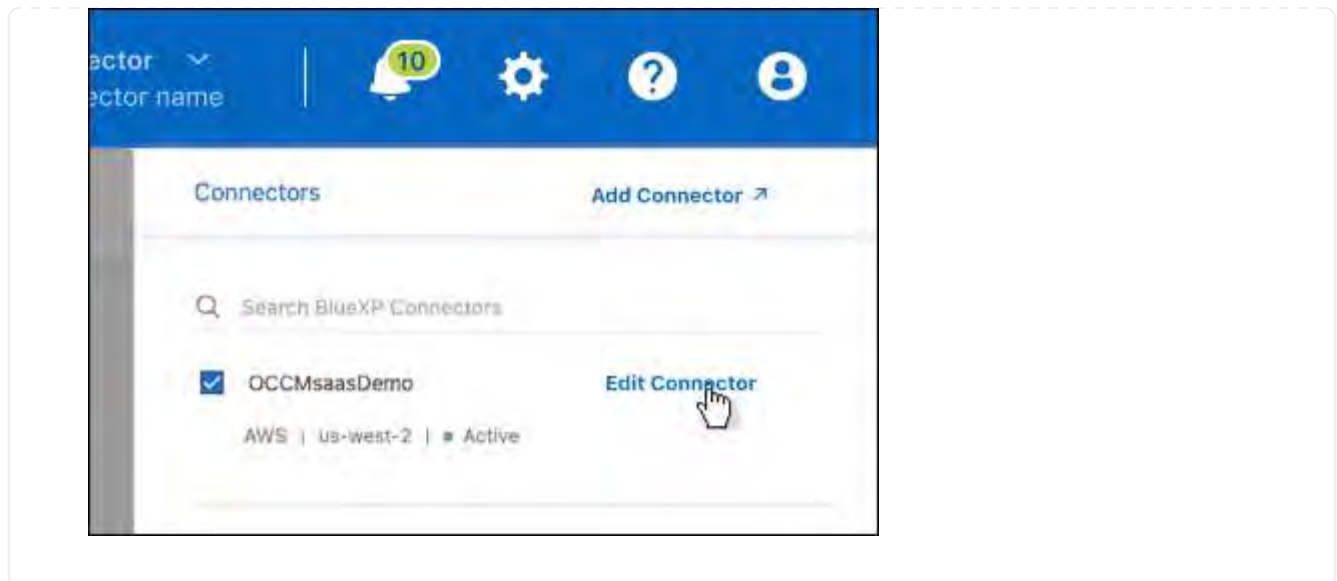
- c. Select the action menu for a Connector and select **Edit Connector**.



**Restricted or private mode**

- a. Select the **Connector** drop-down from the BlueXP header.
- b. Select **Edit Connector**.





2. Select **HTTP Proxy Configuration**.
3. Select **Explicit proxy** in the Configuration type field.
4. Select **Enable Proxy**.
5. Specify the server using the syntax `http://address:port` or `https://address:port`
6. Specify a user name and password if basic authentication is required for the server.

Note the following:

- The user can be a local user or domain user.
- For a domain user, you must enter the ASCII code for the \ as follows: `domain-name%92user-name`  
For example: `netapp%92proxy`
- BlueXP doesn't support passwords that include the @ character.

7. Select **Save**.

### Enable a transparent proxy on a Connector

Only Cloud Volumes ONTAP supports using a transparent proxy on the Connector. If you use BlueXP services in addition to Cloud Volumes ONTAP, you should create a separate Connector to use for data services or to use for Cloud Volumes ONTAP.

Before enabling a transparent proxy, ensure that the following requirements are met:

- The Connector is installed on the same network as the transparent proxy server.
- TLS inspection is enabled on the proxy server.
- You have a certificate in PEM format that matches the one used on the transparent proxy server.
- You do not use the Connector for any NetApp data services other than Cloud Volumes ONTAP.

To configure an existing Connector to use a transparent proxy server, you use the Connector maintenance tool that is available through the command line on the Connector host.

When you configure a proxy server, the Connector restarts. Verify the Connector is idle before proceeding.

## Steps

Ensure that you have a certificate file in PEM format for the proxy server. If you do not have a certificate, contact your network administrator to obtain one.

1. Open a command-line interface on the Connector host.
2. Navigate to the Connector maintenance tool directory: `/opt/application/netapp/service-manager-2/connector-maint-console`
3. Run the following command to enable the transparent proxy, where `/home/ubuntu/<certificate-file>.pem` is the directory and name certificate file that you have for the proxy server:

```
./connector-maint-console proxy add -c /home/ubuntu/<certificate-  
file>.pem
```

Ensure that the certificate file is in PEM format and resides in the same directory as the command or specify the full path to the certificate file.

```
./connector-maint-console proxy add -c /home/ubuntu/<certificate-  
file>.pem
```

## Modify the transparent proxy for the Connector

You can update a Connector's existing transparent proxy server by using the `proxy update` command or remove the transparent proxy server by using the `proxy remove` command. For more information, review the documentation for [Connector maintenance console](#).



Once you have configured a proxy, you cannot change the proxy type. If you need to change the proxy type, you remove the Connector and add a new Connector with the new proxy type.

## Update the Connector proxy if it loses access to the internet

If the proxy configuration for your network changes, your Connector might lose access to the internet. For example, if someone changes the password for the proxy server or updates the certificate. In this case, you'll need to access the UI from the Connector host directly and update the settings. Ensure you have network access to the Connector host and that you can log into the BlueXP UI.

## Enable direct API traffic

If you configured a Connector to use a proxy server, you can enable direct API traffic on the Connector in order to send API calls directly to cloud provider services without going through the proxy. Connectors running in AWS, Azure, or Google Cloud support this option.

If you disable Azure Private Links with Cloud Volumes ONTAP and use service endpoints, enable direct API traffic. Otherwise, the traffic won't be routed properly.

[Learn more about using an Azure Private Link or service endpoints with Cloud Volumes ONTAP](#)

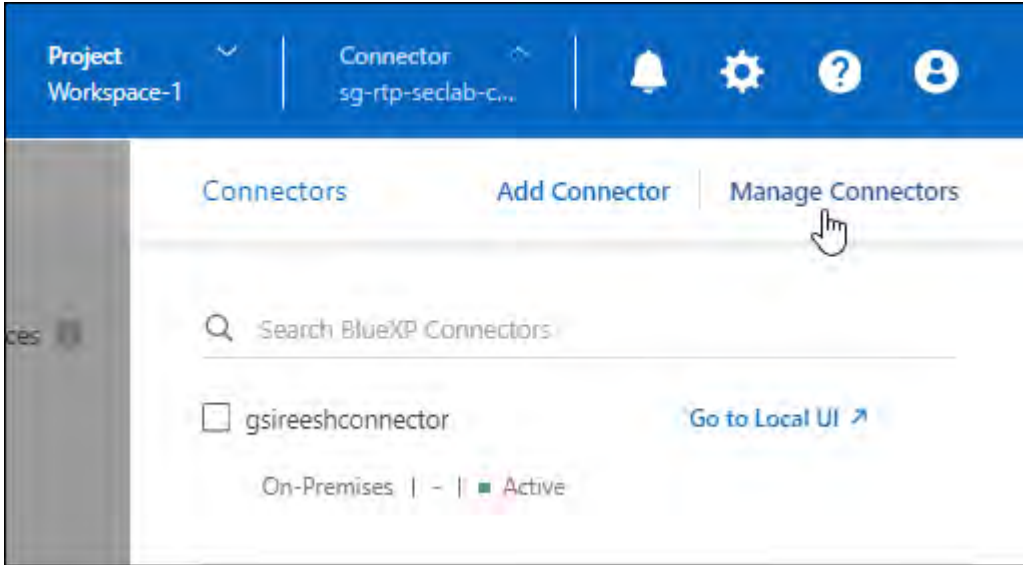
## Steps

1. Navigate to the **Edit BlueXP Connector** page:

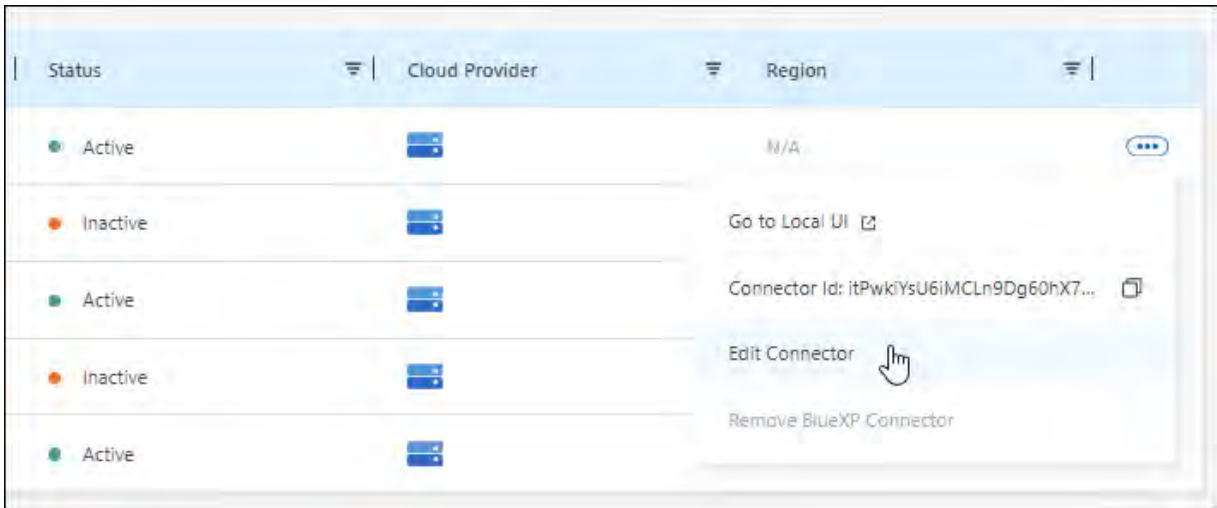
Navigation depends on your BlueXP mode. In standard mode, access the interface from the SaaS website. In restricted or private mode, access it locally from the Connector host.

**Standard mode**

- a. Select the **Connector** drop-down from the BlueXP header.
- b. Select **Manage Connectors**.




- c. Select the action menu for a Connector and select **Edit Connector**.



**Restricted or private mode**

- a. Select the **Connector** drop-down from the BlueXP header.
- b. Select **Edit Connector**.



2. Select **Support Direct API Traffic**.
3. Select the checkbox to enable the option and then select **Save**.

## Require the use of IMDSv2 on Amazon EC2 instances

BlueXP supports the Amazon EC2 Instance Metadata Service Version 2 (IMDSv2) with the Connector and with Cloud Volumes ONTAP (including the mediator for HA deployments). In most cases, IMDSv2 is automatically configured on new EC2 instances. IMDSv1 was enabled prior to March 2024. If required by your security policies, you might need to manually configure IMDSv2 on your EC2 instances.

### Before you begin

- The Connector version must be 3.9.38 or later.
- Cloud Volumes ONTAP must be running one of the following versions:
  - 9.12.1 P2 (or any subsequent patch)
  - 9.13.0 P4 (or any subsequent patch)
  - 9.13.1 or any version after this release
- This change requires you to restart the Cloud Volumes ONTAP instances.
- These steps require the use of the AWS CLI because you must change the response hop limit to 3.

### About this task

IMDSv2 provides enhanced protection against vulnerabilities. [Learn more about IMDSv2 from the AWS Security Blog](#)

The Instance Metadata Service (IMDS) is enabled as follows on EC2 instances:

- For new Connector deployments from BlueXP or using [Terraform scripts](#), IMDSv2 is enabled by default on the EC2 instance.
- If you launch a new EC2 instance in AWS and then manually install the Connector software, IMDSv2 is also enabled by default.
- If you launch the Connector from the AWS Marketplace, IMDSv1 is enabled by default. You can manually

configure IMDSv2 on the EC2 instance.

- For existing Connectors, IMDSv1 is still supported but you can manually configure IMDSv2 on the EC2 instance if you prefer.
- For Cloud Volumes ONTAP, IMDSv1 is enabled by default on new and existing instances. You can manually configure IMDSv2 on the EC2 instances if you prefer.

## Steps

1. Require the use of IMDSv2 on the Connector instance:

a. Connect to the Linux VM for the Connector.

When you created the Connector instance in AWS, you provided an AWS access key and secret key. You can use this key pair to SSH to the instance. The user name for the EC2 Linux instance is ubuntu (for Connectors created prior to May 2023, the user name was ec2-user).

[AWS Docs: Connect to your Linux instance](#)

b. Install the AWS CLI.

[AWS Docs: Install or update to the latest version of the AWS CLI](#)

c. Use the `aws ec2 modify-instance-metadata-options` command to require the use of IMDSv2 and to change the PUT response hop limit to 3.

## Example

```
aws ec2 modify-instance-metadata-options \  
  --instance-id <instance-id> \  
  --http-put-response-hop-limit 3 \  
  --http-tokens required \  
  --http-endpoint enabled
```



The `http-tokens` parameter sets IMDSv2 to required. When `http-tokens` is required, you must also set `http-endpoint` to enabled.

2. Require the use of IMDSv2 on Cloud Volumes ONTAP instances:

a. Go to the [Amazon EC2 console](#)

b. From the navigation pane, select **Instances**.

c. Select a Cloud Volumes ONTAP instance.

d. Select **Actions** > **Instance settings** > **Modify instance metadata options**.

e. On the **Modify instance metadata options** dialog box, select the following:

- For **Instance metadata service**, select **Enable**.
- For **IMDSv2**, select **Required**.
- Select **Save**.

f. Repeat these steps for other Cloud Volumes ONTAP instances, including the HA mediator.

g. [Stop and start the Cloud Volumes ONTAP instances](#)

## Result

The Connector instance and Cloud Volumes ONTAP instances are now configured to use IMDSv2.

## Manage connector upgrades

When you use standard mode or restricted mode, BlueXP automatically upgrades your Connector to the latest release, as long as the Connector has outbound internet access to obtain the software update.

If you need to manually manage when the connector is upgraded, you can disable automatic upgrades for standard mode or restricted mode.



When running BlueXP in private mode, you must always upgrade the connector yourself.

### Disable automatic upgrades

Disabling auto-upgrade for your connector consists of two steps. First you need to ensure that your Connector is healthy and up-to-date. Then you'll edit a configuration file to turn off the automatic upgrade feature.



You can only disable automatic upgrades if you have connector version 3.9.48 or higher.

### Verify the health of your connector

You should verify that your connector is stable and all containers running on your connector VM are healthy and running. After you disable automatic upgrades, your connector VM stops checking for new services or upgrade packages.

Use one of the following commands to verify your connector. All services should have a status of *Running*. If this isn't the case, contact NetApp support before disabling auto-upgrade.

### Docker

```
docker ps -a
```

### Podman

```
podman ps -a
```

### Disable auto-upgrade for the connector

You disable automatic upgrades by setting the *isUpgradeDisabled* flag in the *com/opt/application/netapp/service-manager-2/config.json* file. By default, this flag is set to false and your connector is automatically upgraded. You can set this flag to true to disable automatic upgrades. You should be familiar with JSON syntax before completing this step.

To re-enable auto-upgrade, use these steps and set the *isUpgradeDisabled* flag to false.

### Steps

1. Ensure you have verified that your connector is up-to-date and healthy.
2. Create a backup copy of the */opt/application/netapp/service-manager-2/config.json* file to ensure you can

revert your changes.

3. Edit the `/opt/application/netapp/service-manager-2/config.json` file and change the value of the `isUpgradeDisabled` flag to true.

```
"isUpgradeDisabled": true,
```

4. Save your file.
5. Restart the service manager 2 service by running the following command:

```
systemctl restart netapp-service-manager.service
```

6. Run the following command and verify that the Connector status shows as *active(running)*:

—

```
systemctl status netapp-service-manager.service
```

## Upgrade the connector

The Connector needs to restart during the upgrade process so the web-based console will be unavailable during the upgrade.

### Steps

1. Download the Connector software from the [NetApp Support Site](#).

Be sure to download the offline installer for private networks without internet access.

2. Copy the installer to the Linux host.
3. Assign permissions to run the script.

```
chmod +x /path/BlueXP-Connector-Offline-<version>
```

Where `<version>` is the version of the Connector that you downloaded.

4. Run the installation script:

```
sudo /path/BlueXP-Connector-Offline-<version>
```

Where `<version>` is the version of the Connector that you downloaded.

5. After the upgrade is complete, you can verify the Connector's version by going to **Help > Support > Connector**.



## Work with multiple Connectors

If you use multiple Connectors, BlueXP enables you to switch between those Connectors directly from the console. You can also manage a single working environment with multiple Connectors.

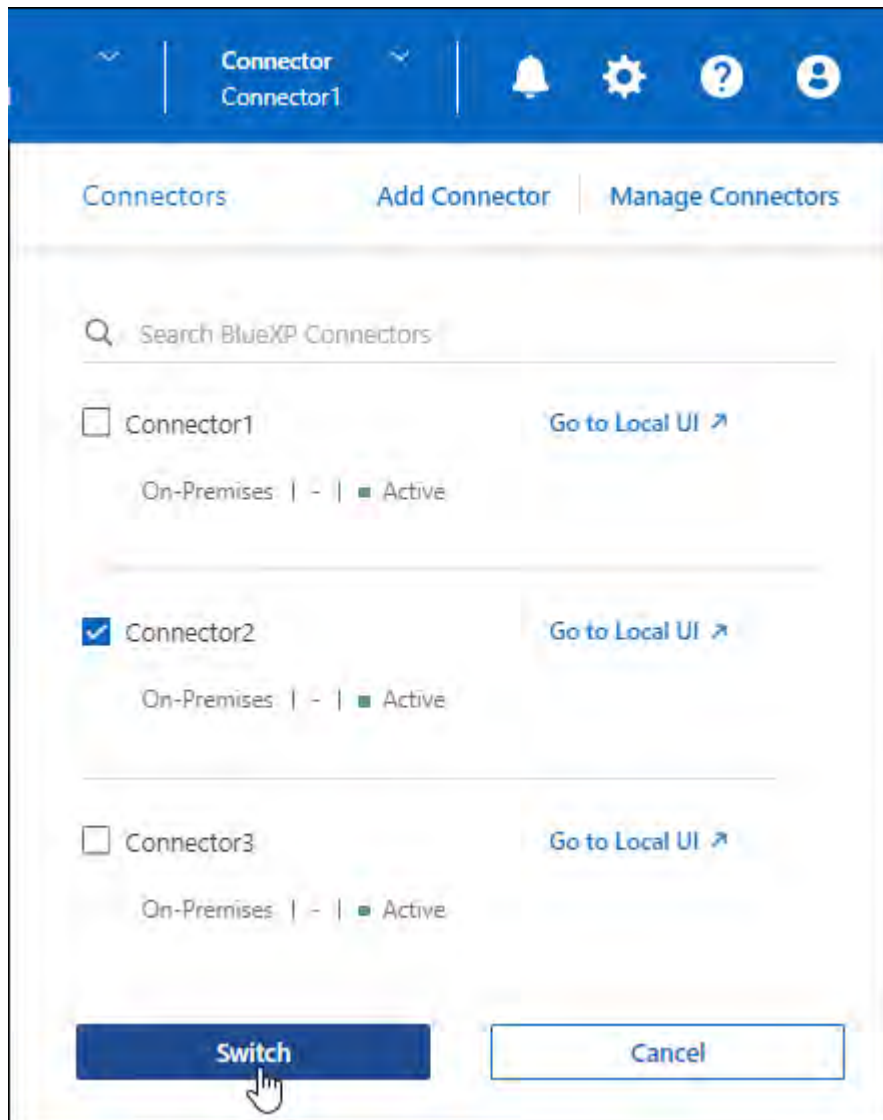
### Switch between Connectors

If you have multiple Connectors, you can switch between them to see the Working Environments that are associated with a specific Connector.

For example, let's say that you're working in a multi-cloud environment. You might have one Connector in AWS and another in Google Cloud. You'd need to switch between those Connectors to manage the Cloud Volumes ONTAP systems running in those clouds.

### Step

1. Select the **Connector** drop-down, select another Connector, and then select **Switch**.



### Result

BlueXP refreshes and shows the Working Environments associated with the selected Connector.

## Set up a disaster recovery configuration

You can manage a working environment with multiple Connectors at the same time for disaster recovery purposes. If one Connector goes down, you can switch to the other Connector to immediately manage the working environment.

### Steps

1. Switch to the other Connector that you want to manage with the working environment.
2. Discover the existing working environment.
  - [Add existing Cloud Volumes ONTAP systems to BlueXP](#)
  - [Discover ONTAP clusters](#)
3. If you're managing a Cloud Volumes ONTAP working environment, select **Settings > Connector Settings** and set the Capacity Management Mode to **Manual Mode**.

To avoid contention issues, only the main Connector should be set to **Automatic Mode**.

[Learn more about the capacity management mode](#)

## Troubleshoot the Connector

To troubleshoot issues with the Connector, you can work with NetApp Support who might ask for your system ID, Connector version, or the latest AutoSupport messages. You can also view the NetApp Knowledge Base to troubleshoot issues yourself.

### Related information

[Get help from NetApp Support.](#)

## Find the system ID for a Connector

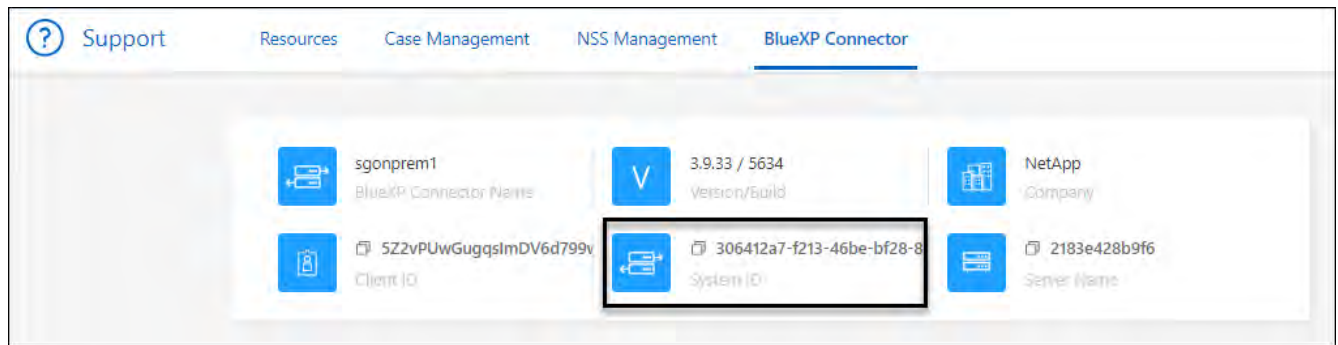
To help you get started, your NetApp representative might ask you for the system ID of your Connector. The ID is typically used for licensing and troubleshooting purposes.

### Steps

1. In the upper right of the BlueXP console, select the Help icon.
2. Select **Support > BlueXP Connector**.

The system ID appears at the top of the page.

### Example



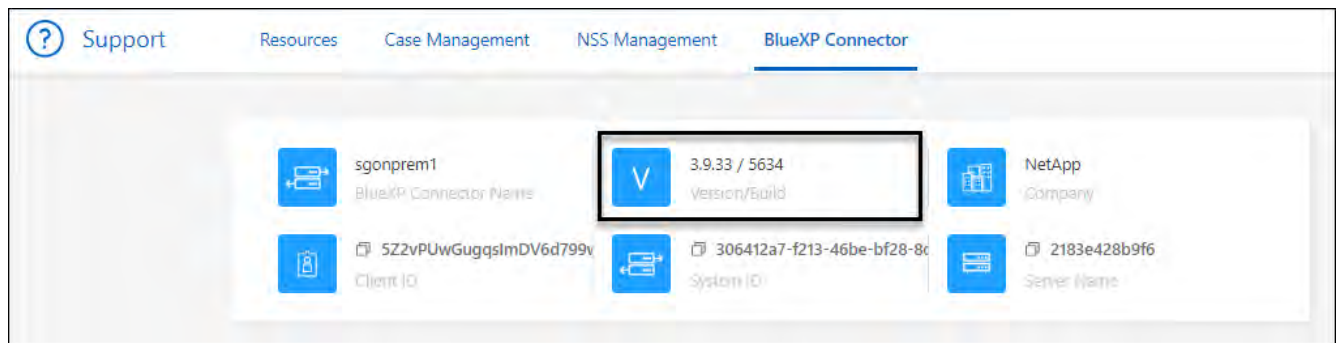
## View a Connector's version

You can view the version of your Connector to verify that the Connector automatically upgraded to the latest release or because you need to share it with your NetApp representative.

### Steps

1. In the upper right of the BlueXP console, select the Help icon.
2. Select **Support > BlueXP Connector**.

The version displays at the top of the page.

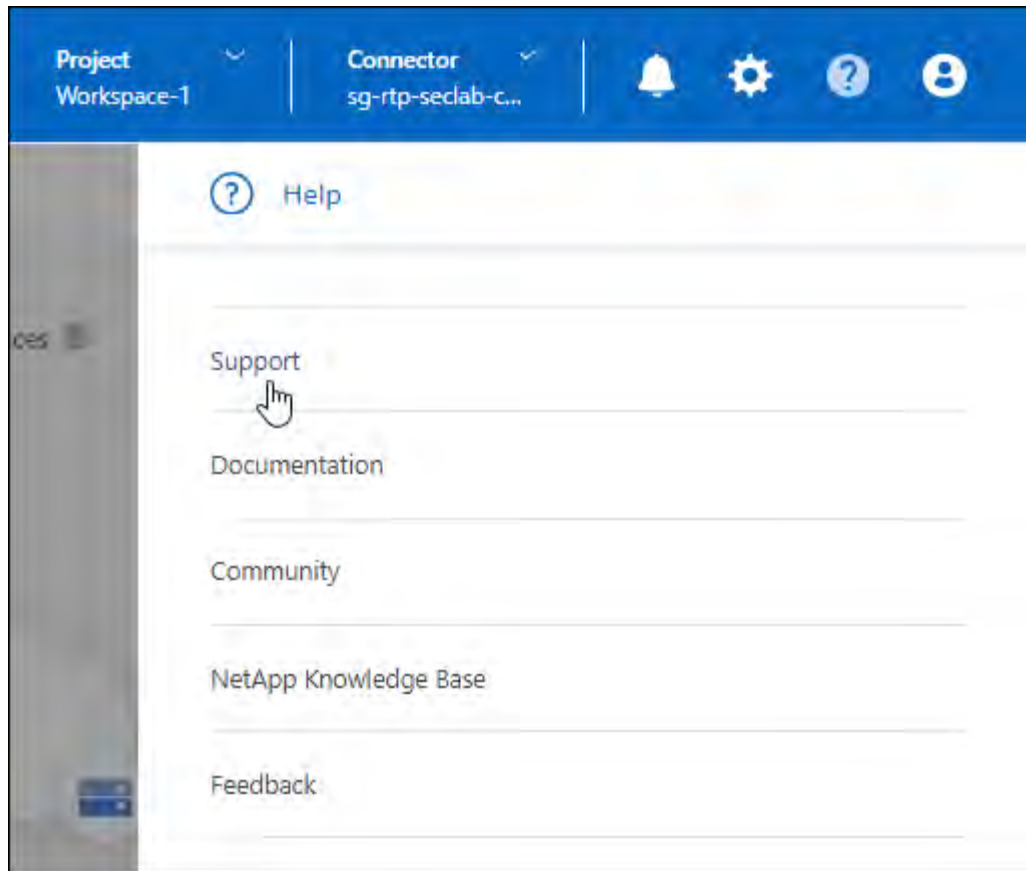


## Download or send an AutoSupport message


If you're having problems, NetApp personnel might ask you to send an AutoSupport message to NetApp support for troubleshooting purposes.

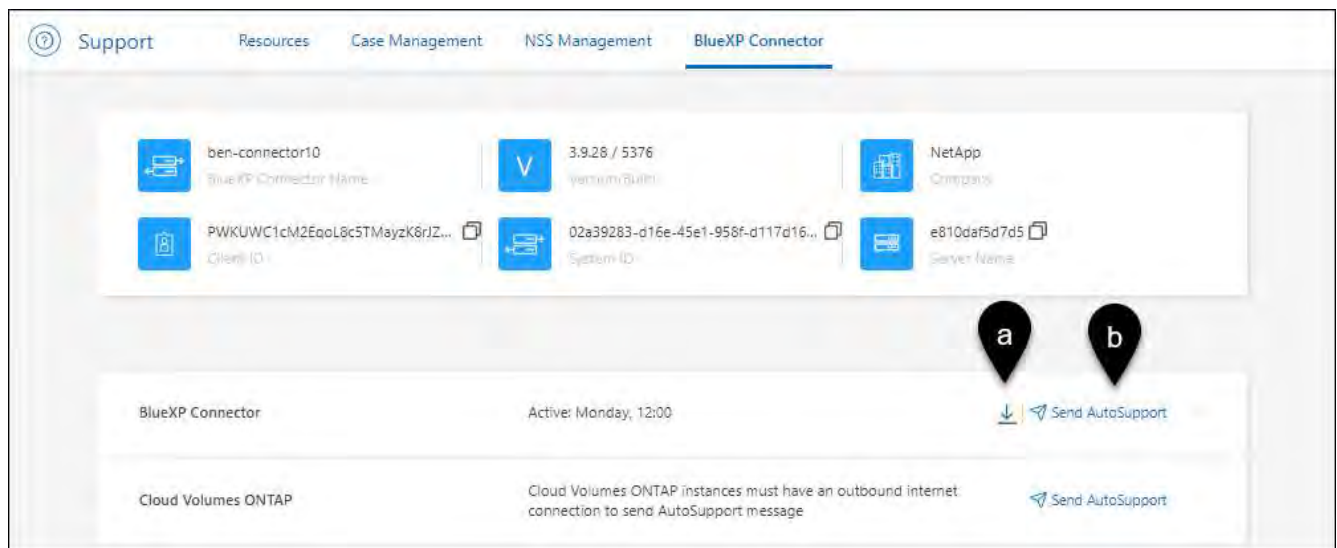
### Steps

1. In the upper right of the BlueXP console, select the Help icon, and select **Support**.



2. Select **BlueXP Connector**.
3. Depending on how you need to send the information to NetApp support, choose one of the following options:
  - a. Select the option to download the AutoSupport message to your local machine. You can then send it to NetApp Support using a preferred method.
  - b. Select **Send AutoSupport** to directly send the message to NetApp Support.

 BlueXP may take up to five hours to send AutoSupport messages due to load balancing. For urgent communication, download the file and send it manually.



## Fix download failures when using a Google Cloud NAT gateway

The Connector automatically downloads software updates for Cloud Volumes ONTAP. The download can fail if your configuration uses a Google Cloud NAT gateway. You can correct this issue by limiting the number of parts that the software image is divided into. This step must be completed by using the BlueXP API.

### Step

1. Submit a PUT request to `/occm/config` with the following JSON as body:

```
{
  "maxDownloadSessions": 32
}
```

The value for `maxDownloadSessions` can be 1 or any integer greater than 1. If the value is 1, then the downloaded image will not be divided.

Note that 32 is an example value. The value that you should use depends on your NAT configuration and the number of sessions that you can have simultaneously.

[Learn more about the /occm/config API call](#)

### Get help from the NetApp Knowledge Base

[View troubleshooting information created by the NetApp Support team.](#)

## Uninstall and remove the Connector

Uninstall the Connector software to troubleshoot issues or to permanently remove the software from the host. The steps that you need to use depends on the deployment mode that you're using. Once a Connector has been removed from your environment, you can remove it from BlueXP.

[Learn about BlueXP deployment modes.](#)

### Uninstall the Connector when using standard or restricted mode

If you're using standard mode or restricted mode (in other words, the Connector host has outbound connectivity), then you should follow the steps below to uninstall the Connector software.

#### Steps

1. Connect to the Linux VM for the Connector.
2. From the Linux host, run the uninstallation script:

```
/opt/application/netapp/service-manager-2/uninstall.sh [silent]
```

`silent` runs the script without prompting you for confirmation.

## Uninstall the Connector when using private mode

If you're using private mode (where the Connector host has *no* outbound connectivity), follow the steps below to uninstall the Connector software.

### Step

1. Connect to the Linux VM for the Connector.
2. From the Linux host, run the following commands:

```
/opt/application/netapp/ds/cleanup.sh  
rm -rf /opt/application/netapp/
```

3. From the Linux host, delete old, unused container image files to free space in the /var directory for re-installation.

#### Podman

```
podman system prune --all
```

#### Docker

```
docker system prune -a
```

## Remove Connectors from BlueXP

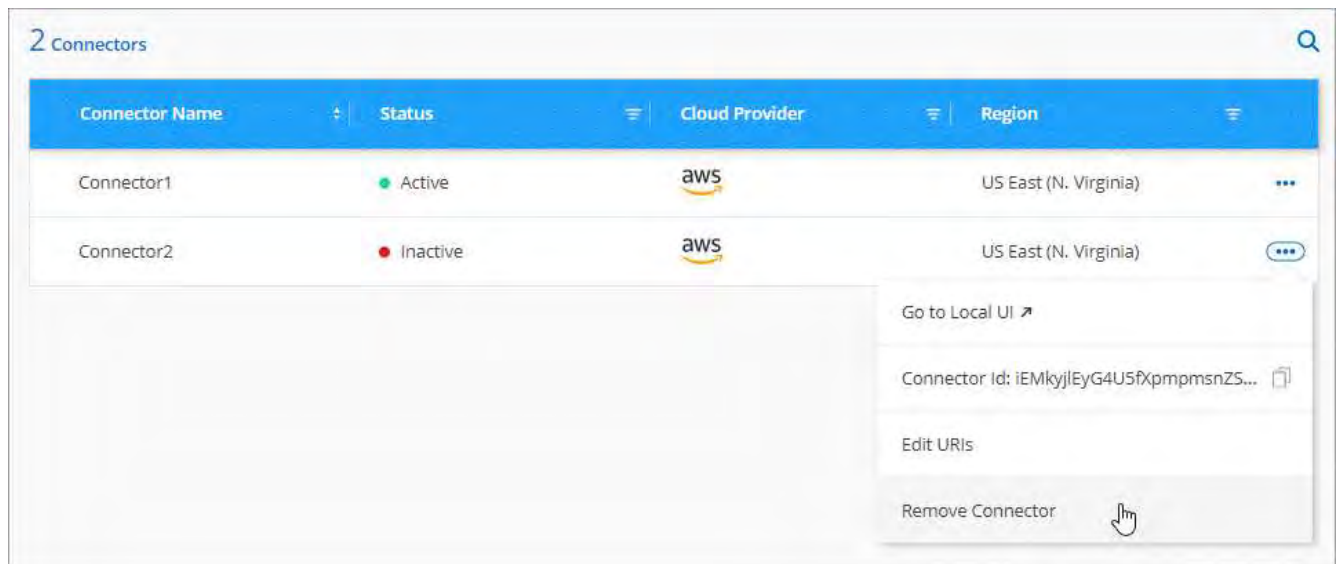
If a Connector is inactive, you can remove it from the list of Connectors in BlueXP. You might do this if you delete the Connector virtual machine or if you uninstall the Connector software.

Note the following about removing a Connector:

- This action doesn't delete the virtual machine.
- This action can't be reverted—once you remove a Connector, you can't add it back.

### Steps

1. Select the **Connector** drop-down from the BlueXP header.
2. Select **Manage Connectors**.
3. Select the action menu for an inactive Connector and select **Remove Connector**.



4. Enter the name of the Connector to confirm and then select **Remove**.

## Default configuration for the Connector

You might want to learn more about the Connector's configuration before you deploy it, or if you need to troubleshoot any issues.

### Default configuration with internet access

The following configuration details apply if you deployed the Connector from BlueXP, from your cloud provider's marketplace, or if you manually installed the Connector on an on-premises Linux host that has internet access.

#### AWS details

If you deployed the Connector from BlueXP or from the cloud provider's marketplace, note the following:

- The EC2 instance type is t3.2xlarge.
- The operating system for the image is Ubuntu 22.04 LTS.

The operating system does not include a GUI. You must use a terminal to access the system.

- The installation includes Docker Engine, which is the required container orchestration tool.
- The user name for the EC2 Linux instance is ubuntu (for Connectors created prior to May 2023, the user name was ec2-user).
- The default system disk is a 100 GiB gp2 disk.

#### Azure details

If you deployed the Connector from BlueXP or from the cloud provider's marketplace, note the following:

- The VM type is Standard\_D8s\_v3.
- The operating system for the image is Ubuntu 22.04 LTS.

The operating system does not include a GUI. You must use a terminal to access the system.

- The installation includes Docker Engine, which is the required container orchestration tool.
- The default system disk is a 100 GiB premium SSD disk.

### Google Cloud details

If you deployed the Connector from BlueXP, note the following:

- The VM instance is n2-standard-8.
- The operating system for the image is Ubuntu 22.04 LTS.

The operating system does not include a GUI. You must use a terminal to access the system.

- The installation includes Docker Engine, which is the required container orchestration tool.
- The default system disk is a 100 GiB SSD persistent disk.

### Installation folder

The Connector installation folder resides in the following location:

```
/opt/application/netapp/cloudmanager
```

### Log files

Log files are contained in the following folders:

- /opt/application/netapp/cloudmanager/log  
or
- /opt/application/netapp/service-manager-2/logs (starting with new 3.9.23 installations)

The logs in these folders provide details about the Connector.

- /opt/application/netapp/cloudmanager/docker\_occm/data/log

The logs in this folder provide details about cloud services and the BlueXP service that runs on the Connector.

### Connector service

- The BlueXP service is named occm.
- The occm service is dependent on the MySQL service.

If the MySQL service is down, then the occm service is down too.

### Ports

The Connector uses the following ports on the Linux host:

- 80 for HTTP access
- 443 for HTTPS access



## Default configuration without internet access

The following configuration applies if you manually installed the Connector on an on-premises Linux host that doesn't have internet access. [Learn more about this installation option.](#)

- The Connector installation folder resides in the following location:

`/opt/application/netapp/ds`

- Log files are contained in the following folders:

`/var/lib/docker/volumes/ds_occmdata/_data/log`

The logs in this folder provide details about the Connector and docker images.

- All services are running inside docker containers

The services are dependent on the docker runtime service running

- The Connector uses the following ports on the Linux host:

- 80 for HTTP access
- 443 for HTTPS access

## Enforce ONTAP permissions for ONTAP Advanced View (ONTAP System Manager)

By default, the Connector credentials allow users to access the Advanced View (ONTAP System Manager). You can prompt users for their ONTAP credentials instead. This ensures that a user's ONTAP permissions are applied when they work with ONTAP clusters in both Cloud Volumes ONTAP and ONTAP on-premises clusters.



You must have the Organization admin role to edit Connector settings.

### Steps

1. Select the **Connector** drop-down from the BlueXP header.
2. Select **Manage Connectors**.
3. Select the action menu in the row that corresponds to the Connector you want to edit.
4. Expand the **Force Credentials** option.
5. Select the checkbox to enable the **Force Credentials** option and then select **Save**.
6. Check if the **Force Credentials** option is enabled.



# Credentials and subscriptions

## AWS


### Learn about AWS credentials and permissions in BlueXP

Learn how BlueXP uses AWS credentials to perform actions on your behalf and how those credentials are associated with marketplace subscriptions. Understanding these details can be helpful as you manage the credentials for one or more AWS accounts in BlueXP. For example, you might want to learn about when to add additional AWS credentials to BlueXP.

#### Initial AWS credentials

When you deploy a Connector from BlueXP, you need to provide the ARN of an IAM role or access keys for an IAM user. The authentication method that you use must have the required permissions to deploy the Connector instance in AWS. The required permissions are listed in the [Connector deployment policy for AWS](#).

When BlueXP launches the Connector instance in AWS, it creates an IAM role and an instance profile for the instance. It also attaches a policy that provides the Connector with permissions to manage resources and processes within that AWS account. [Review how BlueXP uses the permissions](#).



If you create a new working environment for Cloud Volumes ONTAP, BlueXP selects these AWS credentials by default:

Details & Credentials			
Instance Profile Credentials	XXXXXXXXXXXX Account ID	QA Subscription Marketplace Subscription	<a href="#">Edit Credentials</a>

You can deploy all of your Cloud Volumes ONTAP systems using the initial AWS credentials, or you can add additional credentials.

## Additional AWS credentials

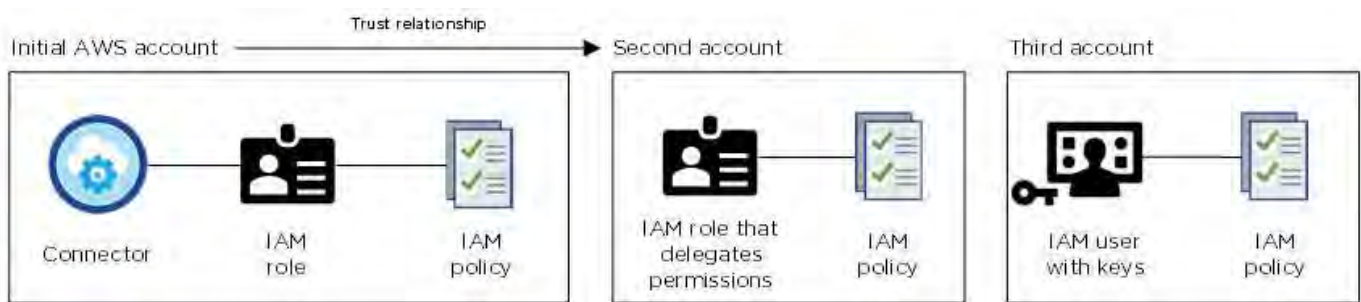
You might add additional AWS credentials to BlueXP in the following cases:

- To use your existing BlueXP Connector with an additional AWS account
- To create a new Connector in a specific AWS account
- To create and manage FSx for ONTAP file systems

Review the sections below for more details.


### Add AWS credentials to use a Connector with another AWS account

If you want to use BlueXP with additional AWS accounts, then you can either provide AWS keys for an IAM user or the ARN of a role in a trusted account. The following image shows two additional accounts, one providing permissions through an IAM role in a trusted account and another through the AWS keys of an IAM user:



You would then add the account credentials to BlueXP by specifying the Amazon Resource Name (ARN) of the IAM role, or the AWS keys for the IAM user.

For example, you can switch between credentials when creating a new Cloud Volumes ONTAP working environment:



[Learn how to add AWS credentials to an existing Connector.](#)

### **Add AWS credentials to create a Connector**

Adding new AWS credentials to BlueXP provides the permissions needed to create a Connector.

[Learn how to add AWS credentials to BlueXP for creating a Connector](#)

### **Add AWS credentials for FSx for ONTAP**

Adding new AWS credentials to BlueXP provides the permissions needed to create and manage an FSx for ONTAP working environment.

[Learn how to add AWS credentials to BlueXP for Amazon FSx for ONTAP](#)

### **Credentials and marketplace subscriptions**

The credentials that you add to a Connector must be associated with an AWS Marketplace subscription so that you can pay for Cloud Volumes ONTAP at an hourly rate (PAYGO) or through an annual contract, and to use other BlueXP services.

[Learn how to associate an AWS subscription.](#)

Note the following about AWS credentials and marketplace subscriptions:

- You can associate only one AWS Marketplace subscription with a set of AWS credentials
- You can replace an existing marketplace subscription with a new subscription

### **FAQ**

The following questions are related to credentials and subscriptions.

## How can I securely rotate my AWS credentials?

As described in the sections above, BlueXP enables you to provide AWS credentials in a few ways: an IAM role associated with the Connector instance, by assuming an IAM role in a trusted account, or by providing AWS access keys.

With the first two options, BlueXP uses the AWS Security Token Service to obtain temporary credentials that rotate constantly. This process is the best practice—it's automatic and it's secure.

If you provide BlueXP with AWS access keys, you should rotate the keys by updating them in BlueXP at a regular interval. This is a completely manual process.

## Can I change the AWS Marketplace subscription for Cloud Volumes ONTAP working environments?

Yes, you can. When you change the AWS Marketplace subscription that's associated with a set of credentials, all existing and new Cloud Volumes ONTAP working environments will be charged against the new subscription.

[Learn how to associate an AWS subscription.](#)

## Can I add multiple AWS credentials, each with different marketplace subscriptions?

All AWS credentials that belong to the same AWS account will be associated with the same AWS Marketplace subscription.

If you have multiple AWS credentials that belong to different AWS accounts, then those credentials can be associated with the same AWS Marketplace subscription or with different subscriptions.

## Can I move existing Cloud Volumes ONTAP working environments to a different AWS account?

No, it's not possible to move the AWS resources associated with your Cloud Volumes ONTAP working environment to a different AWS account.

## How do credentials work for marketplace deployments and on-premises deployments?

The sections above describe the recommended deployment method for the Connector, which is from BlueXP. You can also deploy a Connector in AWS from the AWS Marketplace and you can manually install the Connector software on your own Linux host.

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

For on-premises deployments, you can't set up an IAM role for the BlueXP system, but you can provide permissions using AWS access keys.

To learn how to set up permissions, refer to the following pages:

- Standard mode
  - [Set up permissions for an AWS Marketplace deployment](#)
  - [Set up permissions for on-premises deployments](#)
- [Set up permissions for restricted mode](#)
- [Set up permissions for private mode](#)

## Manage AWS credentials and marketplace subscriptions for BlueXP

Add and manage AWS credentials so that you deploy and manage cloud resources in your AWS accounts from BlueXP. If you manage multiple AWS Marketplace subscriptions, you can assign each one of them to different AWS credentials from the Credentials page.

### Overview

You can add AWS credentials to an existing Connector or directly to BlueXP:

- Add additional AWS credentials to an existing Connector

Add AWS credentials to a Connector to manage resources in your cloud environment. [Learn how to add AWS credentials to a Connector.](#)

- Add AWS credentials to BlueXP for creating a Connector

Adding new AWS credentials to BlueXP gives BlueXP the permissions needed to create a Connector. [Learn how to add AWS credentials to BlueXP.](#)

- Add AWS credentials to BlueXP for FSx for ONTAP

Add new AWS credentials to BlueXP to create and manage FSx for ONTAP. [Learn how to set up permissions for FSx for ONTAP](#)

### How to rotate credentials

BlueXP enables you to provide AWS credentials in a few ways: an IAM role associated with the Connector instance, by assuming an IAM role in a trusted account, or by providing AWS access keys. [Learn more about AWS credentials and permissions.](#)

With the first two options, BlueXP uses the AWS Security Token Service to obtain temporary credentials that rotate constantly. This process is the best practice because it's automatic and it's secure.

Rotate AWS access keys regularly by updating them in BlueXP. This process is manual.

### Add additional credentials to a Connector

Add additional AWS credentials to a Connector so that it has the permissions needed to manage resources and processes within your public cloud environment. You can either provide the ARN of an IAM role in another account or provide AWS access keys.

If you're just getting started with BlueXP, [Learn how BlueXP uses AWS credentials and permissions.](#)

### Grant permissions

Provide required permissions before adding AWS credentials to a Connector. The permissions allow the Connector to manage resources and processes within that AWS account. You can provide the permissions with the the ARN of a role in a trusted account or AWS keys.



If you deployed a Connector from BlueXP, BlueXP automatically added AWS credentials for the account in which you deployed the Connector. This ensures the necessary permissions are in place for managing resources. [Learn about AWS credentials and permissions.](#)

## Choices

- [Grant permissions by assuming an IAM role in another account](#)
- [Grant permissions by providing AWS keys](#)

### Grant permissions by assuming an IAM role in another account

You can set up a trust relationship between the source AWS account in which you deployed the Connector instance and other AWS accounts by using IAM roles. You would then provide BlueXP with the ARN of the IAM roles from the trusted accounts.

If the Connector is installed on-premises, you can't use this authentication method. You must use AWS keys.

#### Steps

1. Go to the IAM console in the target account in which you want to provide the Connector with permissions.
2. Under Access Management, select **Roles > Create Role** and follow the steps to create the role.

Be sure to do the following:

- Under **Trusted entity type**, select **AWS account**.
  - Select **Another AWS account** and enter the ID of the account where the Connector instance resides.
  - Create the required policies by copying and pasting the contents of [the IAM policies for the Connector](#).
3. Copy the Role ARN of the IAM role so that you can paste it in BlueXP later on.

#### Result

The account now has the required permissions. [You can now add the credentials to a Connector](#).

### Grant permissions by providing AWS keys

If you want to provide BlueXP with AWS keys for an IAM user, then you need to grant the required permissions to that user. The BlueXP IAM policy defines the AWS actions and resources that BlueXP is allowed to use.

You must use this authentication method if the Connector is installed on-premises. You can't use an IAM role.

#### Steps

1. From the IAM console, create policies by copying and pasting the contents of [the IAM policies for the Connector](#).

[AWS Documentation: Creating IAM Policies](#)

2. Attach the policies to an IAM role or an IAM user.
  - [AWS Documentation: Creating IAM Roles](#)
  - [AWS Documentation: Adding and Removing IAM Policies](#)

#### Result

The account now has the required permissions. [You can now add the credentials to a Connector](#).

### Add the credentials

After you provide an AWS account with the required permissions, you can add the credentials for that account to an existing Connector. This enables you to launch Cloud Volumes ONTAP systems in that account using the

same Connector.

### Before you begin

If you just created these credentials in your cloud provider, it might take a few minutes until they are available for use. Wait a few minutes and then add you add the credentials.

### Steps

1. Use the top navigation bar to elect the Connector to which you want to add credentials.
2. In the upper right of the console, select the Settings icon, and select **Credentials**.



3. On the **Organization credentials** or **Account credentials** page, select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location:** Select **Amazon Web Services > Connector**.
  - b. **Define Credentials:** Provide the ARN (Amazon Resource Name) of a trusted IAM role, or enter an AWS access key and secret key.
  - c. **Marketplace Subscription:** Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.

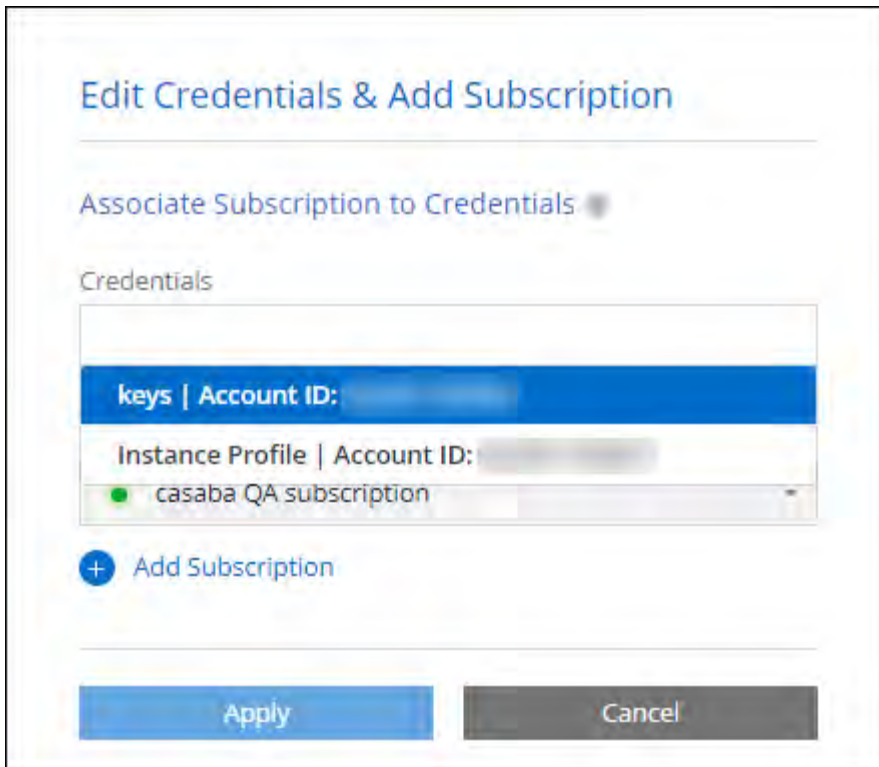
To pay for services at an hourly rate (PAYGO) or with an annual contract, you must associate AWS credentials with your AWS Marketplace subscription.

- d. **Review:** Confirm the details about the new credentials and select **Add**.

### Result

You can now switch to a different set of credentials from the Details and Credentials page when creating a new working environment:





#### Add credentials to BlueXP for creating a Connector

Add AWS credentials by providing the ARN of an IAM role that gives the permissions needed to create a Connector. You can choose these credentials when creating a new Connector.

#### Set up the IAM role

Set up an IAM role that enables the BlueXP software as a service (SaaS) layer to assume the role.

#### Steps

1. Go to the IAM console in the target account.
2. Under Access Management, select **Roles > Create Role** and follow the steps to create the role.

Be sure to do the following:

- Under **Trusted entity type**, select **AWS account**.
- Select **Another AWS account** and enter the ID of the BlueXP SaaS: 952013314444
- For Amazon FSx for NetApp ONTAP specifically, edit the **Trust relationships** policy to include "AWS": "arn:aws:iam::952013314444:root".

For example, the policy should look like this:

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::952013314444:root",
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}

```

Refer to [AWS Identity and Access Management \(IAM\) documentation](#) for more information on cross account resource access in IAM.

- Create a policy that includes the permissions required to create a Connector.
  - [View the permissions needed for FSx for ONTAP](#)
  - [View the Connector deployment policy](#)

3. Copy the Role ARN of the IAM role so that you can paste it in BlueXP in the next step.

## Result

The IAM role now has the required permissions. [You can now add it to BlueXP.](#)

## Add the credentials

After you provide the IAM role with the required permissions, add the role ARN to BlueXP.

## Before you begin

If you just created the IAM role, it might take a few minutes until they are available for use. Wait a few minutes before you add the credentials to BlueXP.

## Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.



2. On the **Organization credentials** or **Account credentials** page, select **Add Credentials** and follow the steps in the wizard.

- a. **Credentials Location:** Select **Amazon Web Services > BlueXP**.
- b. **Define Credentials:** Provide the ARN (Amazon Resource Name) of the IAM role.
- c. **Review:** Confirm the details about the new credentials and select **Add**.

## Add credentials to BlueXP for Amazon FSx for ONTAP

For details, refer to the [BlueXP documentation for Amazon FSx for ONTAP](#)

### Configure an AWS subscription

After you add your AWS credentials, you can configure an AWS Marketplace subscription with those credentials. The subscription enables you to pay for Cloud Volumes ONTAP at an hourly rate (PAYGO) or using an annual contract, and to pay for other data services.

There are two scenarios in which you might configure an AWS Marketplace subscription after you've already added the credentials:

- You didn't configure a subscription when you initially added the credentials.
- You want to change the AWS Marketplace subscription that is configured to the AWS credentials.

Replacing the current marketplace subscription with a new subscription changes the marketplace subscription for any existing Cloud Volumes ONTAP working environments and all new working environments.

### Before you begin

You need to create a Connector before you can configure a subscription. [Learn how to create a Connector](#).

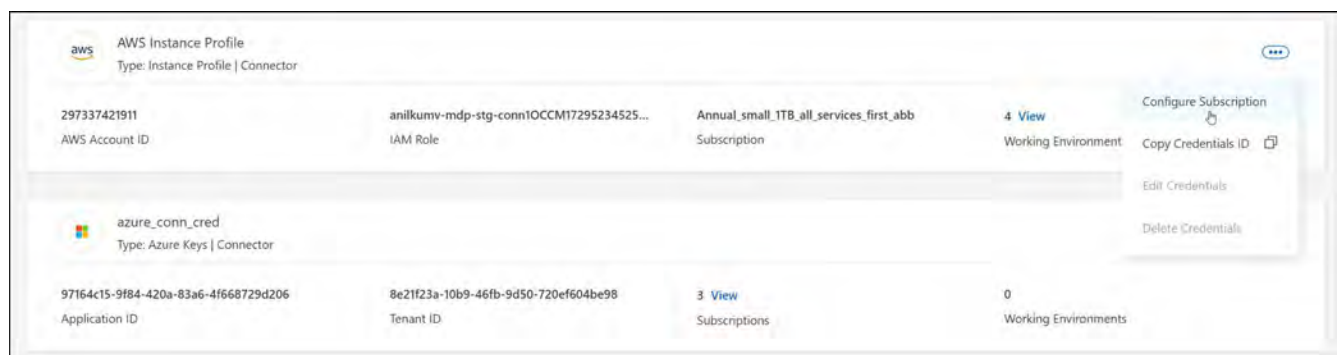
The following video shows the steps to subscribe to NetApp Intelligent Services from the AWS Marketplace:

[Subscribe to NetApp Intelligent Services from the AWS Marketplace](#)

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.
2. Select the action menu for a set of credentials and then select **Configure Subscription**.

You must select credentials that are associated with a Connector. You can't associate a marketplace subscription with credentials that are associated with BlueXP.



3. To associate the credentials with an existing subscription, select the subscription from the down-down list and select **Configure**.
4. To associate the credentials with a new subscription, select **Add Subscription > Continue** and follow the steps in the AWS Marketplace:
  - a. Select **View purchase options**.
  - b. Select **Subscribe**.

c. Select **Set up your account**.

You'll be redirected to the BlueXP website.

d. From the **Subscription Assignment** page:

- Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
- In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

- Select **Save**.

#### **Associate an existing subscription with your organization or account**

When you subscribe to from the AWS Marketplace, the last step in the process is to associate the subscription with your organization. If you didn't complete this step, then you can't use the subscription with your organization or account.

- [Learn about BlueXP deployment modes](#)
- [Learn about BlueXP identity and access management](#)

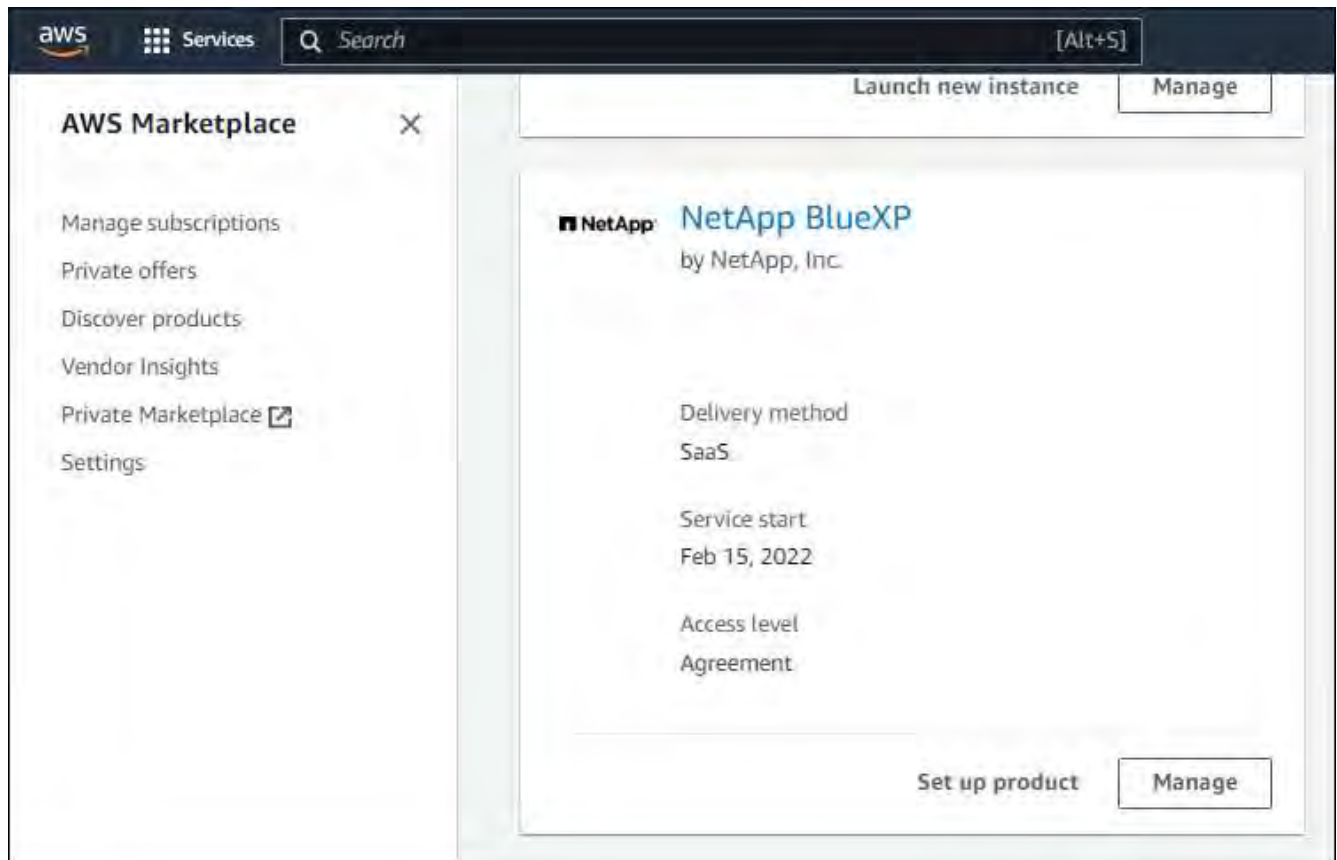
Follow the steps below if you subscribed to NetApp intelligent data services from the AWS Marketplace, but you missed the step to associate the subscription with your account.

#### **Steps**

1. Go to the digital wallet to confirm that you didn't associate your subscription with your BlueXP organization or account.
  - a. From the navigation menu, select **Governance > Digital wallet**.
  - b. Select **Subscriptions**.
  - c. Verify that your subscription doesn't appear.

You'll only see the subscriptions that are associated with the organization or account that you're currently viewing. If you don't see your subscription, proceed with the following steps.

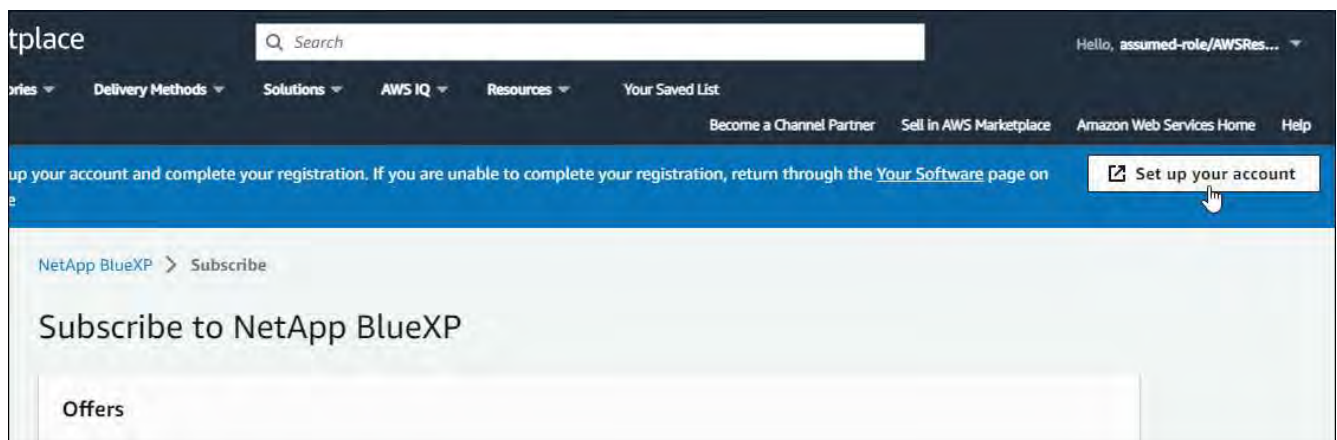
2. Log in to the AWS Console and navigate to **AWS Marketplace Subscriptions**.
3. Find the NetApp Intelligent Data Services subscription.



4. Select **Set up product**.

The subscription offer page should load in a new browser tab or window.

5. Select **Set up your account**.



The **Subscription Assignment** page on netapp.com should load in a new browser tab or window.

Note that you might be prompted to log in to BlueXP first.

6. From the **Subscription Assignment** page:

- Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
- In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

**Subscription Assignment**

✓ Your subscription to BlueXP / Cloud Volumes ONTAP from the AWS Marketplace was created successfully.

Subscription name ?

PayAsYouGo

Select the NetApp accounts that you'd like to associate this subscription with. ?

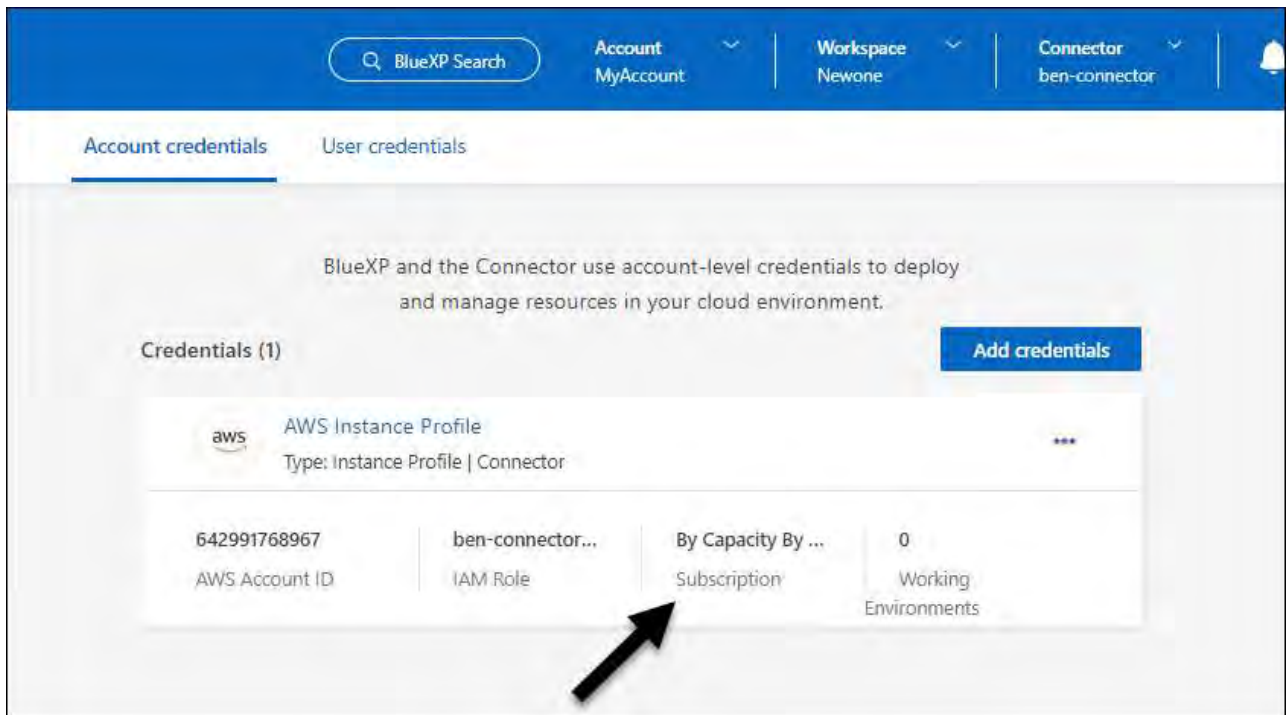
You can automatically replace the existing subscription for one account with this new subscription.

NetApp account	Replace existing subscription
<input checked="" type="checkbox"/> cloudTiering_undefined	<input type="checkbox"/>
<input checked="" type="checkbox"/> CS-HhewH	<input type="checkbox"/>
<input checked="" type="checkbox"/> benAccount	<input checked="" type="checkbox"/>

Save

7. Go to the digital wallet to confirm that the subscription is associated with your organization or account.
  - a. From the navigation menu, select **Governance > Digital wallet**.
  - b. Select **Subscriptions**.
  - c. Verify that your subscription appears.
8. Confirm that the subscription is associated with your AWS credentials.
  - a. In the upper right of the console, select the Settings icon, and select **Credentials**.
  - b. On the **Organization credentials** or **Account credentials** page, verify that the subscription is associated with your AWS credentials.

Here's an example.



### Edit credentials

Edit your AWS credentials by changing the account type (AWS keys or assume role), by editing the name, or by updating the credentials themselves (the keys or the role ARN).



You can't edit the credentials for an instance profile that is associated with a Connector instance or an Amazon FSx for ONTAP instance. You can only rename the credentials for an FSx for ONTAP instance.

### Steps

1. In the upper right of the console, select the Settings icon, and select **Credentials**.
2. On the **Organization credentials** or **Account credentials** page, select the action menu for a set of credentials and then select **Edit Credentials**.
3. Make the required changes and then select **Apply**.

### Delete credentials

If you no longer need a set of credentials, you can delete them. You can only delete credentials that aren't associated with a working environment.



You can't delete the credentials for an instance profile that is associated with a Connector instance.

### Steps

1. In the upper right of the console, select the Settings icon, and select **Credentials**.
2. On the **Organization credentials** or **Account credentials** page, select the action menu for a set of credentials and then select **Delete Credentials**.



3. Select **Delete** to confirm.

## Azure


### Learn about Azure credentials and permissions in BlueXP

Learn how BlueXP uses Azure credentials to perform actions on your behalf and how those credentials are associated with marketplace subscriptions. Understanding these details can be helpful as you manage the credentials for one or more Azure subscriptions. For example, you might want to learn when to add additional Azure credentials to BlueXP.

#### Initial Azure credentials

When you deploy a Connector from BlueXP, you need to use an Azure account or service principal that has permissions to deploy the Connector virtual machine. The required permissions are listed in the [Connector deployment policy for Azure](#).

When BlueXP deploys the Connector virtual machine in Azure, it enables a [system-assigned managed identity](#) on the virtual machine, creates a custom role, and assigns it to the virtual machine. The role provides BlueXP with the permissions required to manage resources and processes within that Azure subscription. [Review how BlueXP uses the permissions](#).



If you create a new working environment for Cloud Volumes ONTAP, BlueXP selects these Azure credentials by default:

Details & Credentials			
Managed Service Ide...	OCCM QA1	<span style="color: orange;">!</span> <i>No subscription is associated</i>	<a href="#">Edit Credentials</a>
Credential Name	Azure Subscription	Marketplace Subscription	

You can deploy all of your Cloud Volumes ONTAP systems using the initial Azure credentials, or you can add additional credentials.

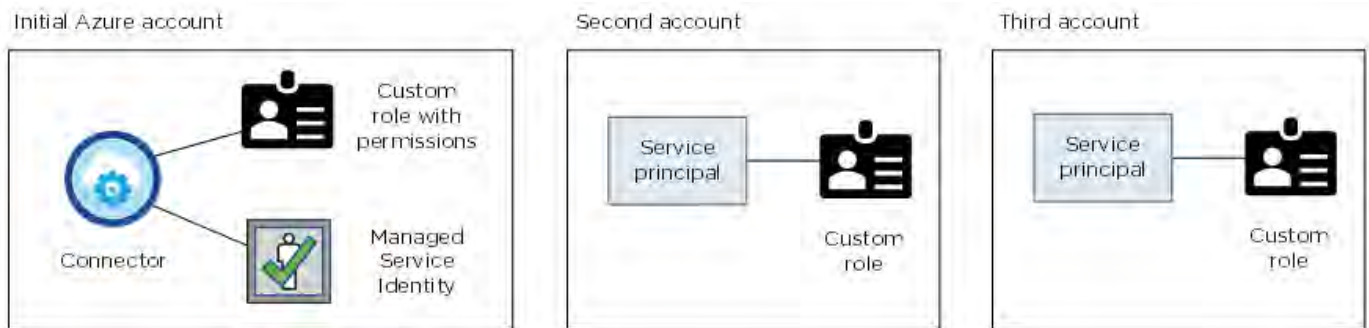


### Additional Azure subscriptions for a managed identity

The system-assigned managed identity assigned to the Connector VM is associated with the subscription in which you launched the Connector. If you want to select a different Azure subscription, then you need to [associate the managed identity with those subscriptions](#).


### Additional Azure credentials

If you want to use different Azure credentials with BlueXP, then you must grant the required permissions by [creating and setting up a service principal in Microsoft Entra ID](#) for each Azure account. The following image shows two additional accounts, each set up with a service principal and custom role that provides permissions:



You would then [add the account credentials to BlueXP](#) by providing details about the AD service principal.

For example, you can switch between credentials when creating a new Cloud Volumes ONTAP working environment:



### Credentials and marketplace subscriptions

The credentials that you add to a Connector must be associated with an Azure Marketplace subscription so that you can pay for Cloud Volumes ONTAP at an hourly rate (PAYGO) or through an annual contract, and to use other BlueXP services.

[Learn how to associate an Azure subscription.](#)

Note the following about Azure credentials and marketplace subscriptions:

- You can associate only one Azure Marketplace subscription with a set of Azure credentials
- You can replace an existing marketplace subscription with a new subscription

## FAQ

The following question is related to credentials and subscriptions.

### **Can I change the Azure Marketplace subscription for Cloud Volumes ONTAP working environments?**

Yes, you can. When you change the Azure Marketplace subscription that's associated with a set of Azure credentials, all existing and new Cloud Volumes ONTAP working environments will be charged against the new subscription.

[Learn how to associate an Azure subscription.](#)

### **Can I add multiple Azure credentials, each with different marketplace subscriptions?**

All Azure credentials that belong to the same Azure subscription will be associated with the same Azure Marketplace subscription.

If you have multiple Azure credentials that belong to different Azure subscriptions, then those credentials can be associated with the same Azure Marketplace subscription or with different marketplace subscriptions.

### **Can I move existing Cloud Volumes ONTAP working environments to a different Azure subscription?**

No, it's not possible to move the Azure resources associated with your Cloud Volumes ONTAP working environment to a different Azure subscription.

### **How do credentials work for marketplace deployments and on-premises deployments?**

The sections above describe the recommended deployment method for the Connector, which is from BlueXP. You can also deploy a Connector in Azure from the Azure Marketplace, and you can install the Connector software on your own Linux host.

If you use the Marketplace, you can provide permissions by assigning a custom role to the Connector VM and to a system-assigned managed identity, or you can use a Microsoft Entra service principal.

For on-premises deployments, you can't set up a managed identity for the Connector, but you can provide permissions by using a service principal.

To learn how to set up permissions, refer to the following pages:

- Standard mode
  - [Set up permissions for an Azure Marketplace deployment](#)
  - [Set up permissions for on-premises deployments](#)
- [Set up permissions for restricted mode](#)
- [Set up permissions for private mode](#)

### **Manage Azure credentials and marketplace subscriptions for BlueXP**

Add and manage Azure credentials so that BlueXP has the permissions that it needs to deploy and manage cloud resources in your Azure subscriptions. If you manage multiple

Azure Marketplace subscriptions, you can assign each one of them to different Azure credentials from the Credentials page.

Follow the steps on this page if you need to use multiple Azure credentials or multiple Azure Marketplace subscriptions for Cloud Volumes ONTAP.

## Overview

There are two ways to add additional Azure subscriptions and credentials in BlueXP.

1. Associate additional Azure subscriptions with the Azure managed identity.
2. If you want to deploy Cloud Volumes ONTAP using different Azure credentials, grant Azure permissions using a service principal and add its credentials to BlueXP.

## Associate additional Azure subscriptions with a managed identity

BlueXP enables you to choose the Azure credentials and Azure subscription in which you want to deploy Cloud Volumes ONTAP. You can't select a different Azure subscription for the managed identity profile unless you associate the [managed identity](#) with those subscriptions.

## About this task

A managed identity is [the initial Azure account](#) when you deploy a Connector from BlueXP. When you deployed the Connector, BlueXP created the BlueXP Operator role and assigned it to the Connector virtual machine.

## Steps

1. Log in to the Azure portal.
2. Open the **Subscriptions** service and then select the subscription in which you want to deploy Cloud Volumes ONTAP.
3. Select **Access control (IAM)**.
  - a. Select **Add > Add role assignment** and then add the permissions:
    - Select the **BlueXP Operator** role.

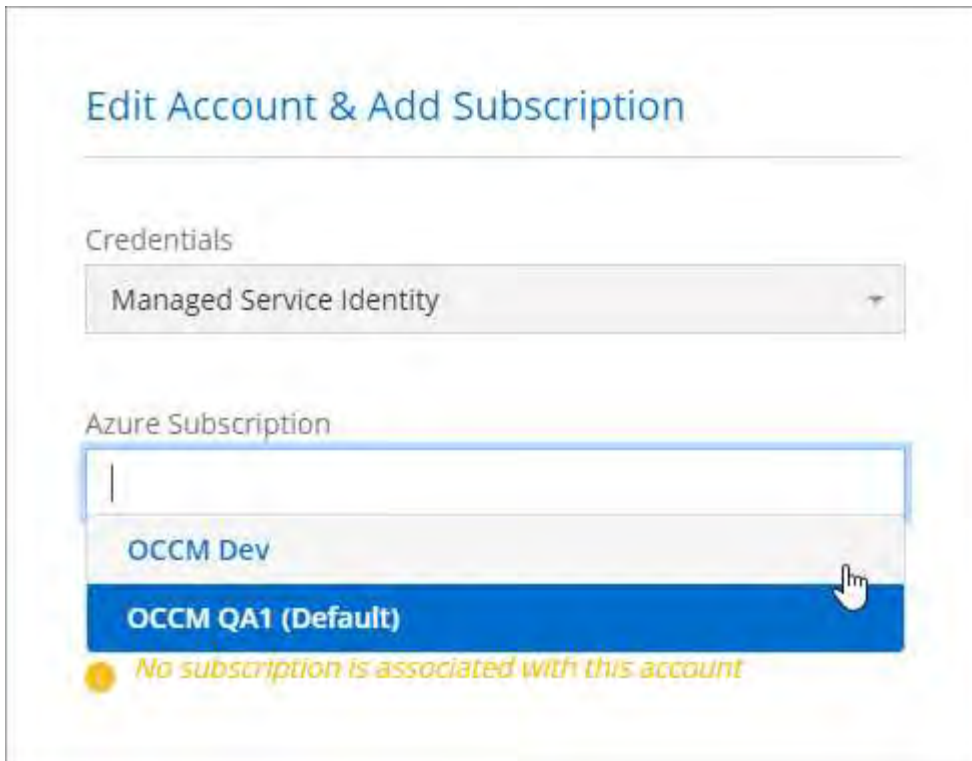


BlueXP Operator is the default name provided in the Connector policy. If you chose a different name for the role, then select that name instead.

- Assign access to a **Virtual Machine**.
  - Select the subscription in which the Connector virtual machine was created.
  - Select the Connector virtual machine.
  - Select **Save**.
4. Repeat these steps for additional subscriptions.

## Result

When you create a new working environment, you should now have the ability to select from multiple Azure subscriptions for the managed identity profile.



### Add additional Azure credentials to BlueXP

When you deploy a Connector from BlueXP, BlueXP enables a system-assigned managed identity on the virtual machine that has the required permissions. BlueXP selects these Azure credentials by default when you create a new working environment for Cloud Volumes ONTAP.



An initial set of credentials isn't added if you manually installed the Connector software on an existing system. [Learn about Azure credentials and permissions.](#)

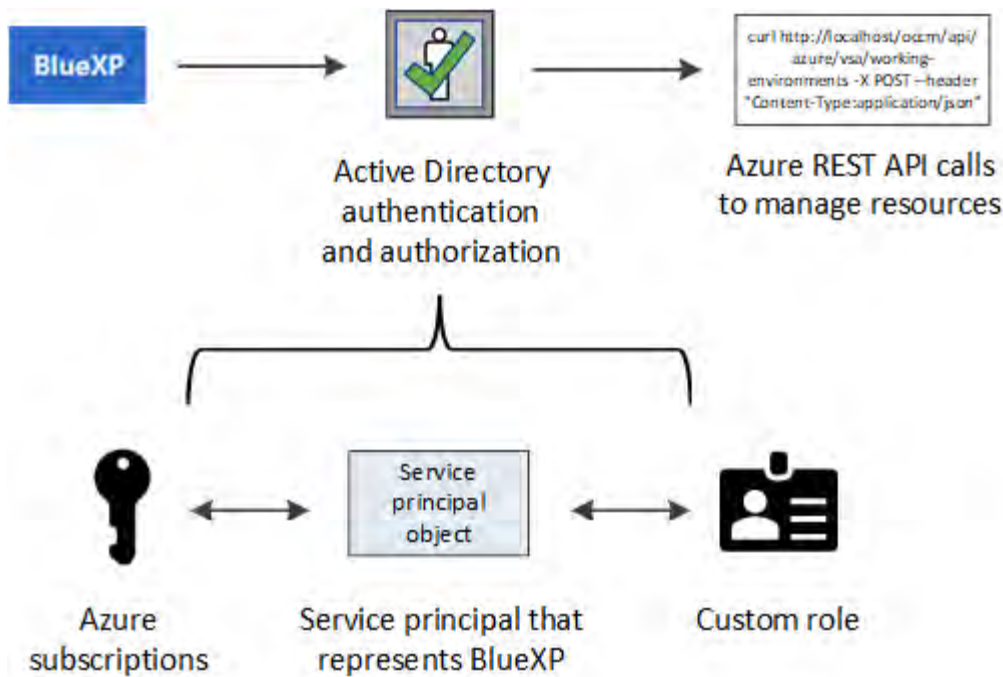
If you want to deploy Cloud Volumes ONTAP using *different* Azure credentials, then you must grant the required permissions by creating and setting up a service principal in Microsoft Entra ID for each Azure account. You can then add the new credentials to BlueXP.

### Grant Azure permissions using a service principal

BlueXP needs permissions to perform actions in Azure. You can grant the required permissions to an Azure account by creating and setting up a service principal in Microsoft Entra ID and by obtaining the Azure credentials that BlueXP needs.

#### About this task

The following image depicts how BlueXP obtains permissions to perform operations in Azure. A service principal object, which is tied to one or more Azure subscriptions, represents BlueXP in Microsoft Entra ID and is assigned to a custom role that allows the required permissions.



### Steps

1. Create a Microsoft Entra application.
2. Assign the application to a role.
3. Add Windows Azure Service Management API permissions.
4. Get the application ID and directory ID.
5. Create a client secret.

### Create a Microsoft Entra application

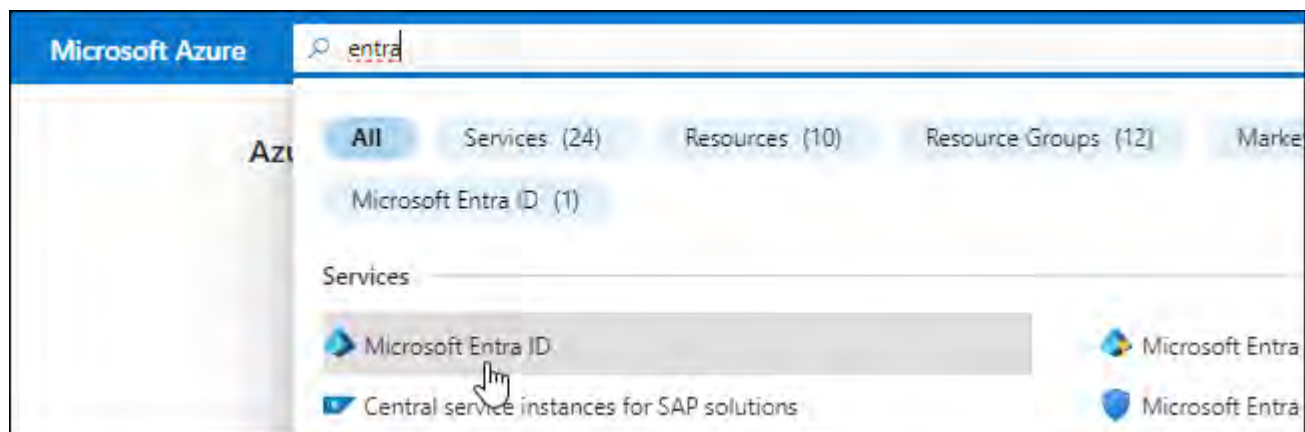
Create a Microsoft Entra application and service principal that BlueXP can use for role-based access control.

### Steps

1. Ensure that you have permissions in Azure to create an Active Directory application and to assign the application to a role.

For details, refer to [Microsoft Azure Documentation: Required permissions](#)

2. From the Azure portal, open the **Microsoft Entra ID** service.



3. In the menu, select **App registrations**.
4. Select **New registration**.
5. Specify details about the application:
  - **Name**: Enter a name for the application.
  - **Account type**: Select an account type (any will work with BlueXP).
  - **Redirect URI**: You can leave this field blank.
6. Select **Register**.

You've created the AD application and service principal.

## Result

You've created the AD application and service principal.

## Assign the application to a role

You must bind the service principal to one or more Azure subscriptions and assign it the custom "BlueXP Operator" role so BlueXP has permissions in Azure.

## Steps

1. Create a custom role:

Note that you can create an Azure custom role using the Azure portal, Azure PowerShell, Azure CLI, or REST API. The following steps show how to create the role using the Azure CLI. If you would prefer to use a different method, refer to [Azure documentation](#)

- a. Copy the contents of the [custom role permissions for the Connector](#) and save them in a JSON file.
- b. Modify the JSON file by adding Azure subscription IDs to the assignable scope.

You should add the ID for each Azure subscription from which users will create Cloud Volumes ONTAP systems.


## Example

```
"AssignableScopes": [  
  "/subscriptions/d333af45-0d07-4154-943d-c25fbzzzzzzz",  
  "/subscriptions/54b91999-b3e6-4599-908e-416e0zzzzzzz",  
  "/subscriptions/398e471c-3b42-4ae7-9b59-ce5bbzzzzzzz"]
```

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

The following steps describe how to create the role by using Bash in Azure Cloud Shell.

- Start [Azure Cloud Shell](#) and choose the Bash environment.
- Upload the JSON file.



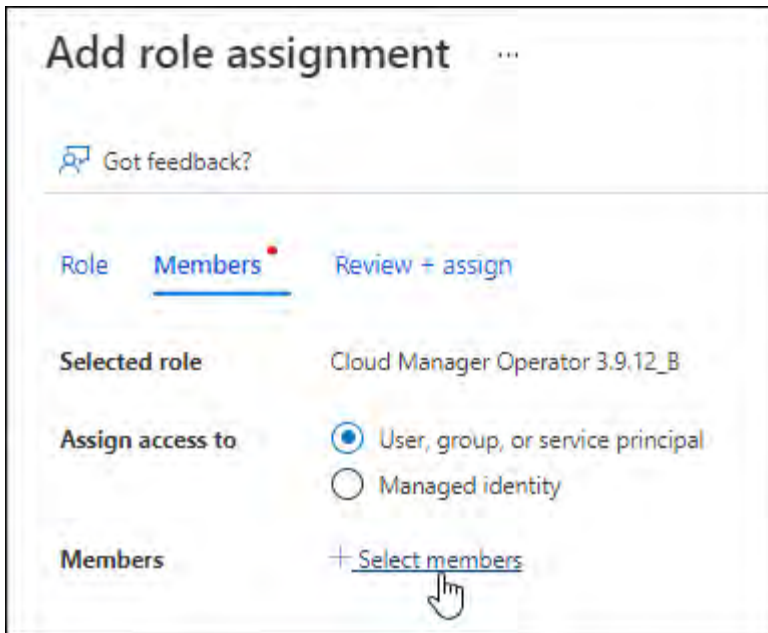
- Use the Azure CLI to create the custom role:

```
az role definition create --role-definition Connector_Policy.json
```

You should now have a custom role called BlueXP Operator that you can assign to the Connector virtual machine.

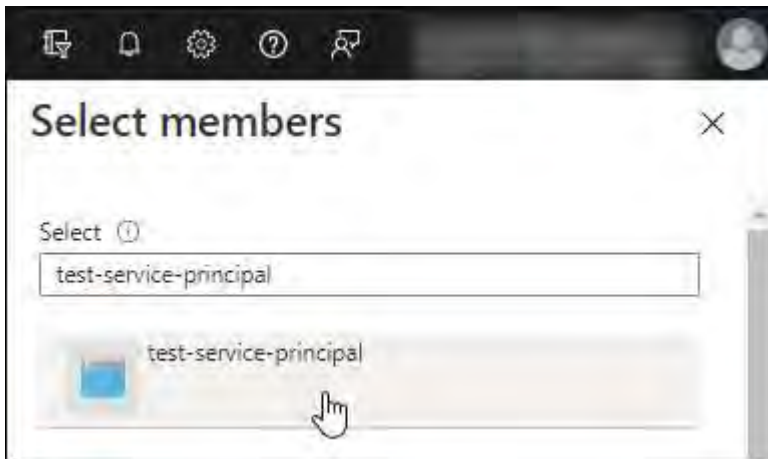
2. Assign the application to the role:

- a. From the Azure portal, open the **Subscriptions** service.
- b. Select the subscription.
- c. Select **Access control (IAM) > Add > Add role assignment**.
- d. In the **Role** tab, select the **BlueXP Operator** role and select **Next**.
- e. In the **Members** tab, complete the following steps:
  - Keep **User, group, or service principal** selected.
  - Select **Select members**.



- Search for the name of the application.

Here's an example:



- Select the application and select **Select**.
  - Select **Next**.
- f. Select **Review + assign**.

The service principal now has the required Azure permissions to deploy the Connector.

If you want to deploy Cloud Volumes ONTAP from multiple Azure subscriptions, then you must bind the service principal to each of those subscriptions. BlueXP enables you to select the subscription that you want to use when deploying Cloud Volumes ONTAP.

### Add Windows Azure Service Management API permissions

The service principal must have "Windows Azure Service Management API" permissions.

#### Steps




1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Select **API permissions > Add a permission**.
3. Under **Microsoft APIs**, select **Azure Service Management**.

**Request API permissions**

Select an API

Microsoft APIs   APIs my organization uses   My APIs


Commonly used Microsoft APIs

<p><b>Microsoft Graph</b></p> <p>Take advantage of the tremendous amount of data in Office 365, Enterprise Mobility + Security, and Windows 10. Access Azure AD, Excel, Intune, Outlook/Exchange, OneDrive, OneNote, SharePoint, Planner, and more through a single endpoint.</p> 		
<p><b>Azure Batch</b></p> <p>Schedule large-scale parallel and HPC applications in the cloud</p>	<p><b>Azure Data Catalog</b></p> <p>Programmatic access to Data Catalog resources to register, annotate and search data assets</p>	<p><b>Azure Data Explorer</b></p> <p>Perform ad-hoc queries on terabytes of data to build near real-time and complex analytics solutions</p>
<p><b>Azure Data Lake</b></p> <p>Access to storage and compute for big data analytic scenarios</p>	<p><b>Azure DevOps</b></p> <p>Integrate with Azure DevOps and Azure DevOps server</p>	<p><b>Azure Import/Export</b></p> <p>Programmatic control of import/export jobs</p>
<p><b>Azure Key Vault</b></p> <p>Manage your key vaults as well as the keys, secrets, and certificates within your Key Vaults</p>	<p><b>Azure Rights Management Services</b></p> <p>Allow validated users to read and write protected content</p>	<p><b>Azure Service Management</b></p> <p>Programmatic access to much of the functionality available through the Azure portal</p>
<p><b>Azure Storage</b></p> <p>Secure, massively scalable object and data lake storage for unstructured and semi-structured data</p>	<p><b>Customer Insights</b></p> <p>Create profile and interaction models for your products</p>	<p><b>Data Export Service for Microsoft Dynamics 365</b></p> <p>Export data from Microsoft Dynamics CRM organization to an external destination</p>

4. Select **Access Azure Service Management as organization users** and then select **Add permissions**.

## Request API permissions

[< All APIs](#)

 Azure Service Management  
<https://management.azure.com/> [Docs](#)

What type of permissions does your application require?

### Delegated permissions


Your application needs to access the API as the signed-in user.

### Application permissions

Your application runs as a background service or daemon without a signed-in user.

Select permissions

[expand all](#)

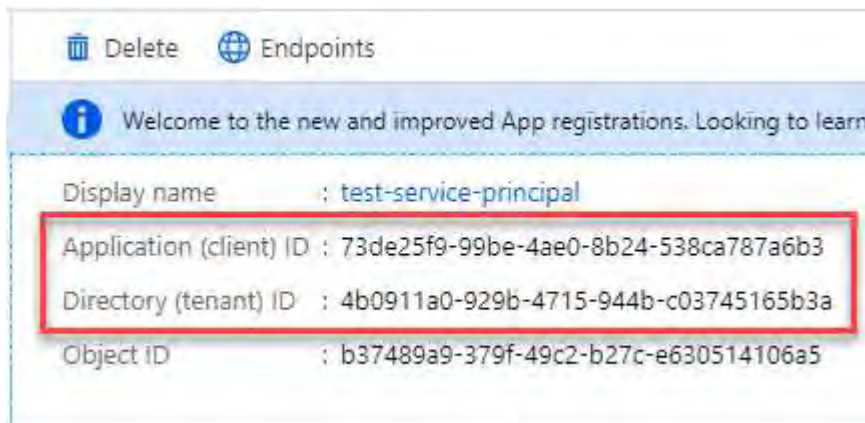
PERMISSION	ADMIN CONSENT REQUIRED
<input checked="" type="checkbox"/> <b>user_impersonation</b> Access Azure Service Management as organization users (preview) 	-

## Get the application ID and directory ID

When you add the Azure account to BlueXP, you need to provide the application (client) ID and the directory (tenant) ID for the application. BlueXP uses the IDs to programmatically sign in.

### Steps

1. In the **Microsoft Entra ID** service, select **App registrations** and select the application.
2. Copy the **Application (client) ID** and the **Directory (tenant) ID**.



When you add the Azure account to BlueXP, you need to provide the application (client) ID and the directory (tenant) ID for the application. BlueXP uses the IDs to programmatically sign in.

## Create a client secret

You need to create a client secret and then provide BlueXP with the value of the secret so BlueXP can use it to authenticate with Microsoft Entra ID.

### Steps

1. Open the **Microsoft Entra ID** service.

2. Select **App registrations** and select your application.
3. Select **Certificates & secrets > New client secret**.
4. Provide a description of the secret and a duration.
5. Select **Add**.
6. Copy the value of the client secret.

### Client secrets

A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.

[+ New client secret](#)

DESCRIPTION	EXPIRES	VALUE
test secret	8/16/2020	*sZ1jSe2By:D*-ZR0V4NLfdAcY7:+0vA

Copy to clipboard

You now have a client secret that BlueXP can use it to authenticate with Microsoft Entra ID.

### Result

Your service principal is now setup and you should have copied the application (client) ID, the directory (tenant) ID, and the value of the client secret. You need to enter this information in BlueXP when you add an Azure account.

### Add the credentials to BlueXP

After you provide an Azure account with the required permissions, you can add the credentials for that account to BlueXP. Completing this step enables you to launch Cloud Volumes ONTAP using different Azure credentials.

### Before you begin

If you just created these credentials in your cloud provider, it might take a few minutes until they are available for use. Wait a few minutes before you add the credentials to BlueXP.

### Before you begin

You need to create a Connector before you can change BlueXP settings. [Learn how to create a Connector.](#)

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.




2. Select **Add Credentials** and follow the steps in the wizard.
  - a. **Credentials Location:** Select **Microsoft Azure > Connector**.
  - b. **Define Credentials:** Enter information about the Microsoft Entra service principal that grants the required permissions:
    - Application (client) ID
    - Directory (tenant) ID

- Client Secret
- c. **Marketplace Subscription:** Associate a Marketplace subscription with these credentials by subscribing now or by selecting an existing subscription.
- d. **Review:** Confirm the details about the new credentials and select **Add**.

## Result

You can now switch to different set of credentials from the Details and Credentials page [when creating a new working environment](#)



## Manage existing credentials

Manage the Azure credentials that you've already added to BlueXP by associating a Marketplace subscription, editing credentials, and deleting them.

## Associate an Azure Marketplace subscription to credentials

After you add your Azure credentials to BlueXP, you can associate an Azure Marketplace subscription to those credentials. The subscription enables you to create a pay-as-you-go Cloud Volumes ONTAP system, and to use other BlueXP services.

There are two scenarios in which you might associate an Azure Marketplace subscription after you've already added the credentials to BlueXP:

- You didn't associate a subscription when you initially added the credentials to BlueXP.
- You want to change the Azure Marketplace subscription that is associated with Azure credentials.

Replacing the current marketplace subscription with a new subscription changes the marketplace subscription for any existing Cloud Volumes ONTAP working environments and all new working environments.

## Before you begin

You need to create a Connector before you can change BlueXP settings. [Learn how to create a Connector.](#)

## Steps

1. In the upper right of the console, select the Settings icon, and select **Credentials**.
2. Select the action menu for a set of credentials and then select **Configure Subscription**.

You must select credentials that are associated with a Connector. You can't associate a marketplace subscription with credentials that are associated with BlueXP.

3. To associate the credentials with an existing subscription, select the subscription from the down-down list and select **Configure**.
4. To associate the credentials with a new subscription, select **Add Subscription > Continue** and follow the steps in the Azure Marketplace:
  - a. If prompted, log in to your Azure account.
  - b. Select **Subscribe**.
  - c. Fill out the form and select **Subscribe**.
  - d. After the subscription process is complete, select **Configure account now**.

You'll be redirected to BlueXP.

- e. From the **Subscription Assignment** page:
  - Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
  - In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

- Select **Save**.

The following video shows the steps to subscribe from the Azure Marketplace:

[Subscribe to NetApp Intelligent Services from the Azure Marketplace](#)

## Edit credentials

Edit your Azure credentials in BlueXP by modifying the details about your Azure service credentials. For example, you might need to update the client secret if a new secret was created for the service principal application.

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.
2. On the **Organization credentials** or **Account credentials** page, select the action menu for a set of credentials and then select **Edit Credentials**.
3. Make the required changes and then select **Apply**.

## Delete credentials

If you no longer need a set of credentials, you can delete them from BlueXP. You can only delete credentials that aren't associated with a working environment.

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.
2. On the **Organization credentials** or **Account credentials** page, select the action menu for a set of credentials and then select **Delete Credentials**.
3. Select **Delete** to confirm.

## Google Cloud

### Learn about Google Cloud projects and permissions

Learn how BlueXP uses Google Cloud credentials to perform actions on your behalf and how those credentials are associated with marketplace subscriptions. Understanding these details can be helpful as you manage the credentials for one or more Google Cloud projects. For example, you might want to learn about the service account that's associated with the Connector VM.

### Project and permissions for BlueXP

Before you can use BlueXP to manage resources in your Google Cloud project, you must first deploy a Connector. The Connector can't be running on your premises, or in a different cloud provider.

Two sets of permissions must be in place before you deploy a Connector directly from BlueXP:

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

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





To learn how to set up permissions, refer to the following pages:

- [Set up Google Cloud permissions for standard mode](#)
- [Set up permissions for restricted mode](#)
- [Set up permissions for private mode](#)

#### Credentials and marketplace subscriptions

When you deploy a Connector in Google Cloud, BlueXP creates a default set of credentials for the Google Cloud service account in the project in which the Connector resides. These credentials must be associated with a Google Cloud Marketplace subscription so that you can pay for Cloud Volumes ONTAP at an hourly rate (PAYGO) and use other BlueXP services.

[Learn how to associate a Google Cloud Marketplace subscription.](#)

Note the following about Google Cloud credentials and marketplace subscriptions:

- Only one set of Google Cloud credentials can be associated with a Connector
- You can associate only one Google Cloud Marketplace subscription with the credentials
- You can replace an existing marketplace subscription with a new subscription

#### Project for Cloud Volumes ONTAP

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

- [Learn how to set up the service account](#)
- [Learn how to deploy Cloud Volumes ONTAP in Google Cloud and select a project](#)

#### Manage Google Cloud credentials and subscriptions for BlueXP

You can manage the Google Cloud credentials that are associated with the Connector

VM instance by associating a marketplace subscription and by troubleshooting the subscription process. Both of these tasks ensure that you can use your marketplace subscription to pay for data services.

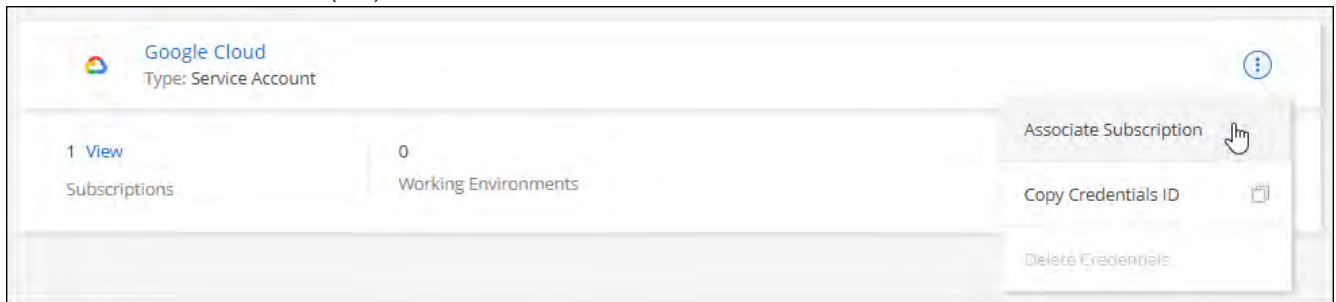
### Associate a Marketplace subscription with Google Cloud credentials

When you deploy a Connector in Google Cloud, BlueXP creates a default set of credentials that are associated with the Connector VM instance. At any time, you can change the Google Cloud Marketplace subscription that is associated with these credentials. The subscription enables you to create a pay-as-you-go Cloud Volumes ONTAP system, and to use other data services.

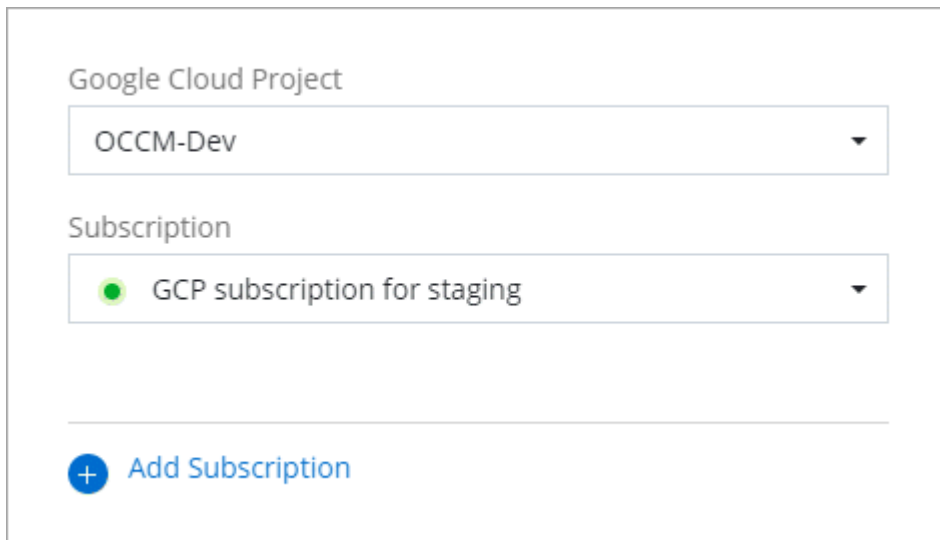
Replacing the current marketplace subscription with a new subscription changes the marketplace subscription for any existing Cloud Volumes ONTAP working environments and all new working environments.

### Steps

1. In the upper right of the console, select the Settings icon, and select **Credentials**.
2. Select the action menu for a set of credentials and then select **Configure Subscription**.  
+new screenshot needed (TS)



3. To configure an existing subscription with the selected credentials, select a Google Cloud project and subscription from the drop-down list, and then select **Configure**.



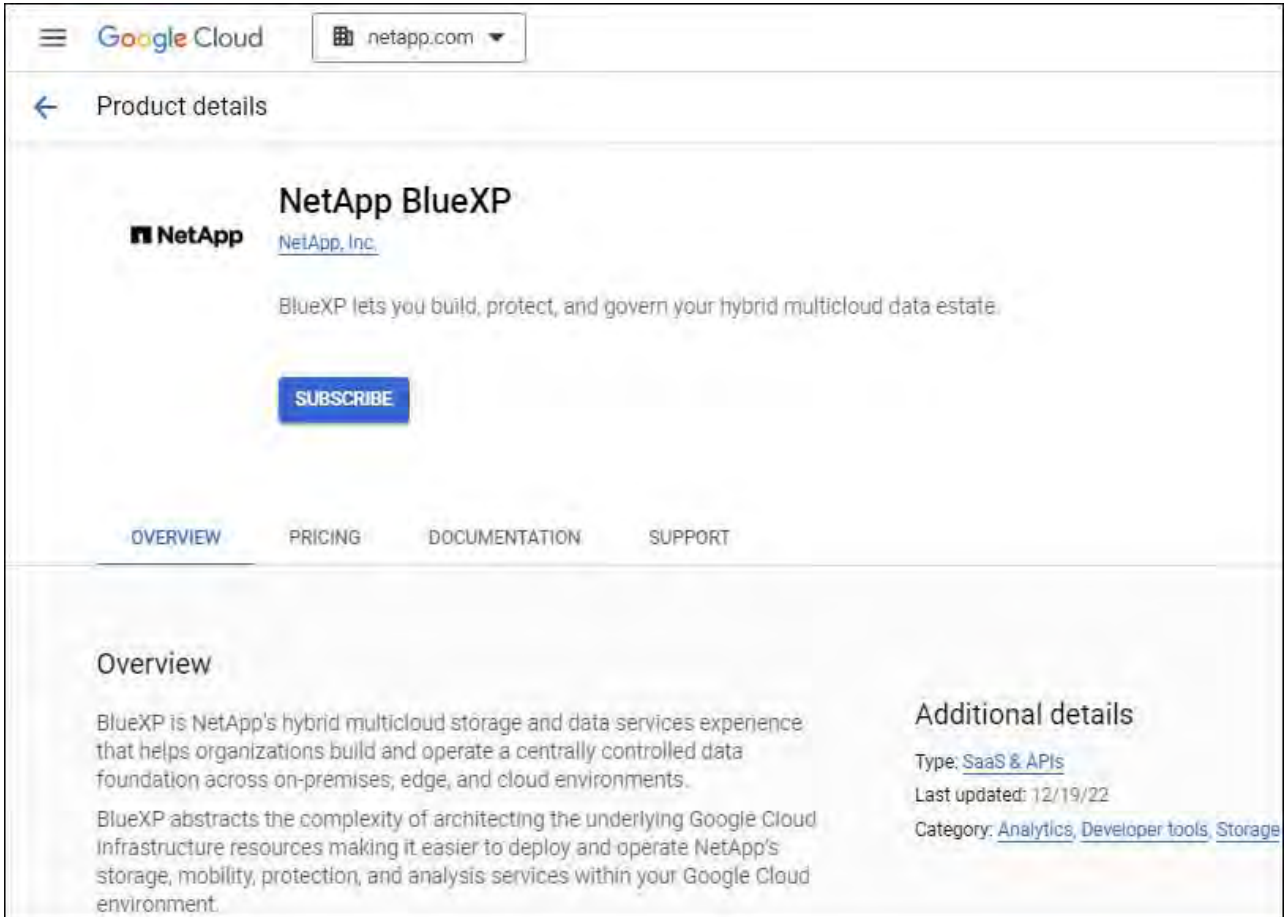
4. If you don't already have a subscription, select **Add Subscription > Continue** and follow the steps in the Google Cloud Marketplace.



Before you complete the following steps, ensure that you have both Billing Admin privileges in your Google Cloud account as well as a BlueXP login.



- a. After you're redirected to the [NetApp Intelligent Services page on the Google Cloud Marketplace](#), ensure that the correct project is selected at the top navigation menu.

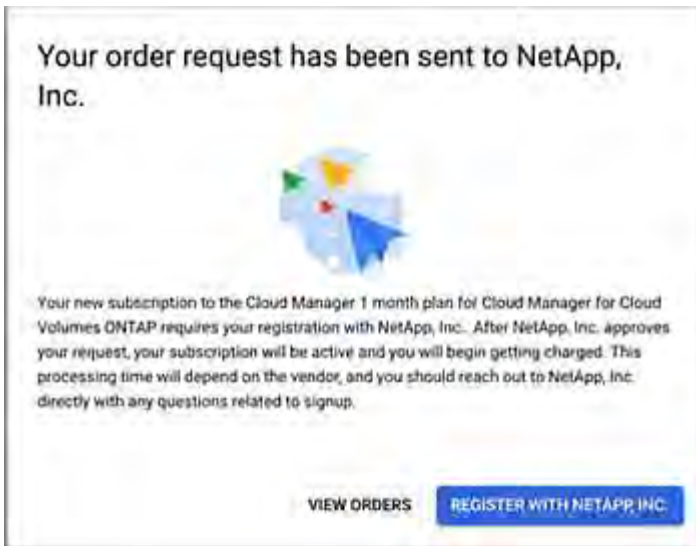


- b. Select **Subscribe**.
- c. Select the appropriate billing account and agree to the terms and conditions.
- d. Select **Subscribe**.

This step sends your transfer request to NetApp.

- e. On the pop-up dialog box, select **Register with NetApp, Inc.**

This step must be completed to link the Google Cloud subscription with your BlueXP organization or account. The process of linking a subscription isn't complete until you are redirected from this page and then sign in to BlueXP.



f. Complete the steps on the **Subscription Assignment** page:



If someone from your organization has already subscribed to the NetApp BlueXP subscription from your billing account, then you will be redirected to [the Cloud Volumes ONTAP page on the BlueXP website](#) instead. If this is unexpected, contact your NetApp sales team. Google enables only one subscription per Google billing account.

- Select the BlueXP organizations or accounts that you'd like to associate this subscription with.
- In the **Replace existing subscription** field, choose whether you'd like to automatically replace the existing subscription for one organization or account with this new subscription.

BlueXP replaces the existing subscription for all credentials in the organization or account with this new subscription. If a set of credentials wasn't ever associated with a subscription, then this new subscription won't be associated with those credentials.

For all other organizations or accounts, you'll need to manually associate the subscription by repeating these steps.

- Select **Save**.

The following video shows the steps to subscribe from the Google Cloud Marketplace:

[Subscribe to BlueXP from the Google Cloud Marketplace](#)

g. Once this process is complete, navigate back to the Credentials page in BlueXP and select this new subscription.

Google Cloud Project

OCCM-Dev

Subscription

GCP subscription for staging

+ Add Subscription

### Troubleshoot the Marketplace subscription process

Sometimes subscribing to NetApp Intelligent Services through the Google Cloud Marketplace can become fragmented due to incorrect permissions or accidentally not following the redirection to the BlueXP website. If this happens, use the following steps to complete the subscription process.

### Steps

1. Navigate to the [NetApp BlueXP page on the Google Cloud Marketplace](#) to check on the state of the order. If the page states **Manage on Provider**, scroll down and select **Manage Orders**.



- If the order shows a green check mark and this is unexpected, somebody else from the organization using the same billing account might already be subscribed. If this is unexpected or you require the details of this subscription, contact your NetApp sales team.

Status	Order number	Plan	Discount	Start date ↓	Plan duration	End date	Payment Schedule	Auto-renew	Next plan
✓	2eebbc...	Cloud Manager	-	10/21/21	1 month	-	Postpay	N/A	N/A

- If the order shows a clock and **Pending** status, go back to the marketplace page and choose **Manage on Provider** to complete the process as documented above.

Status	Order number	Plan	Discount	Start date ↓	Plan duration	End date	Payment Schedule	Auto-renew	Next plan
🕒	d56c66...	Cloud Manager	-	Pending	1 month	Pending	Postpay	N/A	N/A

## Manage NSS credentials associated with BlueXP

Associate a NetApp Support Site account with your BlueXP organization to enable key workflows for Cloud Volumes ONTAP. These NSS credentials are associated with the entire BlueXP organization.

BlueXP also supports associating one NSS account per BlueXP user account. [Learn how to manage user-level credentials.](#)

- [Learn about BlueXP deployment modes](#)
- [Learn about BlueXP identity and access management](#)

### Overview

Associating NetApp Support Site credentials with your specific BlueXP account serial number is required to enable the following tasks in BlueXP:

- 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.

- Registering pay-as-you-go Cloud Volumes ONTAP systems

Providing your NSS account is required to activate support for your system and to gain access to NetApp technical support resources.

- Upgrading Cloud Volumes ONTAP software to the latest release

These credentials are associated with your specific BlueXP account serial number. Users who belong to the BlueXP organization or account can access these credentials from **Support > NSS Management**.

### Add an NSS account

You can add and manage your NetApp Support Site accounts for use with BlueXP from the Support Dashboard within BlueXP.

When you have added your NSS account, BlueXP can use this information for things like license downloads, software upgrade verification, and future support registrations.

You can associate multiple NSS accounts with your BlueXP organization; however, you cannot have customer accounts and partner accounts within the same organization.



NetApp uses Microsoft Entra ID as the identity provider for authentication services specific to support and licensing.

### 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. Select **Continue** to be redirected to a Microsoft login page.
4. At the login page, provide your NetApp Support Site registered email address and password.

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.

### What's next?

Users can now select the account when creating new Cloud Volumes ONTAP systems and when registering existing Cloud Volumes ONTAP systems.

- [Launching Cloud Volumes ONTAP in AWS](#)
- [Launching Cloud Volumes ONTAP in Azure](#)
- [Launching Cloud Volumes ONTAP in Google Cloud](#)
- [Registering pay-as-you-go systems](#)

### Update NSS credentials

For security reasons, you must update your NSS credentials every 90 days. You'll be notified in the BlueXP notification center if your NSS credential has expired. [Learn about the Notification Center](#).

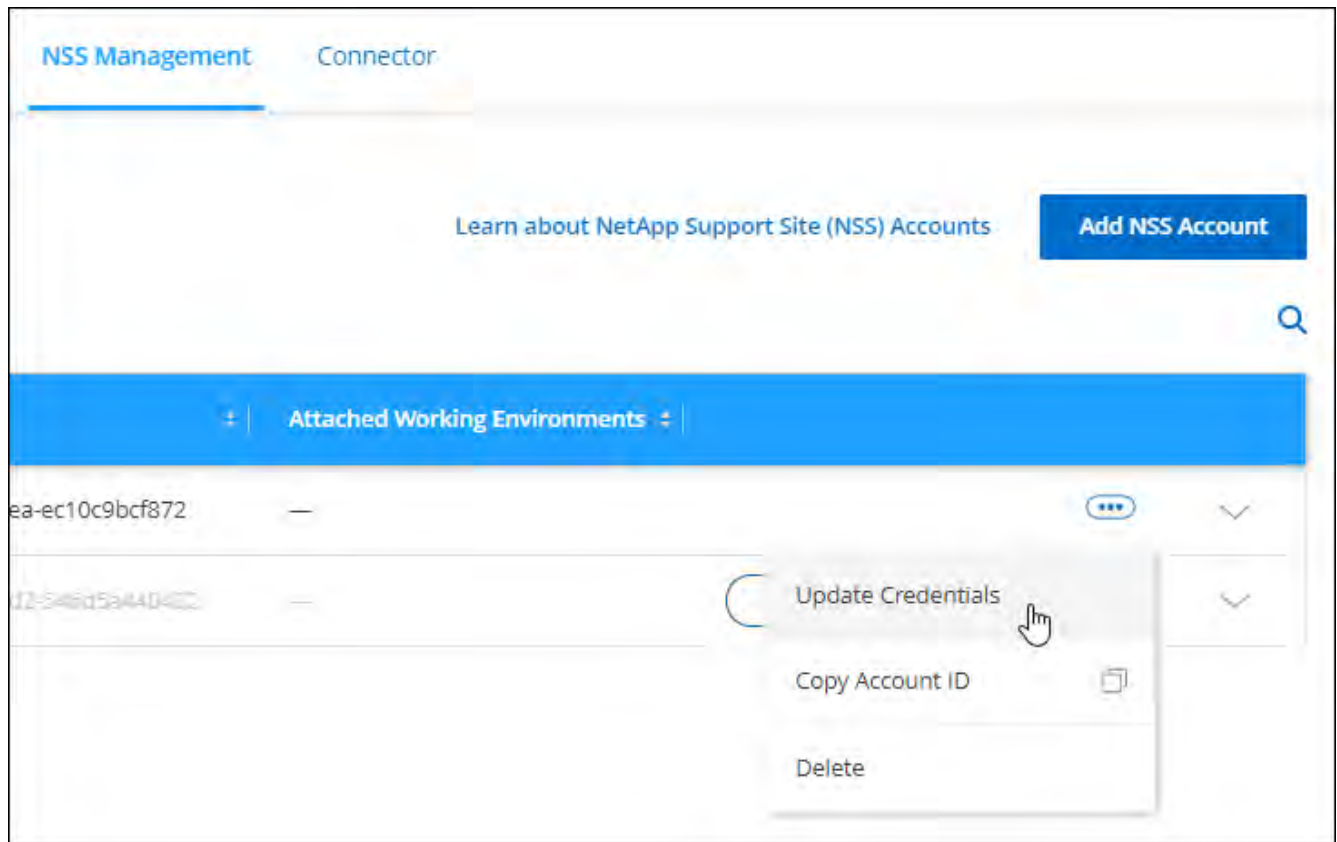
Expired credentials can disrupt the following, but are not limited to:

- License updates in digital wallet, which means you won't be able to take advantage of newly purchased capacity.
- Ability to submit and track support cases.

Additionally, you can update the NSS credentials associated with your organization if you want to change the NSS account associated with your BlueXP organization. For example, if the person associated with your NSS account has left your company.

### Steps

1. In the upper right of the BlueXP console, select the Help icon, and select **Support**.
2. Select **NSS Management**.
3. For the NSS account that you want to update, select **☰** and then select **Update Credentials**.



4. 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 related to support and licensing.

5. At the login page, provide your NetApp Support Site registered email address and password.

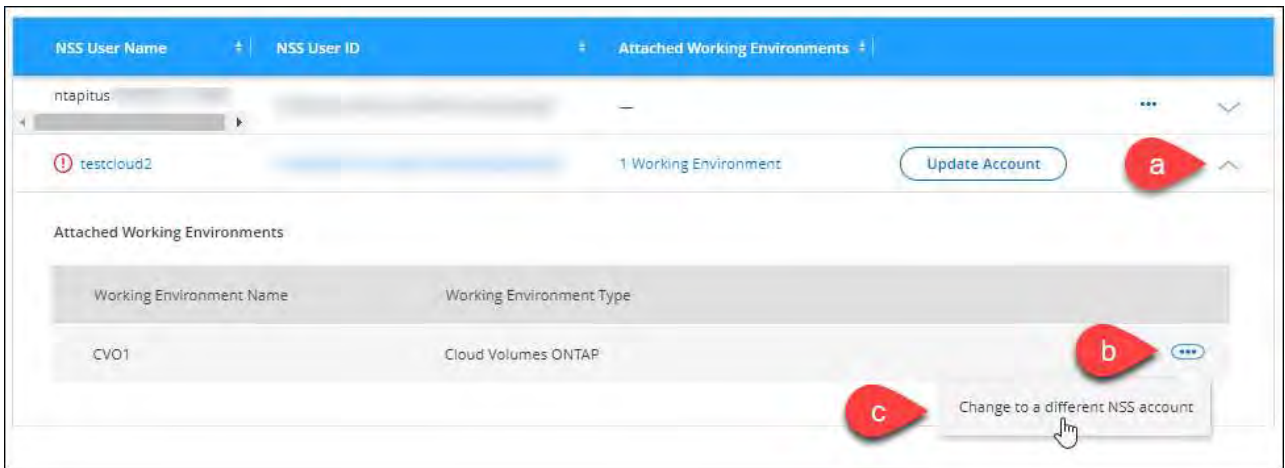
### Attach a working environment to a different NSS account

If your organization has multiple NetApp Support Site accounts, you can change which account is associated with a Cloud Volumes ONTAP system.

You must first have associated the account with BlueXP.

#### Steps

1. In the upper right of the BlueXP console, select the Help icon, and select **Support**.
2. Select **NSS Management**.
3. Complete the following steps to change the NSS account:
  - a. Expand the row for the NetApp Support Site account that the working environment is currently associated with.
  - b. For the working environment that you want to change the association for, select **...**
  - c. Select **Change to a different NSS account**.



d. Select the account and then select **Save**.

### Display the email address for an NSS account

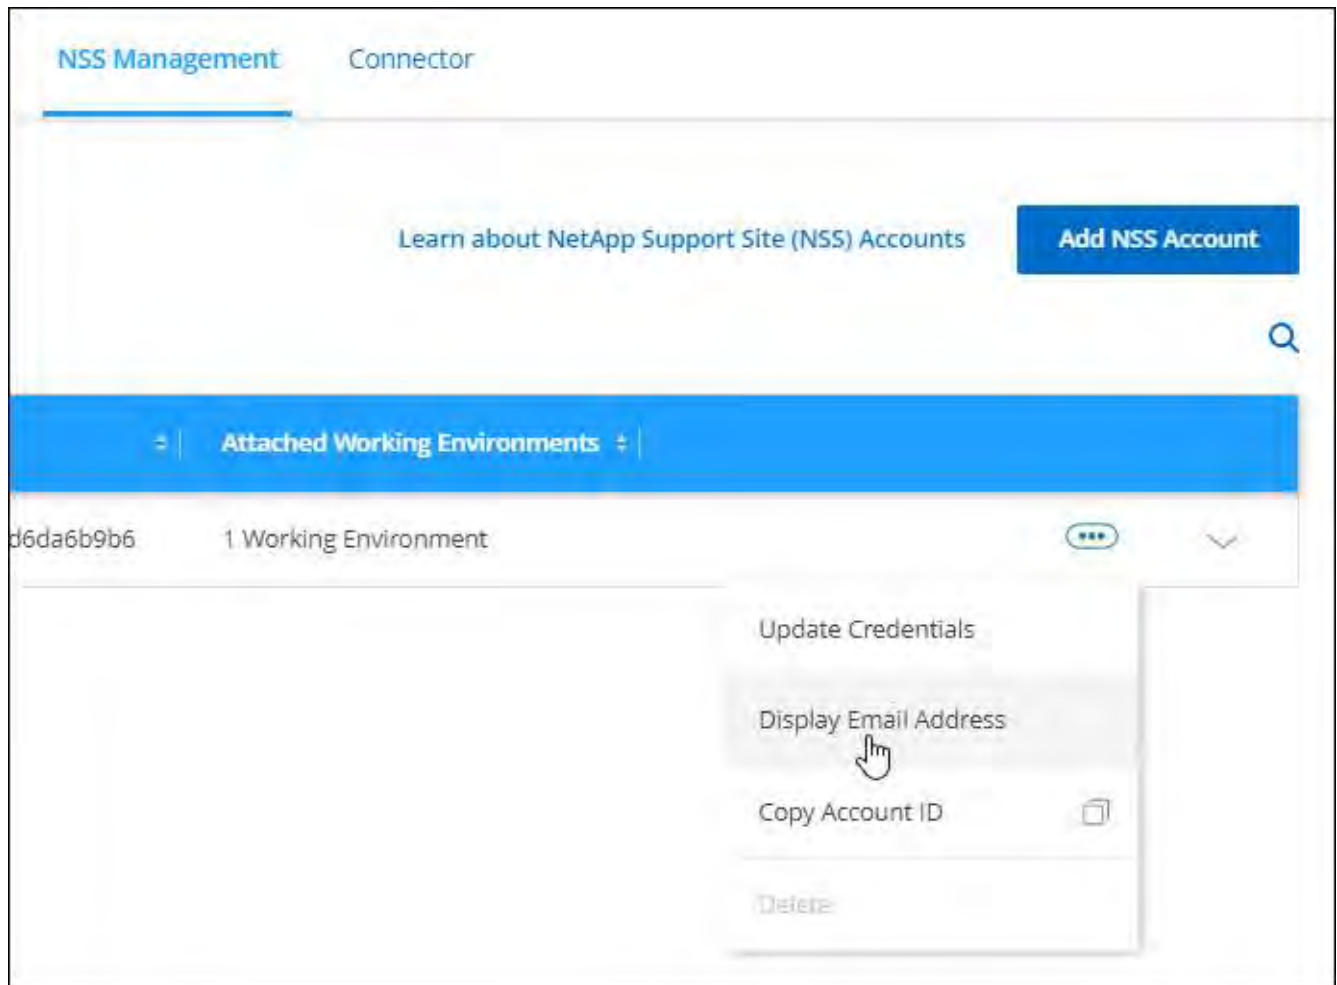
For security, the email address associated with an NSS account is not displayed by default. You can view the email address and associated user name for an NSS account.



When you go to the NSS Management page, BlueXP generates a token for each account in the table. That token includes information about the associated email address. The token is removed when you leave the page. The information is never cached, which helps protect your privacy.

### Steps

1. In the upper right of the BlueXP console, select the Help icon, and select **Support**.
2. Select **NSS Management**.
3. For the NSS account that you want to update, select **...** and then select **Display Email Address**. You can use the copy button to copy the email address.



## Remove an NSS account

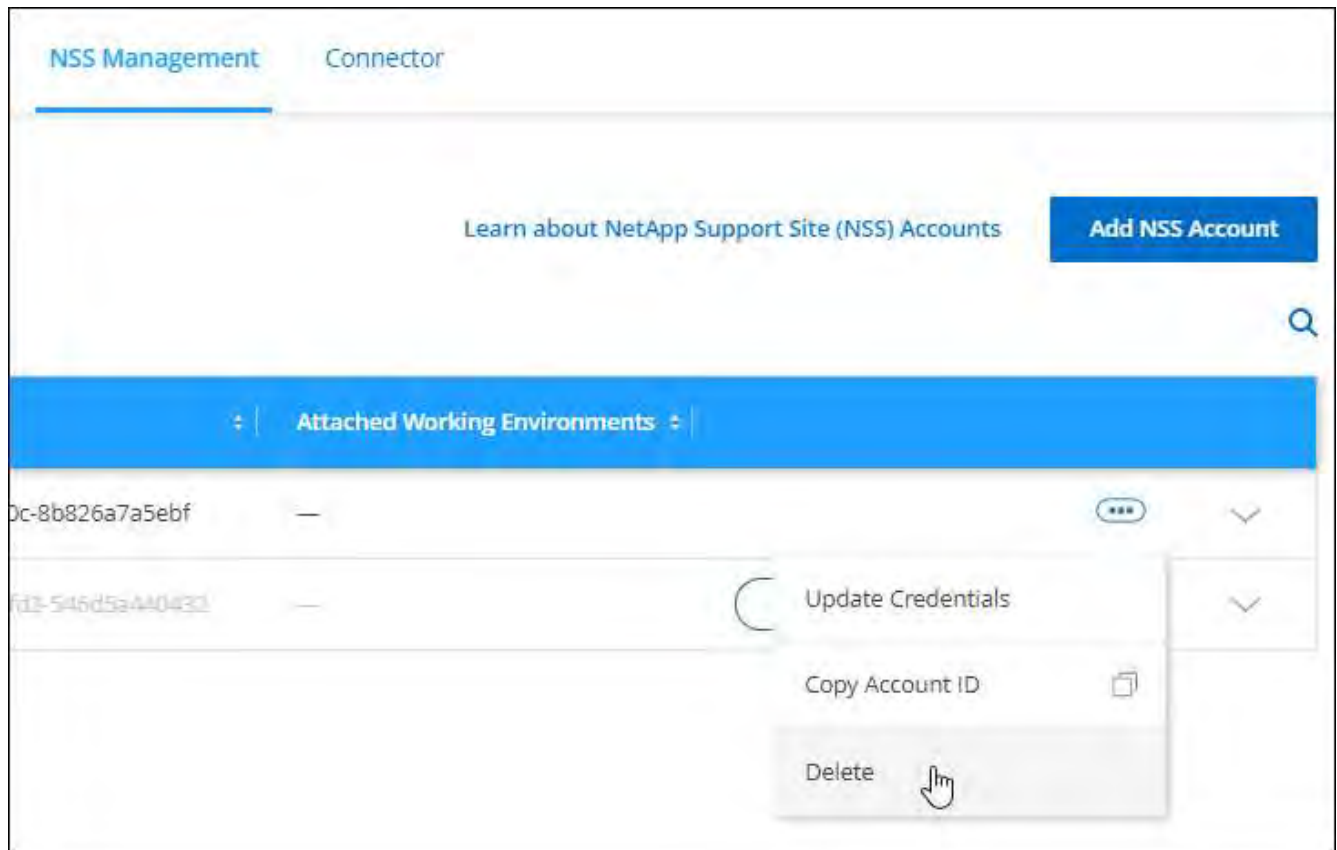
Delete any of the NSS accounts that you no longer want to use with BlueXP.

You can't delete an account that is currently associated with a Cloud Volumes ONTAP working environment. You first need to [attach those working environments to a different NSS account](#).

### Steps

1. In the upper right of the BlueXP console, select the Help icon, and select **Support**.
2. Select **NSS Management**.
3. For the NSS account that you want to delete, select **...** and then select **Delete**.





4. Select **Delete** to confirm.

## Manage credentials associated with your BlueXP login

Depending on the actions that you've taken in BlueXP, you might have associated ONTAP credentials and NetApp Support Site (NSS) credentials with your BlueXP user login. You can view and manage those credentials in BlueXP after you've associated them. For example, if you change the password for these credentials, then you'll need to update the password in BlueXP.

### ONTAP credentials

Users need ONTAP admin credentials to discover ONTAP clusters in BlueXP. However, ONTAP System Manager access depends on whether or not you are using a Connector.

#### Without a Connector

Users are prompted to enter their ONTAP credentials to access ONTAP System Manager for the cluster. Users can choose to save these credentials in BlueXP which means they won't be prompted to enter them each time. User credentials are only visible to the respective user and can be managed from the User credentials page.

#### With a Connector

By default, users are not prompted to enter their ONTAP credentials to access ONTAP System Manager. However, a BlueXP administrator (with the Organization admin role) can configure BlueXP to prompt users to enter their ONTAP credentials. When this setting is enabled, users need enter their ONTAP credentials each time.

[Learn more.](#)

## NSS credentials

The NSS credentials associated with your BlueXP login enable support registration, case management, and access to Digital Advisor.

- When you access **Support > Resources** and register for support, you're prompted to associate NSS credentials with your BlueXP login.

This registers your organization or account for support and activates support entitlement. Only one user in your BlueXP organization or account must associate a NetApp Support Site account with their BlueXP login to register for support and activate support entitlement. After this is completed, the **Resources** page shows that your account is registered for support.

[Learn how to register for support](#)

- When you access **Support > Case Management**, you're prompted to enter your NSS credentials, if you haven't already done so. This page enables you to create and manage the support cases associated with your NSS account and with your company.
- When you access Digital Advisor in BlueXP, you're prompted to log in to Digital Advisor by entering your NSS credentials.

Note the following about the NSS account associated with your BlueXP login:

- The account is managed at the user level, which means it isn't viewable by other users who log in.
- There can be only one NSS account associated with Digital Advisor and support case management, per user.
- If you're trying to associate a NetApp Support Site account with a Cloud Volumes ONTAP working environment, you can only choose from the NSS accounts that were added to the BlueXP organization or account that you are a member of.

NSS account-level credentials are different than the NSS account that's associated with your BlueXP login. NSS account-level credentials enable you to deploy Cloud Volumes ONTAP with BYOL, register PAYGO systems, and upgrade its software.

[Learn more about using NSS credentials with your BlueXP organization or account.](#)

## Manage your user credentials

Manage your user credentials by updating the user name and password or by deleting the credentials.

### Steps

1. In the upper right of the BlueXP console, select the Settings icon, and select **Credentials**.
2. Select **User Credentials**.
3. If you don't have any user credentials yet, you can select **Add NSS credentials** to add your NetApp Support Site account.
4. Manage existing credentials by choosing the following options from the Actions menu:
  - **Update credentials**: Update the user name and password for the account.
  - **Delete credentials**: Remove the account associated with your BlueXP user account.

### Result

BlueXP updates your credentials, and you see the changes when accessing the ONTAP cluster, digital advisor,

or the Case Management page.

## Monitor BlueXP operations

You can monitor the status of the operations that BlueXP is performing to see if there are any issues that you need to address. You can view the status in the Timeline, the Notification Center, or have notifications sent to your email.

The table compares the Timeline and Notification Center to highlight their features.


Notification Center	Timeline
Shows high level status for events and actions	Provides details for each event or action for further investigation
Shows status for the current login session (the information does not appear in the Notification Center after you log off)	Retains status for the last month
Shows only actions initiated in the user interface	Shows all actions from the UI or APIs
Shows user-initiated actions	Shows all actions, whether user-initiated or system-initiated
Filter results by importance	Filter by service, action, user, status, and more
Provides the ability to email notifications to users and to others	No email capability

### Audit user activity from the BlueXP timeline

The Timeline shows the actions that users completed to manage your organization or account. This includes management actions such as associating users, creating working environments, creating Connectors, and more.

The Timeline helps identify who performed an action or its status.

#### Steps

1. In the upper right of the BlueXP console, select  > **Timeline**.
2. Use the filters above the table to change which actions display in the table.

For example, you can use the **Service** filter to show actions related to a specific BlueXP service, or you can use the **User** filter to show actions related to a specific user account.

### Download audit logs from the Timeline

You can download the audit logs from the Timeline to a CSV file. This enables you to keep a record of the actions that users performed in your organization. The downloaded CSV file contains all available columns from the Timeline, regardless of which ones you are filtering or displaying in the Timeline.

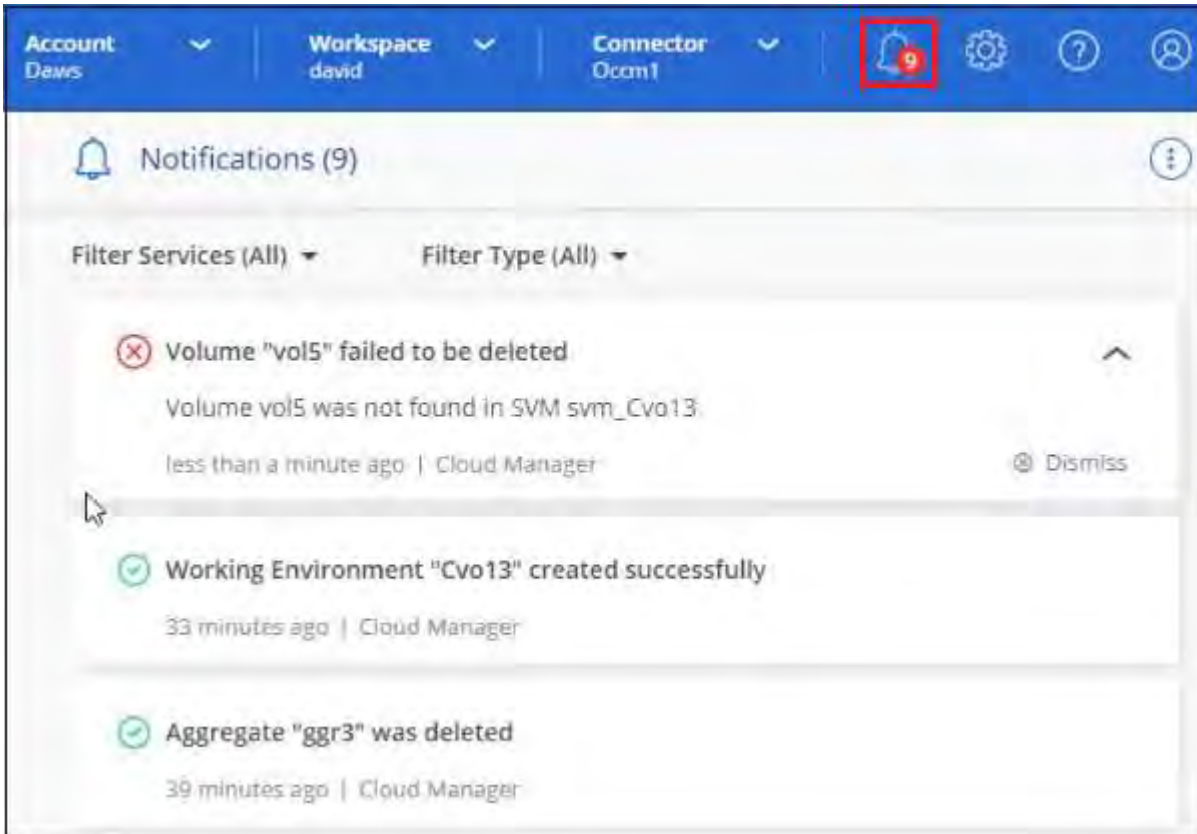
#### Steps

1. In the Timeline, select the download icon in the upper right corner of the table.

## Monitor activities using the Notification Center

Notifications track the progress of your BlueXP operations to verify success. They enable you to view the status for many BlueXP actions that you initiated during your current login session. Not all BlueXP services report information into the Notification Center at this time.

You can display the notifications by selecting the notification bell (🔔) in the menu bar. The color of the little bubble in the bell indicates the highest level severity notification that is active. So if you see a red bubble, it means there's an important notification that you should look at.



You can also configure BlueXP to send certain types of notifications by email so you can be informed of important system activity even when you're not logged into the system. Emails can be sent to any users who are part of your BlueXP organization or account, or to any other recipients who need to be aware of certain types of system activity. See how to [set email notification settings](#).

### Comparing the Notification Center with BlueXP alerts

The Notification Center enables you to view the status of operations you've initiated from BlueXP and set up alert notifications for certain types of system activities. Meanwhile, BlueXP alerts enables you to view issues or potential risks in your ONTAP storage environment related to capacity, availability, performance, protection, and security.

[Learn more about BlueXP alerts](#)

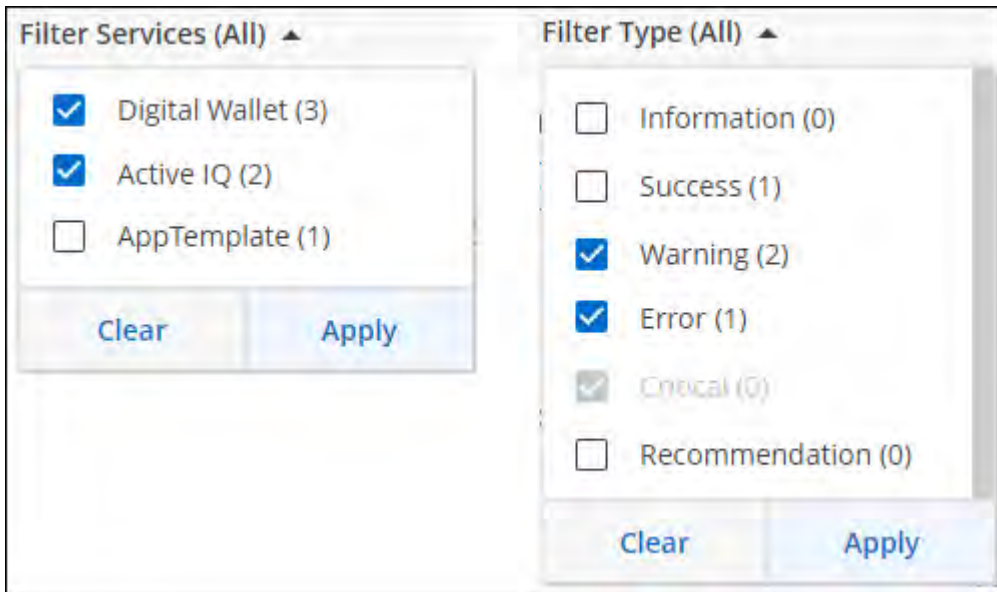
### Notification types

BlueXP classifies notifications into the following categories:

Notification type	Description
Critical	A problem occurred that might lead to service disruption if corrective action is not taken immediately.
Error	An action or process ended with failure, or could lead to failure if corrective action is not taken.
Warning	An issue that you should be aware of to make sure it does not reach the critical severity. Notifications of this severity do not cause service disruption, and immediate corrective action might not be required.
Recommendation	A system recommendation for you to take an action to improve the system or a certain service; for example: costs saving, suggestion for new services, recommended security configuration, etc.
Information	A message that provides additional information about an action or process.
Success	An action or process completed successfully.

### Filter notifications

By default you'll see all active notifications in the Notification Center. You can filter the notifications that you see to show only those notifications that are important to you. You can filter by BlueXP "Service" and by notification "Type".




For example, if you want to see only "Error" and "Warning" notifications for BlueXP operations, select those entries and you'll see only those types of notifications.

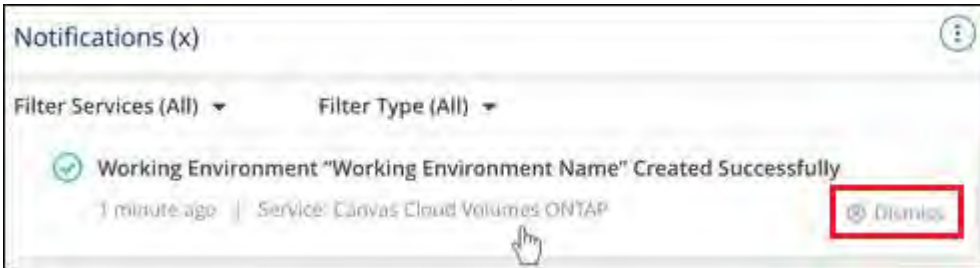
### Dismiss notifications

You can remove notifications from the page if you no longer need to see them. You can dismiss notifications individually or all at once.

To dismiss all notifications, in the Notification Center, select  and select **Dismiss All**.



To dismiss individual notifications, hover your cursor over the notification and select **Dismiss**.



### Set email notification settings

You can send specific types of notifications by email so you can be informed of important system activity even when you're not logged into BlueXP. Emails can be sent to any users who are part of your BlueXP organization or account, or to any other recipients who need to be aware of certain types of system activity.



- BlueXP sends email notifications for the Connector, digital wallet, copy and sync, and backup and recovery.
- Sending email notifications is not supported when the Connector is installed in a site without internet access.

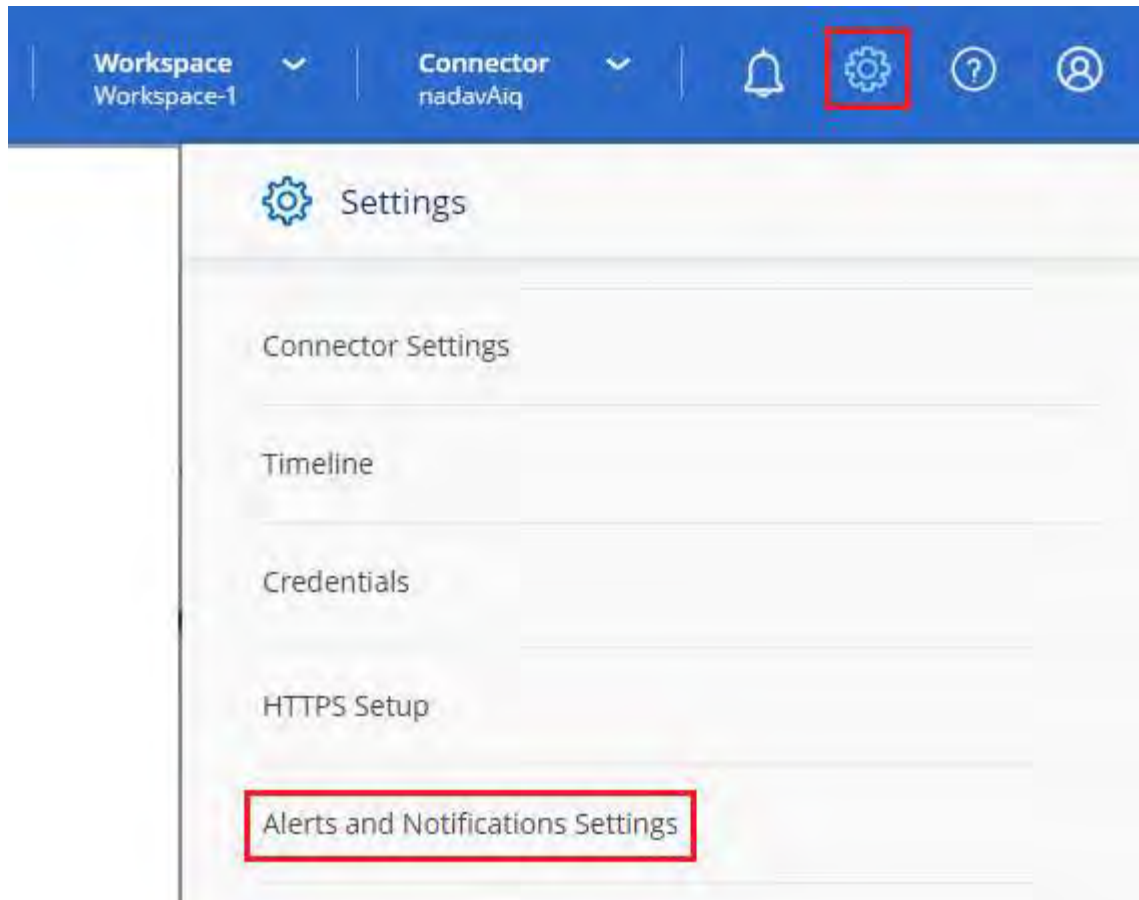
The filters you set in the Notification Center do not determine the types of notifications you'll receive by email. By default, any BlueXP admin will receive emails for all "Critical" and "Recommendation" notifications. These notifications are across all services - you can't choose to receive notifications for only certain services, for example Connectors or BlueXP backup and recovery.

All other users and recipients are configured not to receive any notification emails - so you'll need to configure notification settings for any additional users.

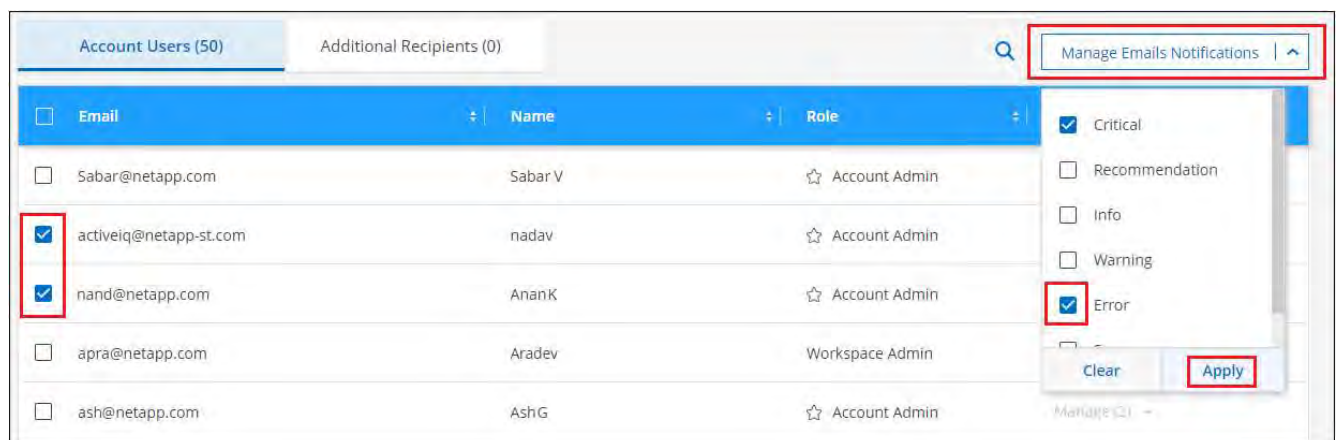
You must have the Organization admin role to customize the notifications settings.

### Steps

1. From the BlueXP menu bar, select **Settings > Alerts and Notifications Settings**.



2. Select a user, or multiple users, from either the *Users* tab or the *Additional Recipients* tab, and choose the type of notifications to be sent:
  - To make changes for a single user, select the menu in the Notifications column for that user, check the types of Notifications to be sent, and select **Apply**.
  - To make changes for multiple users, check the box for each user, select **Manage Email Notifications**, check the types of Notifications to be sent, and select **Apply**.

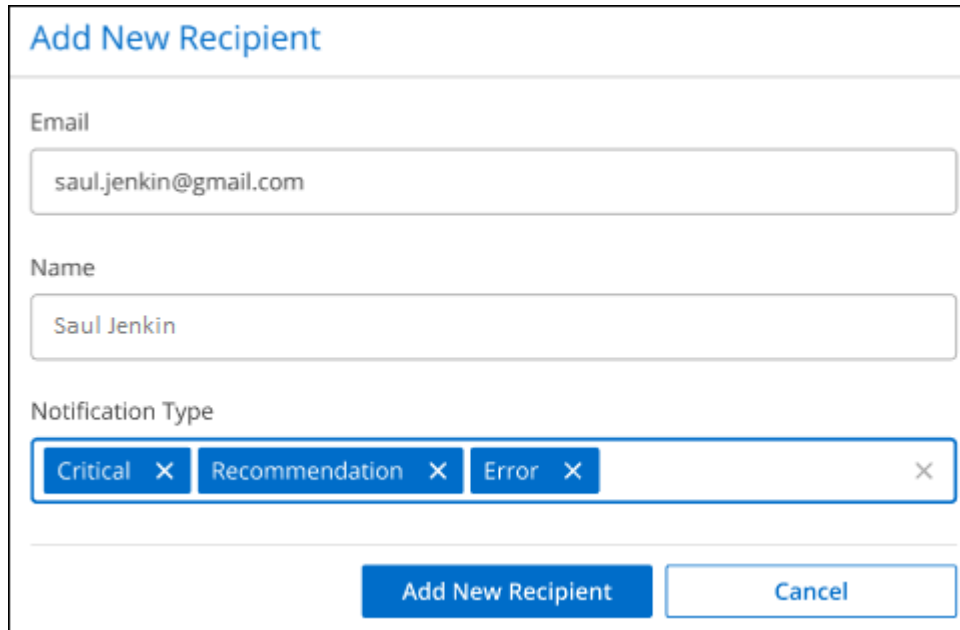


### Add additional email recipients

The users who appear in the *Users* tab are populated automatically from the users in your organization or account. You can add email addresses in the *Additional Recipients* tab for other people, or groups, who do not have access to BlueXP, but who need to be notified about certain types of alerts and notifications.

## Steps

1. From the Alerts and Notifications Settings page, select **Add New Recipients**.



The screenshot shows a form titled "Add New Recipient" with the following fields and controls:

- Email:** A text input field containing "saul.jenkin@gmail.com".
- Name:** A text input field containing "Saul Jenkin".
- Notification Type:** A multi-select dropdown menu with three selected items: "Critical", "Recommendation", and "Error". Each item has a small "x" icon to its right, and there is a larger "x" icon at the end of the menu.
- Buttons:** At the bottom right, there are two buttons: "Add New Recipient" (a solid blue button) and "Cancel" (a white button with a blue border).

2. Enter the name, email address, and select the types of Notifications that recipient will receive, and select **Add New Recipient**.



# Reference

## Connector maintenance console

### Connector maintenance console

You can use the Maintenance Console to configure the Connector to use a transparent proxy server.

#### Access the Maintenance Console

You can access the Maintenance Console from the Connector host. Navigate to the following directory:

```
/opt/application/netapp/service-manager-2/connector-maint-console
```

#### Transparent proxy commands

The Maintenance Console provides commands to configure the Connector to use a transparent proxy server.

#### View the current transparent proxy configuration

To view the current transparent proxy configuration, use the following command:

```
./connector-maint-console proxy get
```

#### Add a transparent proxy server

To add a transparent proxy server, use the following command, where `/home/ubuntu/myCA1.pem` is the path to the certificate file for the proxy server. The certificate file must be in PEM format:

```
./connector-maint-console proxy add -c /home/ubuntu/myCA1.pem
```

Ensure that the certificate file is in the same directory as the command or specify the full path to the certificate file.

#### Update the certificate for a transparent proxy server

To update the certificate for a transparent proxy server, use the following command, where `/home/ubuntu/myCA1.pem` is the path to the new certificate file for the proxy server. The certificate file must be in PEM format:

```
./connector-maint-console proxy update -c /home/ubuntu/myCA1.pem
```

Ensure that the certificate file is in the same directory as the command or specify the full path to the certificate file.

## Remove a transparent proxy server

To remove transparent proxy server, use the following command:

```
./connector-maint-console proxy remove
```

## View help for any command

To view help for any command, append `--help` to the command. For example, to view help for the `proxy add` command, use the following command:

```
./connector-maint-console proxy add --help
```

# Permissions

## Permissions summary for BlueXP

To use BlueXP features and services, you'll need to provide permissions so that BlueXP can perform operations in your cloud environment. Use the links on this page to quickly access the permissions that you need based on your goal.

### AWS permissions

BlueXP requires AWS permissions for the Connector and for individual services.

#### Connectors

Goal	Description	Link
Deploy the Connector from BlueXP	The user who creates a Connector from BlueXP needs specific permissions to deploy the instance in AWS.	<a href="#">Set up AWS permissions</a>
Provide permissions for the Connector	<p>When BlueXP launches the Connector, it attaches a policy to the instance that provides the permissions required to manage resources and processes in your AWS account.</p> <p>You need to set up the policy yourself if you launch a Connector from the AWS Marketplace, if you manually install the Connector, or if you <a href="#">add more AWS credentials to a Connector</a>.</p> <p>You also need to ensure that the policy is up to date as new permissions are added in subsequent releases.</p>	<a href="#">AWS permissions for the Connector</a>

### Backup and recovery

Goal	Description	Link
Back up on-premises ONTAP clusters to Amazon S3 with BlueXP backup and recovery	When activating backups on your ONTAP volumes, BlueXP backup and recovery prompts you to enter an access key and secret for an IAM user that has specific permissions.	<a href="#">Set up S3 permissions for backups</a>

### Cloud Volumes ONTAP

Goal	Description	Link
Provide permissions for Cloud Volumes ONTAP nodes	An IAM role must be attached to each Cloud Volumes ONTAP node in AWS. The same is true for the HA mediator. The default option is to let BlueXP create the IAM roles for you, but you can use your own when creating the working environment.	<a href="#">Learn how to set up the IAM roles yourself</a>

### Copy and sync

Goal	Description	Link
Deploy the data broker in AWS	The AWS user account that you use to deploy the data broker must have specific permissions.	<a href="#">Permissions required to deploy the data broker in AWS</a>
Provide permissions for the data broker	When BlueXP copy and sync deploys the data broker, it creates an IAM role for the data broker instance. You can deploy the data broker using your own IAM role, if you prefer.	<a href="#">Requirements to use your own IAM role with the AWS data broker</a>
Enable AWS access for a manually installed data broker	If you use the data broker with a sync relationship that includes an S3 bucket, then you should prepare the Linux host for AWS access. When you install the data broker, you'll need to provide AWS keys for an IAM user that has programmatic access and specific permissions.	<a href="#">Enabling access to AWS</a>

### FSx for ONTAP

Goal	Description	Link
Create and manage FSx for ONTAP	To create or manage an Amazon FSx for NetApp ONTAP working environment, you need to add AWS credentials to BlueXP by providing the ARN of an IAM role that gives BlueXP the permissions needed to create the working environment.	<a href="#">Learn how to set up AWS credentials for FSx</a>

### Tiering

Goal	Description	Link
Tier on-premises ONTAP clusters to Amazon S3	When you enable BlueXP tiering to AWS, the wizard prompts you to enter an access key and secret key. These credentials are passed to the ONTAP cluster so that ONTAP can tier data to the S3 bucket.	<a href="#">Set up S3 permissions for tiering</a>

## Azure permissions

BlueXP requires Azure permissions for the Connector and for individual services.

### Connectors

Goal	Description	Link
Deploy the Connector from BlueXP	When you deploy a Connector from BlueXP, you need to use an Azure account or service principal that has permissions to deploy the Connector VM in Azure.	<a href="#">Set up Azure permissions</a>
Provide permissions for the Connector	<p>When BlueXP deploys the Connector VM in Azure, it creates a custom role that provides the permissions required to manage resources and processes within that Azure subscription.</p> <p>You need to set up the custom role yourself if you launch a Connector from the marketplace, if you manually install the Connector, or if you <a href="#">add more Azure credentials to a Connector</a>.</p> <p>You also need to ensure that the policy is up to date as new permissions are added in subsequent releases.</p>	<a href="#">Azure permissions for the Connector</a>

### Backup and recovery

Goal	Description	Link
Back up Cloud Volumes ONTAP to Azure blob storage	<p>When using BlueXP backup and recovery to back up Cloud Volumes ONTAP, you need to add permissions to the Connector in the following scenarios:</p> <ul style="list-style-type: none"><li>• You want to use "Search &amp; Restore" functionality</li><li>• You want to use customer-managed encryption keys (CMEK)</li></ul>	<ul style="list-style-type: none"><li>• <a href="#">Back up Cloud Volumes ONTAP data to Azure Blob storage with Backup and Recovery</a></li></ul>
Back up on-premises ONTAP clusters to Azure blob storage	When using BlueXP backup and recovery to back up on-premises ONTAP clusters, you need to add permissions to the Connector in order to use the "Search & Restore" functionality.	<a href="#">Back up on-premises ONTAP data to Azure Blob storage with Backup and Recovery</a>

### Copy and sync

Goal	Description	Link
Deploy the data broker in Azure	The Azure user account that you use to deploy the data broker must have the required permissions.	<a href="#">Permissions required to deploy the data broker in Azure</a>

## Google Cloud permissions

BlueXP requires Google Cloud permissions for the Connector and for individual services.

## Connectors

Goal	Description	Link
Deploy the Connector from BlueXP	The Google Cloud user who deploys a Connector from BlueXP needs specific permissions to deploy the Connector in Google Cloud.	<a href="#">Set up permissions to create the Connector</a>
Provide permissions for the Connector	<p>The service account for the Connector VM instance must have specific permissions for day-to-day operations. You need to associate the service account with the Connector during deployment.</p> <p>You also need to ensure that the policy is up to date as new permissions are added in subsequent releases.</p>	<a href="#">Set up permissions for the Connector</a>

## Backup and recovery

Goal	Description	Link
Back up Cloud Volumes ONTAP to Google Cloud	<p>When using BlueXP backup and recovery to back up Cloud Volumes ONTAP, you need to add permissions to the Connector in the following scenarios:</p> <ul style="list-style-type: none"><li>• You want to use "Search &amp; Restore" functionality</li><li>• You want to use customer-managed encryption keys (CMEK)</li></ul>	<ul style="list-style-type: none"><li>• <a href="#">Back up Cloud Volumes ONTAP data to Google Cloud Storage with Backup and Recovery</a></li><li>• <a href="https://docs.netapp.com/us-en/bluexp-backup-recovery/prev-ontap-backup-cvo-gcp.html">https://docs.netapp.com/us-en/bluexp-backup-recovery/prev-ontap-backup-cvo-gcp.html</a>[Permissions for CMEKs^]</li></ul>
Back up on-premises ONTAP clusters to Google Cloud	When using BlueXP backup and recovery to back up on-premises ONTAP clusters, you need to add permissions to the Connector in order to use the "Search & Restore" functionality.	<a href="#">Permissions for Search &amp; Restore functionality</a>

## Cloud Volumes Service for Google Cloud

Goal	Description	Link
Discover Cloud Volumes Service for Google Cloud	BlueXP needs access to the Cloud Volumes Service API and the right permissions through a Google Cloud service account.	<a href="#">Set up a service account</a>

## Copy and sync

Goal	Description	Link
Deploy the data broker in Google Cloud	Ensure that the Google Cloud user who deploys the data broker has the required permissions.	<a href="#">Permissions required to deploy the data broker in Google Cloud</a>
Enable Google Cloud access for a manually installed data broker	If you plan to use the data broker with a sync relationship that includes a Google Cloud Storage bucket, then you should prepare the Linux host for Google Cloud access. When you install the data broker, you'll need to provide a key for a service account that has specific permissions.	<a href="#">Enabling access to Google Cloud</a>

## StorageGRID permissions

BlueXP requires StorageGRID permissions for two services.

### Backup and recovery

Goal	Description	Link
Back up on-premises ONTAP clusters to StorageGRID	When you prepare StorageGRID as a backup target for ONTAP clusters, BlueXP backup and recovery prompts you to enter an access key and secret for an IAM user that has specific permissions.	<a href="#">Prepare StorageGRID as your backup target</a>

### Tiering

Goal	Description	Link
Tier on-premises ONTAP clusters to StorageGRID	When you set up BlueXP tiering to StorageGRID, you need to provide BlueXP tiering with an S3 access key and secret key. BlueXP tiering uses the keys to access your buckets.	<a href="#">Prepare tiering to StorageGRID</a>

## AWS permissions for the Connector

When BlueXP launches the Connector instance in AWS, it attaches a policy to the instance that provides the Connector with permissions to manage resources and processes within that AWS account. The Connector uses the permissions to make API calls to several AWS services, including EC2, S3, CloudFormation, IAM, the Key Management Service (KMS), and more.

### IAM policies

The IAM policies available below provide the permissions that a Connector needs to manage resources and processes within your public cloud environment based on your AWS region.

Note the following:

- If you create a Connector in a standard AWS region directly from BlueXP, BlueXP automatically applies policies to the Connector.
- You need to set up the policies yourself if you deploy the Connector from the AWS Marketplace, if you manually install the Connector on a Linux host, or if you want to add additional AWS credentials to BlueXP.

- In either case, you need to ensure that the policies are up to date as new permissions are added in subsequent releases. If new permissions are required, they will be listed in the release notes.
- If needed, you can restrict the IAM policies by using the IAM `Condition` element. [AWS documentation: Condition element](#)
- To view step-by-step instructions for using these policies, refer to the following pages:
  - [Set up permissions for an AWS Marketplace deployment](#)
  - [Set up permissions for on-premises deployments](#)
  - [Set up permissions for restricted mode](#)
  - [Set up permissions for private mode](#)

Select your region to view the required policies:

## Standard regions

For standard regions, the permissions are spread across two policies. Two policies are required due to a maximum character size limit for managed policies in AWS.



## Policy #1

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceStatus",
        "ec2:RunInstances",
        "ec2:ModifyInstanceAttribute",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeRouteTables",
        "ec2:DescribeImages",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2:DescribeVolumes",
        "ec2:ModifyVolumeAttribute",
        "ec2:CreateSecurityGroup",
        "ec2:DescribeSecurityGroups",
        "ec2:RevokeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:RevokeSecurityGroupIngress",
        "ec2:CreateNetworkInterface",
        "ec2:DescribeNetworkInterfaces",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcs",
        "ec2:DescribeDhcpOptions",
        "ec2:CreateSnapshot",
        "ec2:DescribeSnapshots",
        "ec2:GetConsoleOutput",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeRegions",
        "ec2:DescribeTags",
        "ec2:AssociateIamInstanceProfile",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:DisassociateIamInstanceProfile",
        "ec2:CreatePlacementGroup",
        "ec2:DescribeReservedInstancesOfferings",
        "ec2:AssignPrivateIpAddresses",
        "ec2:CreateRoute",
        "ec2:DescribeVpcs",
        "ec2:ReplaceRoute",
```

```
"ec2:UnassignPrivateIpAddresses",
"ec2:DeleteSecurityGroup",
"ec2:DeleteNetworkInterface",
"ec2:DeleteSnapshot",
"ec2:DeleteTags",
"ec2:DeleteRoute",
"ec2:DeletePlacementGroup",
"ec2:DescribePlacementGroups",
"ec2:DescribeVolumesModifications",
"ec2:ModifyVolume",
"cloudformation:CreateStack",
"cloudformation:DescribeStacks",
"cloudformation:DescribeStackEvents",
"cloudformation:ValidateTemplate",
"cloudformation>DeleteStack",
"iam:PassRole",
"iam:CreateRole",
"iam:PutRolePolicy",
"iam:CreateInstanceProfile",
"iam:AddRoleToInstanceProfile",
"iam:RemoveRoleFromInstanceProfile",
"iam:ListInstanceProfiles",
"iam>DeleteRole",
"iam>DeleteRolePolicy",
"iam>DeleteInstanceProfile",
"iam:GetRolePolicy",
"iam:GetRole",
"sts:DecodeAuthorizationMessage",
"sts:AssumeRole",
"s3:GetBucketTagging",
"s3:GetBucketLocation",
"s3:ListBucket",
"s3:CreateBucket",
"s3:GetLifecycleConfiguration",
"s3:ListBucketVersions",
"s3:GetBucketPolicyStatus",
"s3:GetBucketPublicAccessBlock",
"s3:GetBucketPolicy",
"s3:GetBucketAcl",
"s3:PutObjectTagging",
"s3:GetObjectTagging",
"s3>DeleteObject",
"s3>DeleteObjectVersion",
"s3:PutObject",
"s3:ListAllMyBuckets",
"s3:GetObject",
```

```

        "s3:GetEncryptionConfiguration",
        "kms:List*",
        "kms:ReEncrypt*",
        "kms:Describe*",
        "kms:CreateGrant",
        "fsx:Describe*",
        "fsx:List*",
        "kms:GenerateDataKeyWithoutPlaintext"
    ],
    "Resource": "*",
    "Effect": "Allow",
    "Sid": "cvoServicePolicy"
},
{
    "Action": [
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceStatus",
        "ec2:RunInstances",
        "ec2:TerminateInstances",
        "ec2:DescribeInstanceAttribute",
        "ec2:DescribeImages",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2:CreateSecurityGroup",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcs",
        "ec2:DescribeRegions",
        "cloudformation:CreateStack",
        "cloudformation>DeleteStack",
        "cloudformation:DescribeStacks",
        "kms:List*",
        "kms:Describe*",
        "ec2:DescribeVpcEndpoints",
        "kms:ListAliases",
        "athena:StartQueryExecution",
        "athena:GetQueryResults",
        "athena:GetQueryExecution",
        "glue:GetDatabase",
        "glue:GetTable",
        "glue:CreateTable",
        "glue:CreateDatabase",
        "glue:GetPartitions",
        "glue:BatchCreatePartition",
        "glue:BatchDeletePartition"
    ]
}

```

```

    ],
    "Resource": "*",
    "Effect": "Allow",
    "Sid": "backupPolicy"
  },
  {
    "Action": [
      "s3:GetBucketLocation",
      "s3:ListAllMyBuckets",
      "s3:ListBucket",
      "s3:CreateBucket",
      "s3:GetLifecycleConfiguration",
      "s3:PutLifecycleConfiguration",
      "s3:PutBucketTagging",
      "s3:ListBucketVersions",
      "s3:GetBucketAcl",
      "s3:PutBucketPublicAccessBlock",
      "s3:GetObject",
      "s3:PutEncryptionConfiguration",
      "s3>DeleteObject",
      "s3>DeleteObjectVersion",
      "s3:ListBucketMultipartUploads",
      "s3:PutObject",
      "s3:PutBucketAcl",
      "s3:AbortMultipartUpload",
      "s3:ListMultipartUploadParts",
      "s3>DeleteBucket",
      "s3:GetObjectVersionTagging",
      "s3:GetObjectVersionAcl",
      "s3:GetObjectRetention",
      "s3:GetObjectTagging",
      "s3:GetObjectVersion",
      "s3:PutObjectVersionTagging",
      "s3:PutObjectRetention",
      "s3>DeleteObjectTagging",
      "s3>DeleteObjectVersionTagging",
      "s3:GetBucketObjectLockConfiguration",
      "s3:GetBucketVersioning",
      "s3:PutBucketObjectLockConfiguration",
      "s3:PutBucketVersioning",
      "s3:BypassGovernanceRetention",
      "s3:PutBucketPolicy",
      "s3:PutBucketOwnershipControls"
    ],
    "Resource": [
      "arn:aws:s3:::netapp-backup-*"
    ]
  }
]

```

```

    ],
    "Effect": "Allow",
    "Sid": "backupS3Policy"
  },
  {
    "Action": [
      "s3:CreateBucket",
      "s3:GetLifecycleConfiguration",
      "s3:PutLifecycleConfiguration",
      "s3:PutBucketTagging",
      "s3:ListBucketVersions",
      "s3:GetBucketPolicyStatus",
      "s3:GetBucketPublicAccessBlock",
      "s3:GetBucketAcl",
      "s3:GetBucketPolicy",
      "s3:PutBucketPublicAccessBlock",
      "s3>DeleteBucket"
    ],
    "Resource": [
      "arn:aws:s3:::fabric-pool*"
    ],
    "Effect": "Allow",
    "Sid": "fabricPools3Policy"
  },
  {
    "Action": [
      "ec2:DescribeRegions"
    ],
    "Resource": "*",
    "Effect": "Allow",
    "Sid": "fabricPoolPolicy"
  },
  {
    "Condition": {
      "StringLike": {
        "ec2:ResourceTag/netapp-adc-manager": "*"
      }
    },
    "Action": [
      "ec2:StartInstances",
      "ec2:StopInstances",
      "ec2:TerminateInstances"
    ],
    "Resource": [
      "arn:aws:ec2:*:*:instance/*"
    ],
  },

```

```

    "Effect": "Allow"
  },
  {
    "Condition": {
      "StringLike": {
        "ec2:ResourceTag/WorkingEnvironment": "*"
      }
    },
    "Action": [
      "ec2:StartInstances",
      "ec2:TerminateInstances",
      "ec2:AttachVolume",
      "ec2:DetachVolume",
      "ec2:StopInstances",
      "ec2>DeleteVolume"
    ],
    "Resource": [
      "arn:aws:ec2:*:*:instance/*"
    ],
    "Effect": "Allow"
  },
  {
    "Action": [
      "ec2:AttachVolume",
      "ec2:DetachVolume"
    ],
    "Resource": [
      "arn:aws:ec2:*:*:volume/*"
    ],
    "Effect": "Allow"
  },
  {
    "Condition": {
      "StringLike": {
        "ec2:ResourceTag/WorkingEnvironment": "*"
      }
    },
    "Action": [
      "ec2>DeleteVolume"
    ],
    "Resource": [
      "arn:aws:ec2:*:*:volume/*"
    ],
    "Effect": "Allow"
  }
]

```

```
}
```

## Policy #2

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "ec2:CreateTags",
        "ec2>DeleteTags",
        "ec2:DescribeTags",
        "tag:getResources",
        "tag:getTagKeys",
        "tag:getTagValues",
        "tag:TagResources",
        "tag:UntagResources"
      ],
      "Resource": "*",
      "Effect": "Allow",
      "Sid": "tagServicePolicy"
    }
  ]
}
```

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:ListInstanceProfiles",
        "iam:CreateRole",
        "iam>DeleteRole",
        "iam:PutRolePolicy",
        "iam:CreateInstanceProfile",
        "iam>DeleteRolePolicy",
        "iam:AddRoleToInstanceProfile",
        "iam:RemoveRoleFromInstanceProfile",
        "iam>DeleteInstanceProfile",
        "ec2:ModifyVolumeAttribute",
        "sts:DecodeAuthorizationMessage",
        "ec2:DescribeImages",
        "ec2:DescribeRouteTables",
        "ec2:DescribeInstances",
        "iam:PassRole",
        "ec2:DescribeInstanceStatus",
        "ec2:RunInstances",
        "ec2:ModifyInstanceAttribute",
        "ec2:CreateTags",
        "ec2:CreateVolume",
        "ec2:DescribeVolumes",
        "ec2>DeleteVolume",
        "ec2:CreateSecurityGroup",
        "ec2>DeleteSecurityGroup",
        "ec2:DescribeSecurityGroups",
        "ec2:RevokeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:RevokeSecurityGroupIngress",
        "ec2:CreateNetworkInterface",
        "ec2:DescribeNetworkInterfaces",
        "ec2>DeleteNetworkInterface",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcs",
        "ec2:DescribeDhcpOptions",
        "ec2:CreateSnapshot",
        "ec2>DeleteSnapshot",

```



```

        "ec2:DescribeSnapshots",
        "ec2:StopInstances",
        "ec2:GetConsoleOutput",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeRegions",
        "ec2>DeleteTags",
        "ec2:DescribeTags",
        "cloudformation:CreateStack",
        "cloudformation>DeleteStack",
        "cloudformation:DescribeStacks",
        "cloudformation:DescribeStackEvents",
        "cloudformation:ValidateTemplate",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:ListAllMyBuckets",
        "s3:GetBucketTagging",
        "s3:GetBucketLocation",
        "s3:CreateBucket",
        "s3:GetBucketPolicyStatus",
        "s3:GetBucketPublicAccessBlock",
        "s3:GetBucketAcl",
        "s3:GetBucketPolicy",
        "kms:List*",
        "kms:ReEncrypt*",
        "kms:Describe*",
        "kms:CreateGrant",
        "ec2:AssociateIamInstanceProfile",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:DisassociateIamInstanceProfile",
        "ec2:DescribeInstanceAttribute",
        "ec2:CreatePlacementGroup",
        "ec2>DeletePlacementGroup"
    ],
    "Resource": "*"
},
{
    "Sid": "fabricPoolPolicy",
    "Effect": "Allow",
    "Action": [
        "s3:DeleteBucket",
        "s3:GetLifecycleConfiguration",
        "s3:PutLifecycleConfiguration",
        "s3:PutBucketTagging",
        "s3:ListBucketVersions",
        "s3:GetBucketPolicyStatus",
        "s3:GetBucketPublicAccessBlock",

```

```

        "s3:GetBucketAcl",
        "s3:GetBucketPolicy",
        "s3:PutBucketPublicAccessBlock"
    ],
    "Resource": [
        "arn:aws-us-gov:s3:::fabric-pool*"
    ]
},
{
    "Sid": "backupPolicy",
    "Effect": "Allow",
    "Action": [
        "s3:DeleteBucket",
        "s3:GetLifecycleConfiguration",
        "s3:PutLifecycleConfiguration",
        "s3:PutBucketTagging",
        "s3:ListBucketVersions",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:ListAllMyBuckets",
        "s3:GetBucketTagging",
        "s3:GetBucketLocation",
        "s3:GetBucketPolicyStatus",
        "s3:GetBucketPublicAccessBlock",
        "s3:GetBucketAcl",
        "s3:GetBucketPolicy",
        "s3:PutBucketPublicAccessBlock"
    ],
    "Resource": [
        "arn:aws-us-gov:s3:::netapp-backup-*"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "ec2:StartInstances",
        "ec2:TerminateInstances",
        "ec2:AttachVolume",
        "ec2:DetachVolume"
    ],
    "Condition": {
        "StringLike": {
            "ec2:ResourceTag/WorkingEnvironment": "*"
        }
    },
    "Resource": [

```

```
        "arn:aws-us-gov:ec2:*:*:instance/*"
    ],
    },
    {
        "Effect": "Allow",
        "Action": [
            "ec2:AttachVolume",
            "ec2:DetachVolume"
        ],
        "Resource": [
            "arn:aws-us-gov:ec2:*:*:volume/*"
        ]
    }
]
}
```

```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeInstances",
      "ec2:DescribeInstanceStatus",
      "ec2:RunInstances",
      "ec2:ModifyInstanceAttribute",
      "ec2:DescribeRouteTables",
      "ec2:DescribeImages",
      "ec2:CreateTags",
      "ec2:CreateVolume",
      "ec2:DescribeVolumes",
      "ec2:ModifyVolumeAttribute",
      "ec2>DeleteVolume",
      "ec2:CreateSecurityGroup",
      "ec2>DeleteSecurityGroup",
      "ec2:DescribeSecurityGroups",
      "ec2:RevokeSecurityGroupEgress",
      "ec2:RevokeSecurityGroupIngress",
      "ec2:AuthorizeSecurityGroupEgress",
      "ec2:AuthorizeSecurityGroupIngress",
      "ec2:CreateNetworkInterface",
      "ec2:DescribeNetworkInterfaces",
      "ec2>DeleteNetworkInterface",
      "ec2:ModifyNetworkInterfaceAttribute",
      "ec2:DescribeSubnets",
      "ec2:DescribeVpcs",
      "ec2:DescribeDhcpOptions",
      "ec2:CreateSnapshot",
      "ec2>DeleteSnapshot",
      "ec2:DescribeSnapshots",
      "ec2:GetConsoleOutput",
      "ec2:DescribeKeyPairs",
      "ec2:DescribeRegions",
      "ec2>DeleteTags",
      "ec2:DescribeTags",
      "cloudformation:CreateStack",
      "cloudformation>DeleteStack",
      "cloudformation:DescribeStacks",
      "cloudformation:DescribeStackEvents",
      "cloudformation:ValidateTemplate",
      "iam:PassRole",

```

```

        "iam:CreateRole",
        "iam>DeleteRole",
        "iam:PutRolePolicy",
        "iam:CreateInstanceProfile",
        "iam>DeleteRolePolicy",
        "iam:AddRoleToInstanceProfile",
        "iam:RemoveRoleFromInstanceProfile",
        "iam>DeleteInstanceProfile",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:GetBucketTagging",
        "s3:GetBucketLocation",
        "s3:ListAllMyBuckets",
        "kms:List*",
        "kms:Describe*",
        "ec2:AssociateIamInstanceProfile",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:DisassociateIamInstanceProfile",
        "ec2:DescribeInstanceAttribute",
        "ec2:CreatePlacementGroup",
        "ec2>DeletePlacementGroup",
        "iam:ListInstanceProfiles"
    ],
    "Resource": "*"
},
{
    "Sid": "fabricPoolPolicy",
    "Effect": "Allow",
    "Action": [
        "s3>DeleteBucket",
        "s3:GetLifecycleConfiguration",
        "s3:PutLifecycleConfiguration",
        "s3:PutBucketTagging",
        "s3:ListBucketVersions"
    ],
    "Resource": [
        "arn:aws-iso-b:s3:::fabric-pool*"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances",
        "ec2:AttachVolume",

```

```

        "ec2:DetachVolume"
    ],
    "Condition": {
        "StringLike": {
            "ec2:ResourceTag/WorkingEnvironment": "*"
        }
    },
    "Resource": [
        "arn:aws-iso-b:ec2:*:*:instance/*"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "ec2:AttachVolume",
        "ec2:DetachVolume"
    ],
    "Resource": [
        "arn:aws-iso-b:ec2:*:*:volume/*"
    ]
}
]
}

```

## Top Secret regions

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeInstances",
      "ec2:DescribeInstanceStatus",
      "ec2:RunInstances",
      "ec2:ModifyInstanceAttribute",
      "ec2:DescribeRouteTables",
      "ec2:DescribeImages",
      "ec2:CreateTags",
      "ec2:CreateVolume",
      "ec2:DescribeVolumes",
      "ec2:ModifyVolumeAttribute",
      "ec2>DeleteVolume",
      "ec2:CreateSecurityGroup",
      "ec2>DeleteSecurityGroup",
      "ec2:DescribeSecurityGroups",
      "ec2:RevokeSecurityGroupEgress",
      "ec2:RevokeSecurityGroupIngress",
      "ec2:AuthorizeSecurityGroupEgress",
      "ec2:AuthorizeSecurityGroupIngress",
      "ec2:CreateNetworkInterface",
      "ec2:DescribeNetworkInterfaces",
      "ec2>DeleteNetworkInterface",
      "ec2:ModifyNetworkInterfaceAttribute",
      "ec2:DescribeSubnets",
      "ec2:DescribeVpcs",
      "ec2:DescribeDhcpOptions",
      "ec2:CreateSnapshot",
      "ec2>DeleteSnapshot",
      "ec2:DescribeSnapshots",
      "ec2:GetConsoleOutput",
      "ec2:DescribeKeyPairs",
      "ec2:DescribeRegions",
      "ec2>DeleteTags",
      "ec2:DescribeTags",
      "cloudformation:CreateStack",
      "cloudformation>DeleteStack",
      "cloudformation:DescribeStacks",
      "cloudformation:DescribeStackEvents",
      "cloudformation:ValidateTemplate",
      "iam:PassRole",
```

```

        "iam:CreateRole",
        "iam>DeleteRole",
        "iam:PutRolePolicy",
        "iam:CreateInstanceProfile",
        "iam>DeleteRolePolicy",
        "iam:AddRoleToInstanceProfile",
        "iam:RemoveRoleFromInstanceProfile",
        "iam>DeleteInstanceProfile",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:GetBucketTagging",
        "s3:GetBucketLocation",
        "s3:ListAllMyBuckets",
        "kms:List*",
        "kms:Describe*",
        "ec2:AssociateIamInstanceProfile",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:DisassociateIamInstanceProfile",
        "ec2:DescribeInstanceAttribute",
        "ec2:CreatePlacementGroup",
        "ec2>DeletePlacementGroup",
        "iam:ListInstanceProfiles"
    ],
    "Resource": "*"
},
{
    "Sid": "fabricPoolPolicy",
    "Effect": "Allow",
    "Action": [
        "s3>DeleteBucket",
        "s3:GetLifecycleConfiguration",
        "s3:PutLifecycleConfiguration",
        "s3:PutBucketTagging",
        "s3:ListBucketVersions"
    ],
    "Resource": [
        "arn:aws-iso:s3:::fabric-pool*"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances",
        "ec2:AttachVolume",

```



```

        "ec2:DetachVolume"
    ],
    "Condition": {
        "StringLike": {
            "ec2:ResourceTag/WorkingEnvironment": "*"
        }
    },
    "Resource": [
        "arn:aws-iso:ec2:*:*:instance/*"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "ec2:AttachVolume",
        "ec2:DetachVolume"
    ],
    "Resource": [
        "arn:aws-iso:ec2:*:*:volume/*"
    ]
}
]
}

```

## How the AWS permissions are used

The following sections describe how the permissions are used for each BlueXP service. This information can be helpful if your corporate policies dictate that permissions are only provided as needed.

### Amazon FSx for ONTAP

The Connector makes the following API requests to manage an Amazon FSx for ONTAP file system:

- ec2:DescribeInstances
- ec2:DescribeInstanceStatus
- ec2:DescribeInstanceAttribute
- ec2:DescribeRouteTables
- ec2:DescribeImages
- ec2:CreateTags
- ec2:DescribeVolumes
- ec2:DescribeSecurityGroups
- ec2:DescribeNetworkInterfaces
- ec2:DescribeSubnets
- ec2:DescribeVpcs

- ec2:DescribeDhcpOptions
- ec2:DescribeSnapshots
- ec2:DescribeKeyPairs
- ec2:DescribeRegions
- ec2:DescribeTags
- ec2:DescribeIamInstanceProfileAssociations
- ec2:DescribeReservedInstancesOfferings
- ec2:DescribeVpcEndpoints
- ec2:DescribeVpcs
- ec2:DescribeVolumesModifications
- ec2:DescribePlacementGroups
- kms:List\*
- kms:Describe\*
- kms:CreateGrant
- kms:ListAliases
- fsx:Describe\*
- fsx:List\*

#### **Amazon S3 bucket discovery**

The Connector makes the following API request to discover Amazon S3 buckets:

s3:GetEncryptionConfiguration

#### **Backup and recovery**

The Connector makes the following API requests to manage backups in Amazon S3:

- s3:GetBucketLocation
- s3:ListAllMyBuckets
- s3:ListBucket
- s3:CreateBucket
- s3:GetLifecycleConfiguration
- s3:PutLifecycleConfiguration
- s3:PutBucketTagging
- s3:ListBucketVersions
- s3:GetBucketAcl
- s3:PutBucketPublicAccessBlock
- kms:List\*
- kms:Describe\*
- s3:GetObject

- ec2:DescribeVpcEndpoints
- kms:ListAliases
- s3:PutEncryptionConfiguration

The Connector makes the following API requests when you use the Search & Restore method to restore volumes and files:

- s3:CreateBucket
- s3>DeleteObject
- s3>DeleteObjectVersion
- s3:GetBucketAcl
- s3:ListBucket
- s3:ListBucketVersions
- s3:ListBucketMultipartUploads
- s3:PutObject
- s3:PutBucketAcl
- s3:PutLifecycleConfiguration
- s3:PutBucketPublicAccessBlock
- s3:AbortMultipartUpload
- s3:ListMultipartUploadParts
- athena:StartQueryExecution
- athena:GetQueryResults
- athena:GetQueryExecution
- athena:StopQueryExecution
- glue:CreateDatabase
- glue:CreateTable
- glue:BatchDeletePartition

The Connector makes the following API requests when you use DataLock and Ransomware protection for your volume backups:

- 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:ListBucketByTags
- 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

The Connector makes the following API requests if you use a different AWS account for your Cloud Volumes ONTAP backups than you're using for the source volumes:

- s3:PutBucketPolicy
- s3:PutBucketOwnershipControls

### **Classification**

The Connector makes the following API requests to deploy the BlueXP classification instance:

- ec2:DescribeInstances
- ec2:DescribeInstanceStatus
- ec2:RunInstances
- ec2:TerminateInstances
- ec2:CreateTags
- ec2:CreateVolume
- ec2:AttachVolume
- ec2:CreateSecurityGroup
- ec2>DeleteSecurityGroup
- ec2:DescribeSecurityGroups
- ec2:CreateNetworkInterface
- ec2:DescribeNetworkInterfaces

- ec2:DeleteNetworkInterface
- ec2:DescribeSubnets
- ec2:DescribeVpcs
- ec2:CreateSnapshot
- ec2:DescribeRegions
- cloudformation:CreateStack
- cloudformation>DeleteStack
- cloudformation:DescribeStacks
- cloudformation:DescribeStackEvents
- iam:AddRoleToInstanceProfile
- ec2:AssociateIamInstanceProfile
- ec2:DescribeIamInstanceProfileAssociations

The Connector makes the following API requests to scan S3 buckets when you use BlueXP classification:

- iam:AddRoleToInstanceProfile
- ec2:AssociateIamInstanceProfile
- ec2:DescribeIamInstanceProfileAssociations
- s3:GetBucketTagging
- s3:GetBucketLocation
- s3:ListAllMyBuckets
- s3:ListBucket
- s3:GetBucketPolicyStatus
- s3:GetBucketPolicy
- s3:GetBucketAcl
- s3:GetObject
- iam:GetRole
- s3>DeleteObject
- s3>DeleteObjectVersion
- s3:PutObject
- sts:AssumeRole

### **Cloud Volumes ONTAP**

The Connector makes the following API requests to deploy and manage Cloud Volumes ONTAP in AWS.

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Create and manage IAM roles and instance profiles for Cloud Volumes ONTAP instances	iam:ListInstanceProfiles	Yes	Yes	No
	iam:CreateRole	Yes	No	No
	iam>DeleteRole	No	Yes	Yes
	iam:PutRolePolicy	Yes	No	No
	iam:CreateInstanceProfile	Yes	No	No
	iam>DeleteRolePolicy	No	Yes	Yes
	iam:AddRoleToInstanceProfile	Yes	No	No
	iam:RemoveRoleFromInstanceProfile	No	Yes	Yes
	iam>DeleteInstanceProfile	No	Yes	Yes
	iam:PassRole	Yes	No	No
	ec2:AssociateIamInstanceProfile	Yes	Yes	No
	ec2:DescribeIamInstanceProfileAssociations	Yes	Yes	No
	ec2:DisassociateIamInstanceProfile	No	Yes	No
Decode authorization status messages	sts:DecodeAuthorizationMessage	Yes	Yes	No
Describe the specified images (AMIs) available to the account	ec2:DescribeImages	Yes	Yes	No
Describe the route tables in a VPC (required for HA pairs only)	ec2:DescribeRouteTables	Yes	No	No

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Stop, start, and monitor instances	ec2:StartInstances	Yes	Yes	No
	ec2:StopInstances	Yes	Yes	No
	ec2:DescribeInstances	Yes	Yes	No
	ec2:DescribeInstanceStatus	Yes	Yes	No
	ec2:RunInstances	Yes	No	No
	ec2:TerminateInstances	No	No	Yes
	ec2:ModifyInstanceAttribute	No	Yes	No
Verify that enhanced networking is enabled for supported instance types	ec2:DescribeInstanceAttribute	No	Yes	No
Tag resources with the "WorkingEnvironment" and "WorkingEnvironmentId" tags which are used for maintenance and cost allocation	ec2:CreateTags	Yes	Yes	No
Manage EBS volumes that Cloud Volumes ONTAP uses as back-end storage	ec2:CreateVolume	Yes	Yes	No
	ec2:DescribeVolumes	Yes	Yes	Yes
	ec2:ModifyVolumeAttribute	No	Yes	Yes
	ec2:AttachVolume	Yes	Yes	No
	ec2>DeleteVolume	No	Yes	Yes
	ec2:DetachVolume	No	Yes	Yes

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Create and manage security groups for Cloud Volumes ONTAP	ec2:CreateSecurityGroup	Yes	No	No
	ec2>DeleteSecurityGroup	No	Yes	Yes
	ec2:DescribeSecurityGroups	Yes	Yes	Yes
	ec2:RevokeSecurityGroupEgress	Yes	No	No
	ec2:AuthorizeSecurityGroupEgress	Yes	No	No
	ec2:AuthorizeSecurityGroupIngress	Yes	No	No
	ec2:RevokeSecurityGroupIngress	Yes	Yes	No
Create and manage network interfaces for Cloud Volumes ONTAP in the target subnet	ec2:CreateNetworkInterface	Yes	No	No
	ec2:DescribeNetworkInterfaces	Yes	Yes	No
	ec2>DeleteNetworkInterface	No	Yes	Yes
	ec2:ModifyNetworkInterfaceAttribute	No	Yes	No
Get the list of destination subnets and security groups	ec2:DescribeSubnets	Yes	Yes	No
	ec2:DescribeVpcs	Yes	Yes	No
Get DNS servers and the default domain name for Cloud Volumes ONTAP instances	ec2:DescribeDhcpOptions	Yes	No	No
Take snapshots of EBS volumes for Cloud Volumes ONTAP	ec2:CreateSnapshot	Yes	Yes	No
	ec2>DeleteSnapshot	No	Yes	Yes
	ec2:DescribeSnapshots	No	Yes	No
Capture the Cloud Volumes ONTAP console, which is attached to AutoSupport messages	ec2:GetConsoleOutput	Yes	Yes	No



Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Get the list of available key pairs	ec2:DescribeKeyPairs	Yes	No	No
Get the list of available AWS regions	ec2:DescribeRegions	Yes	Yes	No
Manage tags for resources associated with Cloud Volumes ONTAP instances	ec2:DeleteTags	No	Yes	Yes
	ec2:DescribeTags	No	Yes	No
Create and manage stacks for AWS CloudFormation templates	cloudformation:CreateStack	Yes	No	No
	cloudformation:DeleteStack	Yes	No	No
	cloudformation:DescribeStacks	Yes	Yes	No
	cloudformation:DescribeStackEvents	Yes	No	No
	cloudformation:ValidateTemplate	Yes	No	No

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Create and manage an S3 bucket that a Cloud Volumes ONTAP system uses as a capacity tier for data tiering	s3:CreateBucket	Yes	Yes	No
	s3>DeleteBucket	No	Yes	Yes
	s3:GetLifecycleConfiguration	No	Yes	No
	s3:PutLifecycleConfiguration	No	Yes	No
	s3:PutBucketTagging	No	Yes	No
	s3:ListBucketVersions	No	Yes	No
	s3:GetBucketPolicyStatus	No	Yes	No
	s3:GetBucketPublicAccessBlock	No	Yes	No
	s3:GetBucketAcl	No	Yes	No
	s3:GetBucketPolicy	No	Yes	No
	s3:PutBucketPublicAccessBlock	No	Yes	No
	s3:GetBucketTagging	No	Yes	No
	s3:GetBucketLocation	No	Yes	No
	s3:ListAllMyBuckets	No	No	No
	s3:ListBucket	No	Yes	No
Enable data encryption of Cloud Volumes ONTAP using the AWS Key Management Service (KMS)	kms:List*	Yes	Yes	No
	kms:ReEncrypt*	Yes	No	No
	kms:Describe*	Yes	Yes	No
	kms:CreateGrant	Yes	Yes	No
	kms:GenerateDataKeyWithoutPlaintext	Yes	Yes	No
Create and manage an AWS spread placement group for two HA nodes and the mediator in a single AWS Availability Zone	ec2:CreatePlacementGroup	Yes	No	No
	ec2>DeletePlacementGroup	No	Yes	Yes

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Create reports	fsx:Describe*	No	Yes	No
	fsx:List*	No	Yes	No
Create and manage aggregates that support the Amazon EBS Elastic Volumes feature	ec2:DescribeVolumesModifications	No	Yes	No
	ec2:ModifyVolume	No	Yes	No
Check whether the Availability Zone is an AWS Local Zone and validates that all deployment parameters are compatible	ec2:DescribeAvailabilityZones	Yes	No	Yes

### Change log

As permissions are added and removed, we'll note them in the sections below.

#### 9 September 2024

Permissions were removed from policy #2 for standard regions because BlueXP no longer supports BlueXP edge caching and discovery and management of Kubernetes clusters.

## View the permissions that were removed from the policy

```
{
  "Action": [
    "ec2:DescribeRegions",
    "eks:ListClusters",
    "eks:DescribeCluster",
    "iam:GetInstanceProfile"
  ],
  "Resource": "*",
  "Effect": "Allow",
  "Sid": "K8sServicePolicy"
},
{
  "Action": [
    "cloudformation:DescribeStacks",
    "cloudwatch:GetMetricStatistics",
    "cloudformation:ListStacks"
  ],
  "Resource": "*",
  "Effect": "Allow",
  "Sid": "GFCservicePolicy"
},
{
  "Condition": {
    "StringLike": {
      "ec2:ResourceTag/GFCInstance": "*"
    }
  },
  "Action": [
    "ec2:StartInstances",
    "ec2:TerminateInstances",
    "ec2:AttachVolume",
    "ec2:DetachVolume"
  ],
  "Resource": [
    "arn:aws:ec2:*:*:instance/*"
  ],
  "Effect": "Allow"
},
```

9 May 2024

The following permissions is now required for Cloud Volumes ONTAP:

ec2:DescribeAvailabilityZones

**6 June 2023**

The following permission is now required for Cloud Volumes ONTAP:

kms:GenerateDataKeyWithoutPlaintext

**14 February 2023**

The following permission is now required for BlueXP tiering:

ec2:DescribeVpcEndpoints

## Azure permissions for the Connector

When BlueXP launches the Connector VM in Azure, it attaches a custom role to the VM that provides the Connector with permissions to manage resources and processes within that Azure subscription. The Connector uses the permissions to make API calls to several Azure services.

Whether or not you need to create this custom role for the Connector depends on how you deployed it.

### Deploying from BlueXP

When you use BlueXP to deploy the Connector virtual machine in Azure, it enables a [system-assigned managed identity](#) on the virtual machine, creates a custom role, and assigns it to the virtual machine. The role provides BlueXP with the permissions required to manage resources and processes within that Azure subscription. The role's permissions are kept up-to-date when the Connector is upgraded. You don't need to create this role for the Connector or manage updates.

### Deploying manually or from Azure marketplace

When you deploy the Connector from the Azure Marketplace or if you manually install the Connector on a Linux host, then you need to set up the custom role yourself and maintain its permissions with any changes.

You'll need to ensure that the role is up to date as new permissions are added in subsequent releases. If new permissions are required, they will be listed in the release notes.

- To view step-by-step instructions for using these policies, refer to the following pages:
  - [Set up permissions for an Azure Marketplace deployment](#)
  - [Set up permissions for on-premises deployments](#)
  - [Set up permissions for restricted mode](#)
  - [Set up permissions for private mode](#)

```
{
  "Name": "BlueXP Operator",
  "Actions": [
    "Microsoft.Compute/disks/delete",
    "Microsoft.Compute/disks/read",
    "Microsoft.Compute/disks/write",
    "Microsoft.Compute/locations/operations/read",
```

```
"Microsoft.Compute/locations/vmSizes/read",
"Microsoft.Resources/subscriptions/locations/read",
"Microsoft.Compute/operations/read",
"Microsoft.Compute/virtualMachines/instanceView/read",
"Microsoft.Compute/virtualMachines/powerOff/action",
"Microsoft.Compute/virtualMachines/read",
"Microsoft.Compute/virtualMachines/restart/action",
"Microsoft.Compute/virtualMachines/deallocate/action",
"Microsoft.Compute/virtualMachines/start/action",
"Microsoft.Compute/virtualMachines/vmSizes/read",
"Microsoft.Compute/virtualMachines/write",
"Microsoft.Compute/images/read",
"Microsoft.Network/locations/operationResults/read",
"Microsoft.Network/locations/operations/read",
"Microsoft.Network/networkInterfaces/read",
"Microsoft.Network/networkInterfaces/write",
"Microsoft.Network/networkInterfaces/join/action",
"Microsoft.Network/networkSecurityGroups/read",
"Microsoft.Network/networkSecurityGroups/write",
"Microsoft.Network/networkSecurityGroups/join/action",
"Microsoft.Network/virtualNetworks/read",

"Microsoft.Network/virtualNetworks/checkIpAddressAvailability/read",
"Microsoft.Network/virtualNetworks/subnets/read",
"Microsoft.Network/virtualNetworks/subnets/write",

"Microsoft.Network/virtualNetworks/subnets/virtualMachines/read",

"Microsoft.Network/virtualNetworks/virtualMachines/read",

"Microsoft.Network/virtualNetworks/subnets/join/action",
"Microsoft.Resources/deployments/operations/read",
"Microsoft.Resources/deployments/read",
"Microsoft.Resources/deployments/write",
"Microsoft.Resources/resources/read",

"Microsoft.Resources/subscriptions/operationresults/read",

"Microsoft.Resources/subscriptions/resourceGroups/delete",

"Microsoft.Resources/subscriptions/resourceGroups/read",

"Microsoft.Resources/subscriptions/resourcegroups/resources/read",

"Microsoft.Resources/subscriptions/resourceGroups/write",
"Microsoft.Storage/checknameavailability/read",
```

```
"Microsoft.Storage/operations/read",
"Microsoft.Storage/storageAccounts/listkeys/action",
"Microsoft.Storage/storageAccounts/read",
"Microsoft.Storage/storageAccounts/delete",
"Microsoft.Storage/storageAccounts/write",

"Microsoft.Storage/storageAccounts/blobServices/containers/read",

"Microsoft.Storage/storageAccounts/listAccountSas/action",
"Microsoft.Storage/usages/read",
"Microsoft.Compute/snapshots/write",
"Microsoft.Compute/snapshots/read",
"Microsoft.Compute/availabilitySets/write",
"Microsoft.Compute/availabilitySets/read",
"Microsoft.Compute/disks/beginGetAccess/action",

"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/read",

"Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/write",

"Microsoft.Network/loadBalancers/read",
"Microsoft.Network/loadBalancers/write",
"Microsoft.Network/loadBalancers/delete",

"Microsoft.Network/loadBalancers/backendAddressPools/read",

"Microsoft.Network/loadBalancers/backendAddressPools/join/action",

"Microsoft.Network/loadBalancers/loadBalancingRules/read",
"Microsoft.Network/loadBalancers/probes/read",
"Microsoft.Network/loadBalancers/probes/join/action",
"Microsoft.Authorization/locks/*",
"Microsoft.Network/routeTables/join/action",
"Microsoft.NetApp/netAppAccounts/read",
"Microsoft.NetApp/netAppAccounts/capacityPools/read",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/write",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/read",

"Microsoft.NetApp/netAppAccounts/capacityPools/volumes/delete",
"Microsoft.Network/privateEndpoints/write",

"Microsoft.Storage/storageAccounts/PrivateEndpointConnectionsApproval/action",
```

```
"Microsoft.Storage/storageAccounts/privateEndpointConnections/read",

"Microsoft.Storage/storageAccounts/managementPolicies/read",

"Microsoft.Storage/storageAccounts/managementPolicies/write",
    "Microsoft.Network/privateEndpoints/read",
    "Microsoft.Network/privateDnsZones/write",

"Microsoft.Network/privateDnsZones/virtualNetworkLinks/write",
    "Microsoft.Network/virtualNetworks/join/action",
    "Microsoft.Network/privateDnsZones/A/write",
    "Microsoft.Network/privateDnsZones/read",

"Microsoft.Network/privateDnsZones/virtualNetworkLinks/read",

"Microsoft.Resources/deployments/operationStatuses/read",
    "Microsoft.Insights/Metrics/Read",
    "Microsoft.Compute/virtualMachines/extensions/write",
    "Microsoft.Compute/virtualMachines/extensions/delete",
    "Microsoft.Compute/virtualMachines/extensions/read",
    "Microsoft.Compute/virtualMachines/delete",
    "Microsoft.Network/networkInterfaces/delete",
    "Microsoft.Network/networkSecurityGroups/delete",
    "Microsoft.Resources/deployments/delete",
    "Microsoft.Compute/diskEncryptionSets/read",
    "Microsoft.Compute/snapshots/delete",
    "Microsoft.Network/privateEndpoints/delete",
    "Microsoft.Compute/availabilitySets/delete",
    "Microsoft.KeyVault/vaults/read",
    "Microsoft.KeyVault/vaults/accessPolicies/write",
    "Microsoft.Compute/diskEncryptionSets/write",
    "Microsoft.KeyVault/vaults/deploy/action",
    "Microsoft.Compute/diskEncryptionSets/delete",
    "Microsoft.Resources/tags/read",
    "Microsoft.Resources/tags/write",
    "Microsoft.Resources/tags/delete",
    "Microsoft.Network/applicationSecurityGroups/write",
    "Microsoft.Network/applicationSecurityGroups/read",

"Microsoft.Network/applicationSecurityGroups/joinIpConfiguration/action",

"Microsoft.Network/networkSecurityGroups/securityRules/write",
    "Microsoft.Network/applicationSecurityGroups/delete",

"Microsoft.Network/networkSecurityGroups/securityRules/delete",
```



```

        "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",

"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action",
        "Microsoft.Compute/images/write",

"Microsoft.Network/loadBalancers/frontendIPConfigurations/read",
        "Microsoft.Compute/virtualMachineScaleSets/write",
        "Microsoft.Compute/virtualMachineScaleSets/read",
        "Microsoft.Compute/virtualMachineScaleSets/delete"
    ],
    "NotActions": [],
    "AssignableScopes": [],
    "Description": "BlueXP Permissions",
    "IsCustom": "true"
}

```

## How Azure permissions are used

The following sections describe how the permissions are used for each BlueXP service. This information can be helpful if your corporate policies dictate that permissions are only provided as needed.

### Azure NetApp Files

The Connector makes the following API requests when you use BlueXP classification to scan Azure NetApp Files data:

- Microsoft.NetApp/netAppAccounts/read
- Microsoft.NetApp/netAppAccounts/capacityPools/read
- Microsoft.NetApp/netAppAccounts/capacityPools/volumes/write
- Microsoft.NetApp/netAppAccounts/capacityPools/volumes/read
- Microsoft.NetApp/netAppAccounts/capacityPools/volumes/delete

### Backup and recovery

The Connector makes the following API requests for BlueXP backup and recovery:

- 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

The Connector makes the following API requests when you use the Search & Restore functionality:

- 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

### Classification

The Connector makes the following API requests when you use BlueXP classification.

Action	Used for set up?	Used for daily operations?
Microsoft.Compute/locations/operations/read	Yes	Yes
Microsoft.Compute/locations/vmSizes/read	Yes	Yes
Microsoft.Compute/operations/read	Yes	Yes
Microsoft.Compute/virtualMachines/instanceView/read	Yes	Yes
Microsoft.Compute/virtualMachines/powerOff/action	Yes	No
Microsoft.Compute/virtualMachines/read	Yes	Yes
Microsoft.Compute/virtualMachines/restart/action	Yes	No
Microsoft.Compute/virtualMachines/start/action	Yes	No
Microsoft.Compute/virtualMachines/vmSizes/read	No	Yes
Microsoft.Compute/virtualMachines/write	Yes	No
Microsoft.Compute/images/read	Yes	Yes
Microsoft.Compute/disks/delete	Yes	No
Microsoft.Compute/disks/read	Yes	Yes
Microsoft.Compute/disks/write	Yes	No
Microsoft.Storage/checknameavailability/read	Yes	Yes
Microsoft.Storage/operations/read	Yes	Yes
Microsoft.Storage/storageAccounts/listkeys/action	Yes	No
Microsoft.Storage/storageAccounts/read	Yes	Yes
Microsoft.Storage/storageAccounts/write	Yes	No
Microsoft.Storage/storageAccounts/blobServices/containers/read	Yes	Yes

Action	Used for set up?	Used for daily operations?
Microsoft.Network/networkInterfaces/read	Yes	Yes
Microsoft.Network/networkInterfaces/write	Yes	No
Microsoft.Network/networkInterfaces/join/action	Yes	No
Microsoft.Network/networkSecurityGroups/read	Yes	Yes
Microsoft.Network/networkSecurityGroups/write	Yes	No
Microsoft.Resources/subscriptions/locations/read	Yes	Yes
Microsoft.Network/locations/operationResults/read	Yes	Yes
Microsoft.Network/locations/operations/read	Yes	Yes
Microsoft.Network/virtualNetworks/read	Yes	Yes
Microsoft.Network/virtualNetworks/checkIpAddressAvailability/read	Yes	Yes
Microsoft.Network/virtualNetworks/subnets/read	Yes	Yes
Microsoft.Network/virtualNetworks/subnets/virtualMachines/read	Yes	Yes
Microsoft.Network/virtualNetworks/virtualMachines/read	Yes	Yes
Microsoft.Network/virtualNetworks/subnets/join/action	Yes	No
Microsoft.Network/virtualNetworks/subnets/write	Yes	No
Microsoft.Network/routeTables/join/action	Yes	No
Microsoft.Resources/deployments/operations/read	Yes	Yes
Microsoft.Resources/deployments/read	Yes	Yes
Microsoft.Resources/deployments/write	Yes	No
Microsoft.Resources/resources/read	Yes	Yes

Action	Used for set up?	Used for daily operations?
Microsoft.Resources/subscriptions/operationresults/read	Yes	Yes
Microsoft.Resources/subscriptions/resourceGroups/delete	Yes	No
Microsoft.Resources/subscriptions/resourceGroups/read	Yes	Yes
Microsoft.Resources/subscriptions/resourcegroups/resources/read	Yes	Yes
Microsoft.Resources/subscriptions/resourceGroups/write	Yes	No

### Cloud Volumes ONTAP

The Connector makes the following API requests to deploy and manage Cloud Volumes ONTAP in Azure.

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Create and manage VMs	Microsoft.Compute/locations/operations/read	Yes	Yes	No
	Microsoft.Compute/locations/vmSizes/read	Yes	Yes	No
	Microsoft.Resources/subscriptions/locations/read	Yes	No	No
	Microsoft.Compute/operations/read	Yes	Yes	No
	Microsoft.Compute/virtualMachines/instanceView/read	Yes	Yes	No
	Microsoft.Compute/virtualMachines/powerOff/action	Yes	Yes	No
	Microsoft.Compute/virtualMachines/read	Yes	Yes	No
	Microsoft.Compute/virtualMachines/restart/action	Yes	Yes	No
	Microsoft.Compute/virtualMachines/start/action	Yes	Yes	No
	Microsoft.Compute/virtualMachines/deallocate/action	No	Yes	Yes
	Microsoft.Compute/virtualMachines/vmSizes/read	No	Yes	No
	Microsoft.Compute/virtualMachines/write	Yes	Yes	No
	Microsoft.Compute/virtualMachines/delete	Yes	Yes	Yes
	Microsoft.Resources/deployments/delete	Yes	No	No
Enable deployment from a VHD	Microsoft.Compute/images/read	Yes	No	No
	Microsoft.Compute/images/write	Yes	No	No

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Create and manage network interfaces in the target subnet	Microsoft.Network/networkInterfaces/read	Yes	Yes	No
	Microsoft.Network/networkInterfaces/write	Yes	Yes	No
	Microsoft.Network/networkInterfaces/join/action	Yes	Yes	No
	Microsoft.Network/networkInterfaces/delete	Yes	Yes	No
Create and manage network security groups	Microsoft.Network/networkSecurityGroups/read	Yes	Yes	No
	Microsoft.Network/networkSecurityGroups/write	Yes	Yes	No
	Microsoft.Network/networkSecurityGroups/join/action	Yes	No	No
	Microsoft.Network/networkSecurityGroups/delete	No	Yes	Yes

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Get network information about regions, the target VNet and subnet, and add the VMs to VNets	Microsoft.Network/locations/operationResults/read	Yes	Yes	No
	Microsoft.Network/locations/operations/read	Yes	Yes	No
	Microsoft.Network/virtualNetworks/read	Yes	No	No
	Microsoft.Network/virtualNetworks/checkIpAddressAvailability/read	Yes	No	No
	Microsoft.Network/virtualNetworks/subnets/read	Yes	Yes	No
	Microsoft.Network/virtualNetworks/subnets/virtualMachines/read	Yes	Yes	No
	Microsoft.Network/virtualNetworks/virtualMachines/read	Yes	Yes	No
	Microsoft.Network/virtualNetworks/subnets/join/action	Yes	Yes	No



<b>Purpose</b>	<b>Action</b>	<b>Used for deployment?</b>	<b>Used for daily operations?</b>	<b>Used for deletion?</b>
Create and manage resource groups	Microsoft.Resources/deployments/operations/read	Yes	Yes	No
	Microsoft.Resources/deployments/read	Yes	Yes	No
	Microsoft.Resources/deployments/write	Yes	Yes	No
	Microsoft.Resources/resources/read	Yes	Yes	No
	Microsoft.Resources/subscriptions/operationresults/read	Yes	Yes	No
	Microsoft.Resources/subscriptions/resourceGroups/delete	Yes	Yes	Yes
	Microsoft.Resources/subscriptions/resourceGroups/read	No	Yes	No
	Microsoft.Resources/subscriptions/resourcegroups/resources/read	Yes	Yes	No
	Microsoft.Resources/subscriptions/resourceGroups/write	Yes	Yes	No

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Manage Azure storage accounts and disks	Microsoft.Compute/disks/read	Yes	Yes	Yes
	Microsoft.Compute/disks/write	Yes	Yes	No
	Microsoft.Compute/disks/delete	Yes	Yes	Yes
	Microsoft.Storage/checknameavailability/read	Yes	Yes	No
	Microsoft.Storage/operations/read	Yes	Yes	No
	Microsoft.Storage/storageAccounts/listkeys/action	Yes	Yes	No
	Microsoft.Storage/storageAccounts/read	Yes	Yes	No
	Microsoft.Storage/storageAccounts/delete	No	Yes	Yes
	Microsoft.Storage/storageAccounts/write	Yes	Yes	No
	Microsoft.Storage/usage/read	No	Yes	No
Enable backups to Blob storage and encryption of storage accounts	Microsoft.Storage/storageAccounts/blobServices/containers/read	Yes	Yes	No
	Microsoft.KeyVault/vaults/read	Yes	Yes	No
	Microsoft.KeyVault/vaults/accessPolicies/write	Yes	Yes	No
Enable VNet service endpoints for data tiering	Microsoft.Network/virtualNetworks/subnets/write	Yes	Yes	No
	Microsoft.Network/routeTables/join/action	Yes	Yes	No

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Create and manage Azure managed snapshots	Microsoft.Compute/snapshots/write	Yes	Yes	No
	Microsoft.Compute/snapshots/read	Yes	Yes	No
	Microsoft.Compute/snapshots/delete	No	Yes	Yes
	Microsoft.Compute/disks/beginGetAccess/action	No	Yes	No
Create and manage availability sets	Microsoft.Compute/availabilitySets/write	Yes	No	No
	Microsoft.Compute/availabilitySets/read	Yes	No	No
Enable programmatic deployments from the marketplace	Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/read	Yes	No	No
	Microsoft.MarketplaceOrdering/offertypes/publishers/offers/plans/agreements/write	Yes	Yes	No

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Manage a load balancer for HA pairs	Microsoft.Network/loadBalancers/read	Yes	Yes	No
	Microsoft.Network/loadBalancers/write	Yes	No	No
	Microsoft.Network/loadBalancers/delete	No	Yes	Yes
	Microsoft.Network/loadBalancers/backendAddressPools/read	Yes	No	No
	Microsoft.Network/loadBalancers/backendAddressPools/join/action	Yes	No	No
	Microsoft.Network/loadBalancers/frontendIPConfigurations/read	Yes	Yes	No
	Microsoft.Network/loadBalancers/loadBalancingRules/read	Yes	No	No
	Microsoft.Network/loadBalancers/probes/read	Yes	No	No
	Microsoft.Network/loadBalancers/probes/join/action	Yes	No	No
Enable management of locks on Azure disks	Microsoft.Authorization/locks/*	Yes	Yes	No

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Enable private endpoints for HA pairs when there's no connectivity outside the subnet	Microsoft.Network/privateEndpoints/write	Yes	Yes	No
	Microsoft.Storage/storageAccounts/PrivateEndpointConnectionsApproval/action	Yes	No	No
	Microsoft.Storage/storageAccounts/privateEndpointConnections/read	Yes	Yes	Yes
	Microsoft.Network/privateEndpoints/read	Yes	Yes	Yes
	Microsoft.Network/privateDnsZones/write	Yes	Yes	No
	Microsoft.Network/privateDnsZones/virtualNetworkLinks/write	Yes	Yes	No
	Microsoft.Network/virtualNetworks/join/action	Yes	Yes	No
	Microsoft.Network/privateDnsZones/A/write	Yes	Yes	No
	Microsoft.Network/privateDnsZones/read	Yes	Yes	No
	Microsoft.Network/privateDnsZones/virtualNetworkLinks/read	Yes	Yes	No
Required for some VM deployments, depending on the underlying physical hardware	Microsoft.Resources/deployments/operationStatuses/read	Yes	Yes	No
Remove resources from a resource group in case of deployment failure or deletion	Microsoft.Network/privateEndpoints/delete	Yes	Yes	No
	Microsoft.Compute/availabilitySets/delete	Yes	Yes	No

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Enable the use of customer-managed encryption keys when using the API	Microsoft.Compute/diskEncryptionSets/read	Yes	Yes	Yes
	Microsoft.Compute/diskEncryptionSets/write	Yes	Yes	No
	Microsoft.KeyVault/vaults/deploy/action	Yes	No	No
	Microsoft.Compute/diskEncryptionSets/delete	Yes	Yes	Yes
Configure an application security group for an HA pair to isolate the HA interconnect and cluster network NICs	Microsoft.Network/applicationSecurityGroups/write	No	Yes	No
	Microsoft.Network/applicationSecurityGroups/read	No	Yes	No
	Microsoft.Network/applicationSecurityGroups/joinIpConfiguration/action	No	Yes	No
	Microsoft.Network/networkSecurityGroups/securityRules/write	Yes	Yes	No
	Microsoft.Network/applicationSecurityGroups/delete	No	Yes	Yes
	Microsoft.Network/networkSecurityGroups/securityRules/delete	No	Yes	Yes
Read, write, and delete tags associated with Cloud Volumes ONTAP resources	Microsoft.Resources/tags/read	No	Yes	No
	Microsoft.Resources/tags/write	Yes	Yes	No
	Microsoft.Resources/tags/delete	Yes	No	No
Encrypt storage accounts during creation	Microsoft.ManagedIdentity/userAssignedIdentities/assign/action	Yes	Yes	No

Purpose	Action	Used for deployment?	Used for daily operations?	Used for deletion?
Use Virtual Machine Scale Sets in Flexible orchestration mode in order to specify specific zones for Cloud Volumes ONTAP	Microsoft.Compute/virtualMachineScaleSets/write	Yes	No	No
	Microsoft.Compute/virtualMachineScaleSets/read	Yes	No	No
	Microsoft.Compute/virtualMachineScaleSets/delete	No	No	Yes

## Tiering

The Connector makes the following API requests when you set up BlueXP tiering.

- Microsoft.Storage/storageAccounts/listkeys/action
- Microsoft.Resources/subscriptions/resourceGroups/read
- Microsoft.Resources/subscriptions/locations/read

The Connector makes the following API requests for daily operations.

- Microsoft.Storage/storageAccounts/blobServices/containers/read
- Microsoft.Storage/storageAccounts/managementPolicies/read
- Microsoft.Storage/storageAccounts/managementPolicies/write
- Microsoft.Storage/storageAccounts/read

## Change log

As permissions are added and removed, we'll note them in the sections below.

### 9 September 2024

The following permissions were removed from the JSON policy because BlueXP no longer supports discovery and management of Kubernetes clusters:

- Microsoft.ContainerService/managedClusters/listClusterUserCredential/action
- Microsoft.ContainerService/managedClusters/read

### 22 August 2024

The following permissions were added to the JSON policy because they are required for Cloud Volumes ONTAP support of Virtual Machine Scale Sets:

- Microsoft.Compute/virtualMachineScaleSets/write
- Microsoft.Compute/virtualMachineScaleSets/read
- Microsoft.Compute/virtualMachineScaleSets/delete

## 5 December 2023

The following permissions are no longer needed for BlueXP backup and recovery when backing up volume data to Azure Blob storage:

- Microsoft.Compute/virtualMachines/read
- Microsoft.Compute/virtualMachines/start/action
- Microsoft.Compute/virtualMachines/deallocate/action
- Microsoft.Compute/virtualMachines/extensions/delete
- Microsoft.Compute/virtualMachines/delete

These permissions are required for other BlueXP storage services, so they'll still remain in the custom role for the Connector if you're using those other storage services.

## 12 May 2023

The following permissions were added to the JSON policy because they are required for Cloud Volumes ONTAP management:

- Microsoft.Compute/images/write
- Microsoft.Network/loadBalancers/frontendIPConfigurations/read

The following permissions were removed from the JSON policy because they are no longer required:

- Microsoft.Storage/storageAccounts/blobServices/containers/write
- Microsoft.Network/publicIPAddresses/delete

## 23 March 2023

The "Microsoft.Storage/storageAccounts/delete" permission is no longer needed for BlueXP classification.

This permission is still required for Cloud Volumes ONTAP.

## 5 January 2023

The following permissions were added to the JSON policy:

- Microsoft.Storage/storageAccounts/listAccountSas/action
- Microsoft.Synapse/workspaces/privateEndpointConnectionsApproval/action

These permissions are required for BlueXP backup and recovery.

- Microsoft.Network/loadBalancers/backendAddressPools/join/action

This permission is required for Cloud Volumes ONTAP deployment.

## Google Cloud permissions for the Connector

BlueXP requires permissions to perform actions in Google Cloud. These permissions are included in a custom role provided by NetApp. You might want to understand what BlueXP does with these permissions.



## Service account permissions

The custom role shown below provides the permissions that a Connector needs to manage resources and processes within your Google Cloud network.

You'll need to apply this custom role to a service account that gets attached to the Connector VM.

- [Set up Google Cloud permissions for standard mode](#)
- [Set up permissions for restricted mode](#)
- [Set up permissions for private mode](#)

You also need to ensure that the role is up to date as new permissions are added in subsequent releases. If new permissions are required, they will be listed in the release notes.

```
title: NetApp BlueXP
description: Permissions for the service account associated with the
Connector instance.
stage: GA
includedPermissions:
- iam.serviceAccounts.actAs
- compute.regionBackendServices.create
- compute.regionBackendServices.get
- compute.regionBackendServices.list
- compute.networks.updatePolicy
- compute.backendServices.create
- compute.addresses.list
- compute.disks.create
- compute.disks.createSnapshot
- compute.disks.delete
- compute.disks.get
- compute.disks.list
- compute.disks.setLabels
- compute.disks.use
- compute.firewalls.create
- compute.firewalls.delete
- compute.firewalls.get
- compute.firewalls.list
- compute.globalOperations.get
- compute.images.get
- compute.images.getFromFamily
- compute.images.list
- compute.images.useReadOnly
- compute.instances.addAccessConfig
- compute.instances.attachDisk
- compute.instances.create
- compute.instances.delete
- compute.instances.detachDisk
```

- `compute.instances.get`
- `compute.instances.getSerialPortOutput`
- `compute.instances.list`
- `compute.instances.setDeletionProtection`
- `compute.instances.setLabels`
- `compute.instances.setMachineType`
- `compute.instances.setMetadata`
- `compute.instances.setTags`
- `compute.instances.start`
- `compute.instances.stop`
- `compute.instances.updateDisplayDevice`
- `compute.instanceGroups.get`
- `compute.addresses.get`
- `compute.instances.updateNetworkInterface`
- `compute.machineTypes.get`
- `compute.networks.get`
- `compute.networks.list`
- `compute.projects.get`
- `compute.regions.get`
- `compute.regions.list`
- `compute.snapshots.create`
- `compute.snapshots.delete`
- `compute.snapshots.get`
- `compute.snapshots.list`
- `compute.snapshots.setLabels`
- `compute.subnetworks.get`
- `compute.subnetworks.list`
- `compute.subnetworks.use`
- `compute.subnetworks.useExternalIp`
- `compute.zoneOperations.get`
- `compute.zones.get`
- `compute.zones.list`
- `compute.instances.setServiceAccount`
- `deploymentmanager.compositeTypes.get`
- `deploymentmanager.compositeTypes.list`
- `deploymentmanager.deployments.create`
- `deploymentmanager.deployments.delete`
- `deploymentmanager.deployments.get`
- `deploymentmanager.deployments.list`
- `deploymentmanager.manifests.get`
- `deploymentmanager.manifests.list`
- `deploymentmanager.operations.get`
- `deploymentmanager.operations.list`
- `deploymentmanager.resources.get`
- `deploymentmanager.resources.list`
- `deploymentmanager.typeProviders.get`

- deploymentmanager.typeProviders.list
- deploymentmanager.types.get
- deploymentmanager.types.list
- logging.logEntries.list
- logging.privateLogEntries.list
- resourcemanager.projects.get
- storage.buckets.create
- storage.buckets.delete
- storage.buckets.get
- storage.buckets.list
- cloudkms.cryptoKeyVersions.useToEncrypt
- cloudkms.cryptoKeys.get
- cloudkms.cryptoKeys.list
- cloudkms.keyRings.list
- storage.buckets.update
- iam.serviceAccounts.getIamPolicy
- iam.serviceAccounts.list
- storage.objects.get
- storage.objects.list
- monitoring.timeSeries.list
- storage.buckets.getIamPolicy
- cloudkms.cryptoKeys.getIamPolicy
- cloudkms.cryptoKeys.setIamPolicy
- cloudkms.keyRings.get
- cloudkms.keyRings.getIamPolicy
- cloudkms.keyRings.setIamPolicy

### How Google Cloud permissions are used

Actions	Purpose
<ul style="list-style-type: none"> <li>- compute.disks.create</li> <li>- compute.disks.createSnapshot</li> <li>- compute.disks.delete</li> <li>- compute.disks.get</li> <li>- compute.disks.list</li> <li>- compute.disks.setLabels</li> <li>- compute.disks.use</li> </ul>	To create and manage disks for Cloud Volumes ONTAP.
<ul style="list-style-type: none"> <li>- compute.firewalls.create</li> <li>- compute.firewalls.delete</li> <li>- compute.firewalls.get</li> <li>- compute.firewalls.list</li> </ul>	To create firewall rules for Cloud Volumes ONTAP.
<ul style="list-style-type: none"> <li>- compute.globalOperations.get</li> </ul>	To get the status of operations.
<ul style="list-style-type: none"> <li>- compute.images.get</li> <li>- compute.images.getFromFamily</li> <li>- compute.images.list</li> <li>- compute.images.useReadOnly</li> </ul>	To get images for VM instances.

<b>Actions</b>	<b>Purpose</b>
- compute.instances.attachDisk - compute.instances.detachDisk	To attach and detach disks to Cloud Volumes ONTAP.
- compute.instances.create - compute.instances.delete	To create and delete Cloud Volumes ONTAP VM instances.
- compute.instances.get	To list VM instances.
- compute.instances.getSerialPortOutput	To get console logs.
- compute.instances.list	To retrieve the list of instances in a zone.
- compute.instances.setDeletionProtection	To set deletion protection on the instance.
- compute.instances.setLabels	To add labels.
- compute.instances.setMachineType - compute.instances.setMinCpuPlatform	To change the machine type for Cloud Volumes ONTAP.
- compute.instances.setMetadata	To add metadata.
- compute.instances.setTags	To add tags for firewall rules.
- compute.instances.start - compute.instances.stop - compute.instances.updateDisplayDevice	To start and stop Cloud Volumes ONTAP.
- compute.machineTypes.get	To get the numbers of cores to check quotas.
- compute.projects.get	To support multi-projects.
- compute.snapshots.create - compute.snapshots.delete - compute.snapshots.get - compute.snapshots.list - compute.snapshots.setLabels	To create and manage persistent disk snapshots.
- compute.networks.get - compute.networks.list - compute.regions.get - compute.regions.list - compute.subnetworks.get - compute.subnetworks.list - compute.zoneOperations.get - compute.zones.get - compute.zones.list	To get the networking information needed to create a new Cloud Volumes ONTAP virtual machine instance.

Actions	Purpose
<ul style="list-style-type: none"> <li>- deploymentmanager.compositeTypes.get</li> <li>- deploymentmanager.compositeTypes.list</li> <li>- deploymentmanager.deployments.create</li> <li>- deploymentmanager.deployments.delete</li> <li>- deploymentmanager.deployments.get</li> <li>- deploymentmanager.deployments.list</li> <li>- deploymentmanager.manifests.get</li> <li>- deploymentmanager.manifests.list</li> <li>- deploymentmanager.operations.get</li> <li>- deploymentmanager.operations.list</li> <li>- deploymentmanager.resources.get</li> <li>- deploymentmanager.resources.list</li> <li>- deploymentmanager.typeProviders.get</li> <li>- deploymentmanager.typeProviders.list</li> <li>- deploymentmanager.types.get</li> <li>- deploymentmanager.types.list</li> </ul>	<p>To deploy the Cloud Volumes ONTAP virtual machine instance using Google Cloud Deployment Manager.</p>
<ul style="list-style-type: none"> <li>- logging.logEntries.list</li> <li>- logging.privateLogEntries.list</li> </ul>	<p>To get stack log drives.</p>
<ul style="list-style-type: none"> <li>- resourcemanager.projects.get</li> </ul>	<p>To support multi-projects.</p>
<ul style="list-style-type: none"> <li>- storage.buckets.create</li> <li>- storage.buckets.delete</li> <li>- storage.buckets.get</li> <li>- storage.buckets.list</li> <li>- storage.buckets.update</li> </ul>	<p>To create and manage a Google Cloud Storage bucket for data tiering.</p>
<ul style="list-style-type: none"> <li>- cloudkms.cryptoKeyVersions.useToEncrypt</li> <li>- cloudkms.cryptoKeys.get</li> <li>- cloudkms.cryptoKeys.list</li> <li>- cloudkms.keyRings.list</li> </ul>	<p>To use customer-managed encryption keys from the Cloud Key Management Service with Cloud Volumes ONTAP.</p>
<ul style="list-style-type: none"> <li>- compute.instances.setServiceAccount</li> <li>- iam.serviceAccounts.actAs</li> <li>- iam.serviceAccounts.getIamPolicy</li> <li>- iam.serviceAccounts.list</li> <li>- storage.objects.get</li> <li>- storage.objects.list</li> </ul>	<p>To set a service account on the Cloud Volumes ONTAP instance. This service account provides permissions for data tiering to a Google Cloud Storage bucket.</p>
<ul style="list-style-type: none"> <li>- compute.addresses.list</li> </ul>	<p>To retrieve the addresses in a region when deploying an HA pair.</p>
<ul style="list-style-type: none"> <li>- compute.backendServices.create</li> <li>- compute.regionBackendServices.create</li> <li>- compute.regionBackendServices.get</li> <li>- compute.regionBackendServices.list</li> </ul>	<p>To configure a backend service for distributing traffic in an HA pair.</p>
<ul style="list-style-type: none"> <li>- compute.networks.updatePolicy</li> </ul>	<p>To apply firewall rules on the VPCs and subnets for an HA pair.</p>
<ul style="list-style-type: none"> <li>- compute.subnetworks.use</li> <li>- compute.subnetworks.useExternalIp</li> <li>- compute.instances.addAccessConfig</li> </ul>	<p>To enable BlueXP classification.</p>

Actions	Purpose
<ul style="list-style-type: none"> <li>- compute.instanceGroups.get</li> <li>- compute.addresses.get</li> <li>- compute.instances.updateNetworkInterface</li> </ul>	To create and manage storage VMs on Cloud Volumes ONTAP HA pairs.
<ul style="list-style-type: none"> <li>- monitoring.timeSeries.list</li> <li>- storage.buckets.getIamPolicy</li> </ul>	To discover information about Google Cloud Storage buckets.
<ul style="list-style-type: none"> <li>- cloudkms.cryptoKeys.get</li> <li>- cloudkms.cryptoKeys.getIamPolicy</li> <li>- cloudkms.cryptoKeys.list</li> <li>- cloudkms.cryptoKeys.setIamPolicy</li> <li>- cloudkms.keyRings.get</li> <li>- cloudkms.keyRings.getIamPolicy</li> <li>- cloudkms.keyRings.list</li> <li>- cloudkms.keyRings.setIamPolicy</li> </ul>	To select your own customer-managed keys in the BlueXP backup and recovery activation wizard instead of using the default Google-managed encryption keys.

## Change log

As permissions are added and removed, we'll note them in the sections below.

### 6 February, 2023

The following permission was added to this policy:

- compute.instances.updateNetworkInterface

This permission is required for Cloud Volumes ONTAP.

### 27 January, 2023

The following permissions were added to the policy:

- cloudkms.cryptoKeys.getIamPolicy
- cloudkms.cryptoKeys.setIamPolicy
- cloudkms.keyRings.get
- cloudkms.keyRings.getIamPolicy
- cloudkms.keyRings.setIamPolicy

These permissions are required for BlueXP backup and recovery.

## Ports

### Connector security group rules in AWS

The AWS security group for the Connector requires both inbound and outbound rules. BlueXP automatically creates this security group when you create a Connector from BlueXP. You need to set up this security group for all other installation options.

## Inbound rules

Protocol	Port	Purpose
SSH	22	Provides SSH access to the Connector host
HTTP	80	<ul style="list-style-type: none"><li>• Provides HTTP access from client web browsers to the local user interface</li><li>• Used during the Cloud Volumes ONTAP upgrade process</li></ul>
HTTPS	443	Provides HTTPS access from client web browsers to the local user interface, and connections from the BlueXP classification instance
TCP	3128	Provides Cloud Volumes ONTAP with internet access. You must manually open this port after deployment.

## Outbound rules

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

### Basic outbound rules

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

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

### Advanced outbound rules

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



The source IP address is the Connector host.

Service	Protocol	Port	Destination	Purpose
API calls and AutoSupport	HTTPS	443	Outbound internet and ONTAP cluster management LIF	API calls to AWS, to ONTAP, to BlueXP classification, and sending AutoSupport messages to NetApp
API calls	TCP	3000	ONTAP HA mediator	Communication with the ONTAP HA mediator
	TCP	8080	BlueXP classification	Probe to BlueXP classification instance during deployment
DNS	UDP	53	DNS	Used for DNS resolve by BlueXP

## Connector security group rules in Azure

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

BlueXP automatically creates this security group when you create a Connector from BlueXP. You need to set up this security group for all other installation options.

### Inbound rules

Protocol	Port	Purpose
SSH	22	Provides SSH access to the Connector host
HTTP	80	<ul style="list-style-type: none"> <li>Provides HTTP access from client web browsers to the local user interface</li> <li>Used during the Cloud Volumes ONTAP upgrade process</li> </ul>
HTTPS	443	Provides HTTPS access from client web browsers to the local user interface, and connections from the BlueXP classification instance
TCP	3128	Provides Cloud Volumes ONTAP with internet access to send AutoSupport messages to NetApp Support. You must manually open this port after deployment. <a href="#">Learn how the Connector is used as a proxy for AutoSupport messages</a>

### Outbound rules

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

#### Basic outbound rules

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

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

#### Advanced outbound rules

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



The source IP address is the Connector host.

Service	Protocol	Port	Destination	Purpose
API calls and AutoSupport	HTTPS	443	Outbound internet and ONTAP cluster management LIF	API calls to Azure, to ONTAP, to BlueXP classification, and sending AutoSupport messages to NetApp
API calls	TCP	8080	BlueXP classification	Probe to BlueXP classification instance during deployment
DNS	UDP	53	DNS	Used for DNS resolve by BlueXP



## Connector firewall rules in Google Cloud

The Google Cloud firewall rules for the Connector requires both inbound and outbound rules. BlueXP automatically creates this security group when you create a Connector from BlueXP. You need to set up this security group for all other installation options.

### Inbound rules

Protocol	Port	Purpose
SSH	22	Provides SSH access to the Connector host
HTTP	80	<ul style="list-style-type: none"><li>• Provides HTTP access from client web browsers to the local user interface</li><li>• Used during the Cloud Volumes ONTAP upgrade process</li></ul>
HTTPS	443	Provides HTTPS access from client web browsers to the local user interface
TCP	3128	Provides Cloud Volumes ONTAP with internet access. You must manually open this port after deployment.

### Outbound rules

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

#### Basic outbound rules

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

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

#### Advanced outbound rules

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



The source IP address is the Connector host.

Service	Protocol	Port	Destination	Purpose
API calls and AutoSupport	HTTPS	443	Outbound internet and ONTAP cluster management LIF	API calls to Google Cloud, to ONTAP, to BlueXP classification, and sending AutoSupport messages to NetApp
API calls	TCP	80 80	BlueXP classification	Probe to BlueXP classification instance during deployment
DNS	UDP	53	DNS	Used for DNS resolve by BlueXP

## Ports for the on-premisesConnector

The Connector uses *inbound* ports when installed manually on an on-premises Linux host. You might need to refer to these ports for planning purposes.

These inbound rules apply to all BlueXP deployment models.

Protocol	Port	Purpose
HTTP	80	<ul style="list-style-type: none"><li>• Provides HTTP access from client web browsers to the local user interface</li><li>• Used during the Cloud Volumes ONTAP upgrade process</li></ul>
HTTPS	443	Provides HTTPS access from client web browsers to the local user interface

# 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](#)
- [Google Cloud NetApp Volumes](#)

## Support registration overview

There are two forms of registration to activate support entitlement:

- Registering your BlueXP account serial number (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 BlueXP for NetApp support

To register for support and activate support entitlement, one user in your BlueXP organization (or 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 BlueXP organization 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 organization is not registered for support. As long as one user in the organization has followed these steps, then your organization 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.



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 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 organization 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 organization is different than the NSS account that is associated with a BlueXP user login.

These NSS credentials are associated with your specific BlueXP organization ID. Users who belong to the BlueXP organization 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 24/7, 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](#)
- [Google Cloud NetApp Volumes](#)

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 organization (or account), project (or 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 Working Enviroment


Select 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) Upload 

No files selected  

### 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.



- Filter the contents of the columns.



- Change the columns that appear in the table by selecting  and then choosing the columns that you'd like to display.




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](#)

## Copyright information

Copyright © 2025 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.



# **BlueXP backup and recovery documentation**

BlueXP backup and recovery

NetApp  
August 19, 2025

# Table of Contents

- BlueXP backup and recovery documentation ..... 1
- Release notes ..... 2
  - What's new in BlueXP backup and recovery ..... 2
    - 12 August 2025 ..... 2
    - 28 July 2025 ..... 4
    - 14 July 2025 ..... 5
    - 09 June 2025 ..... 6
    - 13 May 2025 ..... 7
    - 16 April 2025 ..... 8
    - 17 March 2025 ..... 9
    - 21 February 2025 ..... 10
    - 13 February 2025 ..... 11
    - 22 November 2024 ..... 11
    - 27 September 2024 ..... 12
- Known limitations with BlueXP backup and recovery for Microsoft SQL Server workloads ..... 13
  - Clone lifecycle support ..... 13
  - Standard deployment mode only ..... 13
  - Windows cluster name restriction ..... 13
  - SnapCenter migration issues ..... 13
- Known limitations with BlueXP backup and recovery for ONTAP volumes ..... 14
  - Replication limitations for ONTAP volumes ..... 15
  - Backup-to-object limitations for ONTAP volumes ..... 15
  - Restore limitations for ONTAP volumes ..... 16
- Get started ..... 18
  - Learn about BlueXP backup and recovery ..... 18
    - What you can do with BlueXP backup and recovery ..... 18
    - Benefits of using BlueXP backup and recovery ..... 18
    - Cost ..... 19
    - Licensing ..... 20
    - Supported data sources, working environments, and backup targets ..... 21
    - BlueXP backup and recovery uses the Plug-in for Microsoft SQL Server ..... 21
    - How BlueXP backup and recovery works ..... 21
    - Terms that might help you with BlueXP backup and recovery ..... 23
- BlueXP backup and recovery prerequisites ..... 23
  - For ONTAP 9.8 and later ..... 23
  - Prerequisites for backups to object storage ..... 23
  - Microsoft SQL Server workload requirements ..... 23
  - Kubernetes workload requirements ..... 24
  - In BlueXP ..... 25
- Set up licensing for BlueXP backup and recovery ..... 25
  - 30-day free trial ..... 26
  - Use a BlueXP backup and recovery PAYGO subscription ..... 26
  - Use an annual contract ..... 27



- Use a BlueXP backup and recovery BYOL license . . . . . 28
- Set up backup destinations before you use BlueXP backup and recovery . . . . . 29
  - Prepare the backup destination . . . . . 29
  - Set up S3 permissions . . . . . 29
- Log in to BlueXP backup and recovery . . . . . 32
- Discover offsite backup targets in BlueXP backup and recovery . . . . . 33
  - Discover a backup target . . . . . 33
  - Add a bucket for a backup target . . . . . 35
  - Change credentials for a backup target . . . . . 36
- Switch to different BlueXP backup and recovery workloads . . . . . 36
  - Switch to a different workload . . . . . 37
- Configure BlueXP backup and recovery settings . . . . . 37
  - Add credentials for host resources . . . . . 37
  - Maintain VMware vCenter settings . . . . . 40
  - Import and manage SnapCenter host resources . . . . . 41
  - Configure log directories in snapshots for Windows hosts . . . . . 47
- Use BlueXP backup and recovery . . . . . 49
  - View protection health on the BlueXP backup and recovery Dashboard . . . . . 49
    - View the overall system summary . . . . . 49
    - View the Protection summary . . . . . 50
    - View the Job summary . . . . . 50
    - View the Restore summary . . . . . 50
- Create and manage policies to govern backups in BlueXP backup and recovery . . . . . 50
  - View policies . . . . . 51
  - Create a policy . . . . . 51
  - Edit a policy . . . . . 59
  - Delete a policy . . . . . 60
- Protect ONTAP volume workloads . . . . . 60
  - Protect your ONTAP volume data using BlueXP backup and recovery . . . . . 60
  - Plan your protection journey with BlueXP backup and recovery . . . . . 69
  - Manage backup policies for ONTAP volumes with BlueXP backup and recovery . . . . . 75
  - Backup-to-object policy options in BlueXP backup and recovery . . . . . 80
  - Manage backup-to-object storage options in BlueXP backup and recovery Advanced Settings . . . . . 88
  - Back up Cloud Volumes ONTAP data to Amazon S3 with BlueXP backup and recovery . . . . . 92
  - Back up Cloud Volumes ONTAP data to Azure Blob storage with BlueXP backup and recovery . . . . . 101
  - Back up Cloud Volumes ONTAP data to Google Cloud Storage with BlueXP backup and recovery . . . . . 110
  - Back up on-premises ONTAP data to Amazon S3 with BlueXP backup and recovery . . . . . 120
  - Back up on-premises ONTAP data to Azure Blob storage with BlueXP backup and recovery . . . . . 132
  - Back up on-premises ONTAP data to Google Cloud Storage with BlueXP backup and recovery . . . . . 142
  - Back up on-premises ONTAP data to ONTAP S3 with BlueXP backup and recovery . . . . . 153
  - Back up on-premises ONTAP data to StorageGRID with BlueXP backup and recovery . . . . . 162
  - Migrate volumes using SnapMirror to Cloud Resync with BlueXP backup and recovery . . . . . 171
  - Restore BlueXP backup and recovery configuration data in a dark site . . . . . 176
  - Manage backups for your ONTAP systems with BlueXP backup and recovery . . . . . 180
  - Restore ONTAP data from backup files with BlueXP backup and recovery . . . . . 195

- Protect Microsoft SQL Server workloads . . . . . 214
  - Protect Microsoft SQL workloads overview with BlueXP backup and recovery . . . . . 215
  - Prerequisites for importing from the Plug-in service into BlueXP backup and recovery . . . . . 216
  - Discover Microsoft SQL Server workloads and optionally import from SnapCenter in BlueXP backup and recovery . . . . . 219
  - Back up Microsoft SQL Server workloads with BlueXP backup and recovery . . . . . 228
  - Restore Microsoft SQL Server workloads with BlueXP backup and recovery . . . . . 235
  - Clone Microsoft SQL Server workloads with BlueXP backup and recovery . . . . . 244
  - Manage Microsoft SQL Server inventory with BlueXP backup and recovery . . . . . 251
  - Manage Microsoft SQL Server snapshots with BlueXP backup and recovery . . . . . 259
  - Create reports for Microsoft SQL Server workloads in BlueXP backup and recovery . . . . . 260
- Protect virtual machine workloads . . . . . 261
  - Protect virtual machines workloads in BlueXP backup and recovery overview . . . . . 261
  - Prerequisites for virtual machines workloads in BlueXP backup and recovery . . . . . 261
  - Register SnapCenter Plug-in for VMware vSphere host to use with BlueXP backup and recovery . . . . . 263
  - Create a policy to back up datastores in BlueXP backup and recovery . . . . . 264
  - Back up datastores to Amazon Web Services in BlueXP backup and recovery . . . . . 265
  - Back up datastores to Microsoft Azure with BlueXP backup and recovery . . . . . 266
  - Back up datastores to Google Cloud Platform with BlueXP backup and recovery . . . . . 267
  - Back up datastores to StorageGRID with BlueXP backup and recovery . . . . . 268
  - Manage protection of datastores and VMs in BlueXP backup and recovery . . . . . 268
  - Restore virtual machines data with BlueXP backup and recovery . . . . . 270
- Protect Kubernetes workloads (Preview) . . . . . 273
  - Manage Kubernetes workloads overview . . . . . 273
  - Discover Kubernetes workloads in BlueXP backup and recovery . . . . . 274
  - Add and protect Kubernetes applications . . . . . 275
  - Restore Kubernetes applications . . . . . 277
  - Manage Kubernetes clusters . . . . . 279
  - Manage Kubernetes applications . . . . . 280
  - Manage BlueXP backup and recovery execution hook templates for Kubernetes workloads . . . . . 281
- Monitor jobs in BlueXP backup and recovery . . . . . 283
  - View job status on the Job Monitor . . . . . 284
  - Review retention (backup lifecycle) jobs . . . . . 286
  - Review backup and restore alerts in the BlueXP Notification Center . . . . . 286
  - Review operation activity in the BlueXP Timeline . . . . . 287
- Restart the BlueXP backup and recovery service . . . . . 287
- Automate with BlueXP backup and recovery REST APIs . . . . . 289
  - API reference . . . . . 289
  - Getting started . . . . . 289
  - Example using the APIs . . . . . 291
- Reference . . . . . 294
  - Policies in SnapCenter compared to those in BlueXP backup and recovery . . . . . 294
    - Schedule tiers . . . . . 294
    - Multiple policies in SnapCenter with the same schedule tier . . . . . 294
    - Imported SnapCenter daily schedules . . . . . 294

Imported SnapCenter hourly schedules . . . . .	295
Log retention from SnapCenter policies . . . . .	295
Log backup retention . . . . .	295
Retention count from SnapCenter policies . . . . .	295
SnapMirror labels from SnapCenter policies . . . . .	296
BlueXP backup and recovery identity and access management to features . . . . .	296
Supported AWS archive storage tiers with BlueXP backup and recovery . . . . .	298
Supported S3 archival storage classes for BlueXP backup and recovery . . . . .	298
Restore data from archival storage . . . . .	298
Supported Azure archive access tiers with BlueXP backup and recovery . . . . .	299
Supported Azure Blob access tiers for BlueXP backup and recovery . . . . .	299
Restore data from archival storage . . . . .	300
Supported Google archive storage tiers with BlueXP backup and recovery . . . . .	300
Supported Google archival storage classes for BlueXP backup and recovery . . . . .	301
Restore data from archival storage . . . . .	301
Legal notices . . . . .	302
Copyright . . . . .	302
Trademarks . . . . .	302
Patents . . . . .	302
Privacy policy . . . . .	302
Open source . . . . .	302

# BlueXP backup and recovery documentation

# Release notes

## What's new in BlueXP backup and recovery

Learn what's new in BlueXP backup and recovery.

### 12 August 2025

This BlueXP backup and recovery release includes the following updates.

#### Microsoft SQL Server workload supported in General Availability (GA)

Microsoft SQL Server workload support is now generally available (GA) in BlueXP backup and recovery. Organizations using an MSSQL environment on ONTAP, Cloud Volumes ONTAP, and Amazon FSx for NetApp ONTAP storage can now take advantage of this new backup and recovery service to protect their data.

This release includes the following enhancements to the Microsoft SQL Server workload support from the previous preview version:

- **SnapMirror active sync:** This version now supports SnapMirror active sync (also referred to as SnapMirror Business Continuity [SM-BC]), which enables business services to continue operating even through a complete site failure, supporting applications to fail over transparently using a secondary copy. BlueXP backup and recovery now supports protection of Microsoft SQL Server databases in a SnapMirror active sync and Metrocluster configuration. The information appears in the **Storage and relationship status** section of the Protection details page. The relationship information is displayed in the updated **Secondary settings** section of the Policy page.

Refer to [Use policies to protect your workloads](#).

Microsoft SQL Server workload > Database\_name

View protection details

Database name: Database | Instance name: Instance | Host name: Database host | Microsoft SQL Server: Location | Ransomware protection: [Icons] | Protection health: Healthy

3-2-1 fan-out data flow

ONTAP Secondary | ONTAP Primary | Object Store

Protection

Policy name: PROD\_BKP  
Local schedules: cLUSTER\_NAME: PRIMARY\_SVM2  
LUN: LUN\_1, LUN\_2, LUN\_3  
Object store schedules: Daily, Weekly  
Availability group settings: Preferred replica  
Storage & relationship status: View

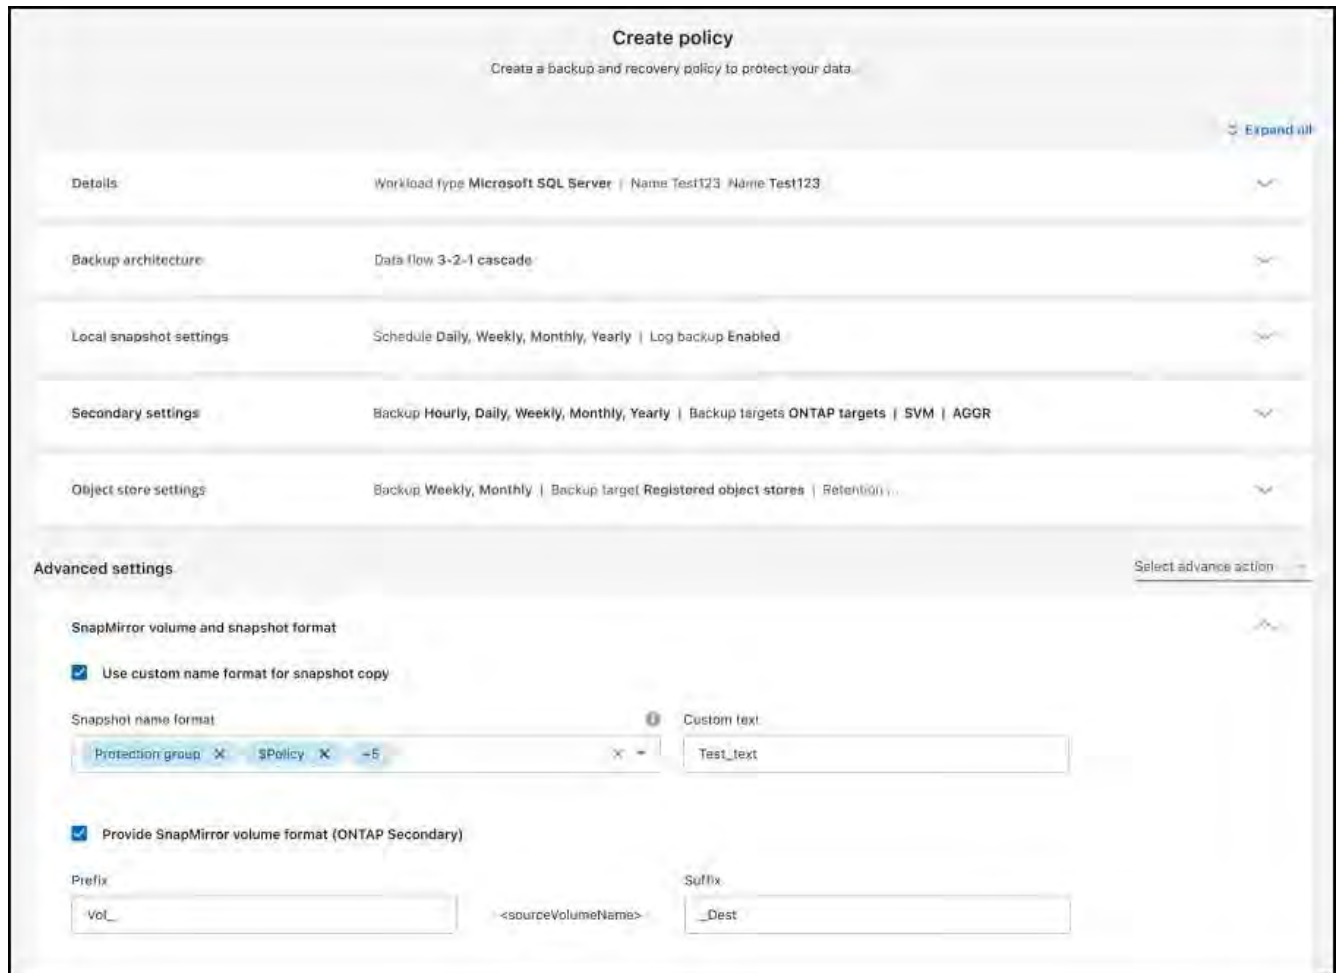
Recovery points (14)

Name	Backup type	Size	Location
SnapshotName_1	Full	25,125 GiB	[Icons]
SnapshotName_1	Log	25,125 GiB	[Icons]
SnapshotName_1	Log	25,125 GiB	[Icons]

- **Multi-bucket support:** You can now protect the volumes within a working environment with up to 6 buckets per working environment across different cloud providers.
- **Licensing and free trial updates** for SQL Server workloads: You can now use the existing BlueXP backup a recovery licensing model to protect SQL Server workloads. There is no separate licensing requirement for SQL Server workloads.

For details, refer to [Set up licensing for BlueXP backup and recovery](#).

- **Custom snapshot name:** You can now use your own snapshot name in a policy that governs the backups for Microsoft SQL Server workloads. Enter this information in the **Advanced settings** section of the Policy page.



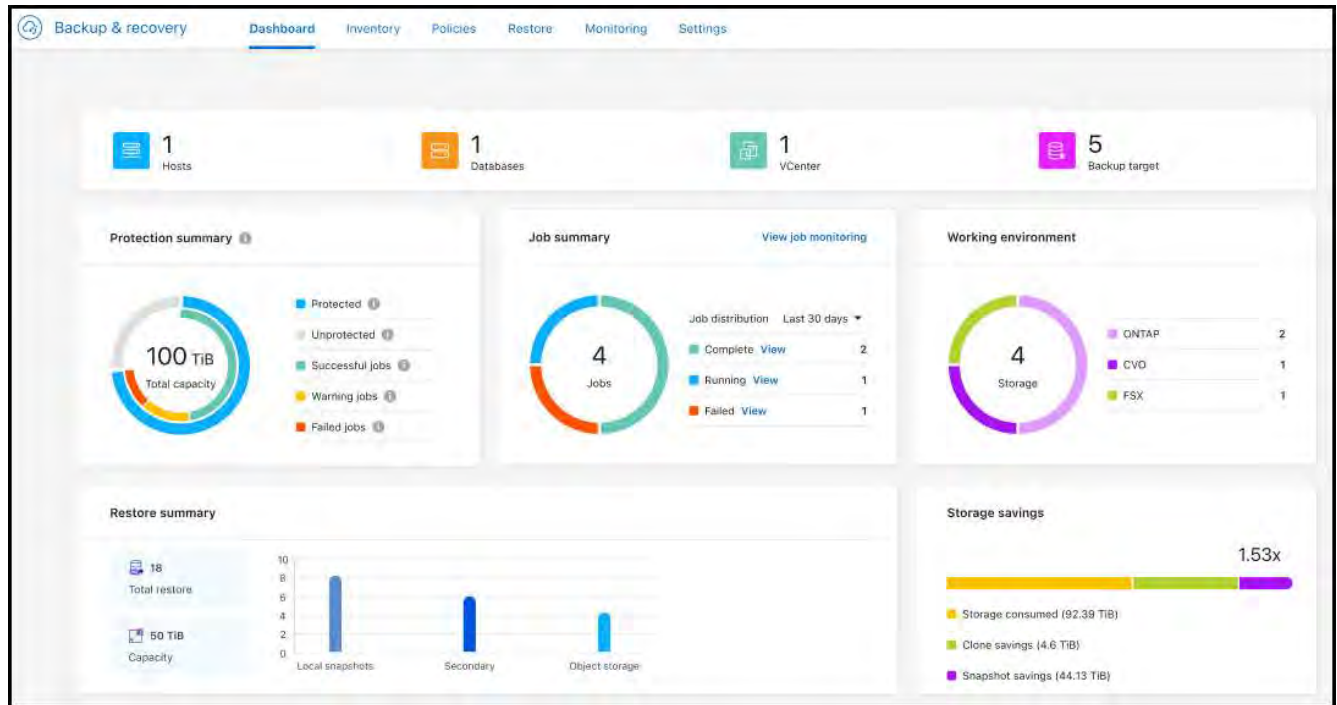
Refer to [Use policies to protect your workloads](#).

- **Secondary volume prefix and suffix:** You can enter a custom prefix and suffix in the **Advanced settings** section of the Policy page.
- **Identity and access management (IAM):** You can now control users' access to features.

Refer to [Log in to BlueXP backup and recovery](#) and [BlueXP backup and recovery access to features](#).

- **Restore from object storage to an alternate host:** You can now restore from object storage to an alternate host even if the primary storage is down.
- **Log backup data:** The database protection details page now shows log backups. You can see the Backup type column that shows whether the backup is a full backup or a log backup.

- **Enhanced Dashboard:** The Dashboard now shows Storage and Clone savings.



## ONTAP volume workload enhancements

- **Multi-folder restore for ONTAP volumes:** Until now, you could restore either one folder or multiple files at a time from the Browse and restore feature. BlueXP backup and recovery now provides the ability to select multiple folders at a time using the Browse and restore feature.
- **View and manage backups of deleted volumes:** The BlueXP backup and recovery Dashboard now gives an option to show and manage volumes that are deleted from ONTAP. With this, you can view and delete backups from volumes that no longer exist in ONTAP.
- **Force delete backups:** In some extreme cases, you might want BlueXP backup and recovery not to have access to backups any longer. This might happen for example, if the service no longer has access to the backup bucket or backups are DataLock protected but you don't want them anymore. Previously, you could not delete these yourself and needed to call NetApp Support. With this release, you can use the option to force delete backups (at volume and work environment levels).



Use this option carefully and only in extreme cleanup needs. BlueXP backup and recovery will not have access to these backups any longer even if they are not deleted in the object storage. You will need to go to your cloud provider and manually delete the backups.

Refer to [Protect ONTAP workloads](#).

## 28 July 2025

This BlueXP backup and recovery release includes the following updates.

### Kubernetes workload support as a Preview

This release of BlueXP backup and recovery introduces support for discovering and managing Kubernetes workloads:

- Discover Red Hat OpenShift and open-source Kubernetes clusters, backed by NetApp ONTAP, without sharing kubeconfig files.
- Discover, manage, and protect applications across multiple Kubernetes clusters using a unified control plane.
- Offload data movement operations for backup and recovery of Kubernetes applications to NetApp ONTAP.
- Orchestrate local and object-storage-based application backups.
- Back up and restore entire applications and individual resources to any Kubernetes clusters.
- Work with containers and virtual machines running on Kubernetes.
- Create application-consistent backups using execution hooks and templates.

For details about protecting Kubernetes workloads, refer to [Protect Kubernetes workloads overview](#).

## 14 July 2025

This BlueXP backup and recovery release includes the following updates.

### Enhanced ONTAP volume Dashboard

In April 2025, we launched a preview of an enhanced ONTAP volume Dashboard that is much faster and more efficient.

This dashboard was designed to help enterprise customers with a high number of workloads. Even for customers with 20,000 volumes, the new dashboard loads in <10 seconds.

After a successful preview and great feedback from preview customers, we are now making it the default experience for all our customers. Be ready for a blazingly fast dashboard.

For details, see [View protection health in the Dashboard](#).

### Microsoft SQL Server workload support as a Public Technology Preview

This release of BlueXP backup and recovery provides an updated user interface that enables you to manage Microsoft SQL Server workloads using a 3-2-1 protection strategy, familiar in the BlueXP backup and recovery service. With this new version, you can back up these workloads to primary storage, replicate them to secondary storage, and back them up to cloud object storage.

You can sign up for the preview by completing this [Preview Signup Form](#).



This documentation about protecting Microsoft SQL Server workloads is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before general availability.

This version of BlueXP backup and recovery includes the following updates:

- **3-2-1 backup capability:** This version integrates SnapCenter capabilities, enabling you to manage and protect your SnapCenter resources with a 3-2-1 data protection strategy from the BlueXP backup and recovery user interface.
- **Import from SnapCenter:** You can import SnapCenter backup data and policies into BlueXP backup and recovery.
- **A redesigned user interface** provides a more intuitive experience for managing your backup and recovery



tasks.

- **Backup targets:** You can add buckets in Amazon Web Services (AWS), Microsoft Azure Blob Storage, StorageGRID, and ONTAP S3 environments to use as backup targets for your Microsoft SQL Server workloads.
- **Workload support:** This version enables you to back up, restore, verify, and clone Microsoft SQL Server databases and availability groups. (Support for other workloads will be added in future releases.)
- **Flexible restore options:** This version enables you to restore databases to both original and alternate locations in case of corruption or accidental data loss.
- **Instant production copies:** Generate space-efficient production copies for development, testing, or analytics in minutes instead of hours or days.
- This version includes the ability to create detailed reports.

For details about protecting Microsoft SQL Server workloads, see [Protect Microsoft SQL Server workloads overview](#).

## 09 June 2025

This BlueXP backup and recovery release includes the following updates.

### Indexed catalog support updates

In February 2025, we introduced the updated indexing feature (Indexed Catalog v2) that you use during the Search & Restore method of restoring data. The previous release significantly improved data indexing performance in on-premises environments. With this release, the indexing catalog is now available with Amazon Web Services, Microsoft Azure, and Google Cloud Platform (GCP) environments.

If you are a new customer, the Indexed Catalog v2 is enabled by default for all new environments. If you are an existing customer, you can re-index your environment to leverage the Indexed Catalog v2.

### How do you enable indexing?

Before you can use the Search & Restore method of restoring data, you need to enable "Indexing" on each source working environment from which you're planning to restore volumes or files. Select the **Enable Indexing** option when you are performing a Search & Restore.

The Indexed Catalog can then track every volume and backup file, making your searches quick and efficient.

For more information, refer to [Enable indexing for Search & Restore](#).

### Azure private link endpoints and service endpoints

Typically, BlueXP backup and recovery establishes a private endpoint with the cloud provider to handle protection tasks. This release introduces an optional setting that lets you enable or disable BlueXP backup and recovery from automatically creating a private endpoint. This might be useful to you if you want more control over the private endpoint creation process.

You can enable or disable this option when you enable protection or start the restore process.

If you disable this setting, you must manually create the private endpoint for BlueXP backup and recovery to function properly. Without proper connectivity, you might not be able to perform backup and recovery tasks successfully.

## Support for SnapMirror to Cloud Resync on ONTAP S3

The previous release introduced support for SnapMirror to Cloud Resync (SM-C Resync). The feature streamlines data protection during volume migration in NetApp environments. This release add support for SM-C Resync on ONTAP S3 as well as other S3-compatible providers such as Wasabi and MinIO.

### Bring your own bucket for StorageGRID

When you create backup files in object storage for a working environment, by default, BlueXP backup and recovery creates the container (bucket or storage account) for the backup files in the object storage account that you configured. Previously, you could override this and specify your own container for Amazon S3, Azure Blob Storage, and Google Cloud Storage. With this release, you can now bring your own StorageGRID object storage container.

See [Create your own object storage container](#).

## 13 May 2025

This BlueXP backup and recovery release includes the following updates.

### SnapMirror to Cloud Resync for volume migrations

The SnapMirror to Cloud Resync feature streamlines data protection and continuity during volume migrations in NetApp environments. When a volume is migrated using SnapMirror Logical Replication (LRSE), from one on-premises NetApp deployment to another, or to a cloud-based solution such as Cloud Volumes ONTAP or Cloud Volumes Service, SnapMirror to Cloud Resync ensures that existing cloud backups remain intact and operational.

This feature eliminates the need for a time-consuming and resource-intensive re-baseline operation, enabling backup operations to continue post-migration. This feature is valuable in workload migration scenarios, supporting both FlexVols and FlexGroups, and is available starting with ONTAP version 9.16.1.

By maintaining backup continuity across environments, SnapMirror to Cloud Resync enhances operational efficiency and reduces the complexity of hybrid and multi-cloud data management.

For details on how to perform the resync operation, see [Migrate volumes using SnapMirror to Cloud Resync](#).

### Support for third-party MinIO object store (Preview)

BlueXP backup and recovery now extends its support to third-party object stores with a primary focus on MinIO. This new preview feature enables you to leverage any S3-compatible object store for your backup and recovery needs.

With this preview version, we hope to ensure robust integration with third-party object stores before the full functionality is rolled out. You are encouraged to explore this new capability and provide feedback to help enhance the service.



This feature should not be used in production.

### Preview mode limitations

While this feature is in preview, there are certain limitations:

- Bring Your Own Bucket (BYOB) is not supported.

- Enabling DataLock in the policy is not supported.
- Enabling Archival mode in the policy is not supported.
- Only on-premises ONTAP environments are supported.
- MetroCluster is not supported.
- Options to enable bucket-level encryption are not supported.

## Getting started

To begin using this preview feature, you must enable a flag on the BlueXP Connector. You can then enter the connection details of your MinIO third-party object store in the protection workflow by choosing **Third party Compatible** object store in the backup section.

## 16 April 2025

This BlueXP backup and recovery release includes the following updates.

### UI improvements

This release enhances your experience by simplifying the interface:

- The removal of the Aggregate column from the Volumes tables, along with the Snapshot Policy, Backup Policy, and Replication Policy columns from the Volume table in the V2 Dashboard, results in a more streamlined layout.
- Excluding non-activated working environments from the drop-down list makes the interface less cluttered, the navigation more efficient, and loading faster.
- While sorting on the Tags column is disabled, you can still view the tags, ensuring that important information remains easily accessible.
- The removal of labels on protection icons contributes to a cleaner look and decreases loading time.
- During the working environment activation process, a dialog box displays a loading icon to provide feedback until the discovery process is complete, enhancing transparency and confidence in the system's operations.

### Enhanced Volume Dashboard (Preview)

The Volume Dashboard now loads in under 10 seconds, providing a much faster and more efficient interface. This preview version is available to select customers, offering them an early look at these improvements.

### Support for third-party Wasabi object store (Preview)

BlueXP backup and recovery now extends its support to third-party object stores with a primary focus on Wasabi. This new preview feature enables you leverage any S3-compatible object store for your backup and recovery needs.

#### Getting started with Wasabi

To begin using third-party storage as an object store, you must enable a flag within the BlueXP Connector. Then, you can enter the connection details for your third-party object store and integrate it into your backup and recovery workflows.

#### Steps

1. SSH into your connector.
2. Go into the BlueXP backup and recovery cbs server container:

```
docker exec -it cloudmanager_cbs sh
```

3. Open the `default.json` file inside the `config` folder via VIM or any other editor:

```
vi default.json
```

4. Modify `allow-s3-compatible: false` to `allow-s3-compatible: true`.
5. Save the changes.
6. Exit from the container.
7. Restart the BlueXP backup and recovery cbs server container.

## Result

After the container is ON again, open the BlueXP backup and recovery UI. When you initiate a backup or edit a backup strategy, you will see the new provider "S3 Compatible" listed along with other backup providers of AWS, Microsoft Azure, Google Cloud, StorageGRID, and ONTAP S3.

## Preview mode limitations

While this feature is in preview, consider the following limitations:

- Bring Your Own Bucket (BYOB) is not supported.
- Enabling DataLock in a policy is not supported.
- Enabling Archival mode in a policy is not supported.
- Only on-premises ONTAP environments are supported.
- MetroCluster is not supported.
- Options to enable bucket-level encryption are not supported.

During this preview, we encourage you to explore this new feature and provide feedback about integration with third-party object stores before the full functionality is rolled out.

## 17 March 2025

This BlueXP backup and recovery release includes the following updates.

### SMB snapshot browsing

This BlueXP backup and recovery update resolved an issue that prevented customers from browsing local snapshots in an SMB environment.

### AWS GovCloud environment update

This BlueXP backup and recovery update fixed an issue that prevented the UI from connecting to an AWS GovCloud environment due to TLS certificate errors. The issue was resolved by using the BlueXP Connector

host name instead of the IP address.

### **Backup policy retention limits**

Previously, the BlueXP backup and recovery UI limited backups to 999 copies, while the CLI allowed more. Now, you can attach up to 4,000 volumes to a backup policy and include 1,018 volumes not attached to a backup policy. This update includes additional validations that prevent exceeding these limits.

### **SnapMirror Cloud resync**

This update ensures that SnapMirror Cloud resync cannot be started from BlueXP backup and recovery for unsupported ONTAP versions after a SnapMirror relationship has been deleted.

## **21 February 2025**

This BlueXP backup and recovery release includes the following updates.

### **High performance indexing**

BlueXP backup and recovery introduces an updated indexing feature that makes the indexing of data on the source working environment more efficient. The new indexing feature includes updates to the UI, improved performance of the Search & Restore method of restoring data, upgrades to global search capabilities, and better scalability.

Here's a breakdown of the improvements:

- **Folder consolidation:** The updated version groups folders together using names that include specific identifiers, making the indexing process smoother.
- **Parquet file compaction:** The updated version reduces the number of files used for indexing each volume, simplifying the process and removing the need for an extra database.
- **Scale-out with more sessions:** The new version adds more sessions to handle indexing tasks, speeding up the process.
- **Support for multiple index containers:** The new version uses multiple containers to better manage and distribute indexing tasks.
- **Split index workflow:** The new version divides the indexing process into two parts, enhancing efficiency.
- **Improved concurrency:** The new version makes it possible to delete or move directories at the same time, speeding up the indexing process.

### **Who benefits from this feature?**

The new indexing feature is available to all new customers.

### **How do you enable indexing?**

Before you can use the Search & Restore method of restoring data, 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 quick and efficient.

Enable indexing on the source working environment by selecting the "Enable Indexing" option when you are performing a Search & Restore.

For more information, see the documentation [how to restore ONTAP data using Search & Restore](#).

### **Supported scale**

The new indexing feature supports the following:

- Global search efficiency in less than 3 minutes
- Up to 5 billion files
- Up to 5000 volumes per cluster
- Up to 100K snapshots per volume
- Maximum time for baseline indexing is less than 7 days. The actual time will vary depending on your environment.

### **Global search performance improvements**

This release also includes enhancements to global search performance. You will now see progress indicators and more detailed search results, including the count of files and the time taken for the search. Dedicated containers for search and indexing ensure that global searches are completed in under five minutes.

Note these considerations related to global search:

- The new index is not performed on snapshots labeled as hourly.
- The new indexing feature works only on snapshots on FlexVols, and not for snapshots on FlexGroups.

## **13 February 2025**

This BlueXP backup and recovery release includes the following updates.

### **BlueXP backup and recovery Preview Release**

This Preview release of BlueXP backup and recovery provides an updated user interface that enables you to manage Microsoft SQL Server workloads using a 3-2-1 protection strategy, familiar in the BlueXP backup and recovery service. With this new version, you can back up these workloads to primary storage, replicate them to secondary storage, and back them up to cloud object storage.



This documentation is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

This version of BlueXP backup and recovery Preview 2025 includes the following updates.

- A redesigned user interface that provides a more intuitive experience for managing your backup and recovery tasks.
- The Preview version enables you to back up and restore Microsoft SQL Server databases. (Support for other workloads will be added in future releases.)
- This version integrates SnapCenter capabilities, enabling you to manage and protect your SnapCenter resources with a 3-2-1 data protection strategy from the BlueXP backup and recovery user interface.
- This version enables you to import SnapCenter workloads into BlueXP backup and recovery.

## **22 November 2024**

This BlueXP backup and recovery release includes the following updates.

## SnapLock Compliance and SnapLock Enterprise protection modes

BlueXP backup and recovery now can back up both FlexVol and FlexGroup on-premises 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.

See the complete list of supported volumes in the [Learn about BlueXP backup and recovery](#).

## Indexing for Search & Restore process on Volumes page

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 enables the Indexed Catalog to track the backup files for every volume. The Volumes page now shows the indexing status:

- Indexed: Volumes have been indexed.
- In-progress
- Not Indexed
- Indexing paused
- Error
- Not Enabled

## 27 September 2024

This BlueXP backup and recovery release includes the following updates.

### Podman support on RHEL 8 or 9 with Browse and Restore

BlueXP backup and recovery now supports file and folder restores on Red Hat Enterprise Linux (RHEL) versions 8 and 9 using the Podman engine. This applies to the BlueXP backup and recovery Browse and Restore method.

BlueXP Connector version 3.9.40 supports certain versions of Red Hat Enterprise Linux versions 8 and 9 for any manual installation of the Connector software on a RHEL 8 or 9 host, regardless of the location in addition to the operating systems mentioned in the [host requirements](#). These newer RHEL versions require the Podman engine instead of the Docker engine. Previously, BlueXP backup and recovery had two limitations when using the Podman engine. These limitations have been removed.

[Learn more about restoring ONTAP data from backup files.](#)

### Faster catalog indexing improves Search and Restore

This release includes an improved catalog index that completes the baseline indexing much faster. Faster indexing enables you to use the Search and Restore feature more quickly.

[Learn more about restoring ONTAP data from backup files.](#)

# Known limitations with BlueXP backup and recovery for Microsoft SQL Server workloads

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

## Clone lifecycle support

- Cloning from object storage is not supported.
- Bulk clone operations are not supported for on-demand clones.
- Choosing I-groups is not supported.
- Choosing QOS (maximum throughput) options is not supported.

## Standard deployment mode only

The BlueXP backup and recovery version works only in standard deployment mode, not restricted or private modes.

## Windows cluster name restriction

The Windows cluster name cannot contain an underscore (\_) character.

## SnapCenter migration issues

The migration of resources from SnapCenter into BlueXP backup and recovery has the following limitations.

For details about how SnapCenter policies migrate to BlueXP backup and recovery policies, see [Policies in SnapCenter compared to those in BlueXP backup and recovery](#).

## Resource group limitations

If all the resources in a resource group are protected and one of those resources is also protected outside of the resource group, the migration from SnapCenter is blocked.

**Workaround:** Protect the resource either in a resource group or by itself, but not in both.

## Resources with multiple policies using the same schedule tier not supported

You cannot have assign multiple policies that use the same schedule tier (for example, hourly, daily, weekly, etc.) to a resource. BlueXP backup and recovery will not import those resources from SnapCenter.

**Workaround:** Attach only one policy using the same schedule tier to a resource.

## Hourly policies must begin at the start of the hour

If you have a SnapCenter policy that repeats every hours, but the hours are not at intervals at the start of the hour, BlueXP backup and recovery will not import the resource. For example, policies with schedules of 1:30, 2:30, 3:30, etc. are not supported, while policies with schedules of 1:00, 2:00, 3:00, etc. are supported.

**Workaround:** Use a policy that repeats in 1-hour intervals starting at the top of the hour.



## Both daily and monthly policies attached to one resource not supported

If a SnapCenter policy repeats both in day and month intervals, BlueXP backup and recovery will not import the policy.

For example, you cannot attach a daily policy (with less than or equal to 7 days or greater than 7 days) to a resource and also attach a monthly policy to the same resource.

**Workaround:** Use a policy that uses a daily or a monthly interval, but not both.

## On demand backup policies not migrated

BlueXP backup and recovery does not import on demand backup policies from SnapCenter.

## Log-only backup policies not migrated

BlueXP backup and recovery does not import log-only backup policies from SnapCenter. If a SnapCenter policy includes log-only backups, BlueXP backup and recovery will not import the resource.

**Workaround:** Use a policy in SnapCenter that uses more than just log-only backups.

## Host mapping

SnapCenter does not have map storage clusters or SVMs for the resources to hosts, but BlueXP backup and recovery does. The on-premises ONTAP cluster or SVM will not be mapped to a host in BlueXP backup and recovery Preview version. Additionally, BlueXP does not support SVMs.

**Workaround:** Before importing resources from SnapCenter, create a working environment in BlueXP backup and recovery for all the on-premises ONTAP storage systems that are registered in on-premises SnapCenter. Then, import the resources for that cluster from SnapCenter into BlueXP backup and recovery.

## Schedules not in 15-minute intervals

If you have a SnapCenter policy schedule that starts at a certain time and repeats every so many minutes but the minutes are not in 15-minute intervals, BlueXP backup and recovery will not import the schedule.

**Workaround:** Use SnapCenter to adjust the policy so that it repeats in 15-minute intervals.

# Known limitations with BlueXP backup and recovery for ONTAP volumes

Known limitations identify functions that are not supported by this release of BlueXP backup and recovery, or that do not interoperate correctly with it. Review these limitations carefully.

- BlueXP backup and recovery backing up Cloud Volume ONTAP to an object store in the AWS China regions (including Beijing and Ningxia); however, you might need to manually modify Identity and Access Management (IAM) policies first.

For details about creating a Connector in AWS, refer to [Installing a Connector in AWS](#).

For additional details in a blog post, refer to [BlueXP backup and recovery Feature Blog May 2023](#).

- BlueXP backup and recovery does not support Microsoft Azure China regions.

For details about creating a Connector in Azure, refer to [Installing a Connector in Azure](#).

- BlueXP backup and recovery does not support backups of FlexCache volumes.

## Replication limitations for ONTAP volumes

- 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 for ONTAP volumes

- When backing up data, BlueXP backup and recovery will not maintain NetApp Volume Encryption (NVE). This means that encrypted data on the NVE volume will be decrypted while the data is being transferred to the destination and the encryption will not be maintained.

For an explanation about these encryption types, refer to [Configure NetApp Volume Encryption overview](#).

- If long-term retention snapshots are enabled on a SnapMirror destination volume using the schedule in the SnapMirror policy, snapshots are created directly on the destination volume. In this case, you should not back up those volumes using BlueXP backup and recovery because those snapshots will not be moved to object storage.
- When backing up data, BlueXP backup and recovery will not maintain NetApp Volume Encryption (NVE). This means that encrypted data on the NVE volume will be decrypted while the data is being transferred to the destination and the encryption will not be maintained.

For an explanation about these encryption types, refer to [Configure NetApp Volume Encryption overview](#).

- If long-term retention snapshots are enabled on a SnapMirror destination volume using the schedule in the SnapMirror policy, snapshots are created directly on the destination volume. In this case, you should not back up those volumes using BlueXP backup and recovery because those snapshots will not be moved to object storage.
- 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.
  - SVM-DR and volume backup and recovery work fully independently when the backup is taken from either the source or destination. The only restriction is that SVM-DR does not replicate the SnapMirror cloud relationship. In the DR scenario when the SVM goes online in the secondary location, you must manually update the SnapMirror cloud relationship.
- 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 for ONTAP volumes

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 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.
- If you back up a DP volume and then decide to break the SnapMirror relationship to that volume, you cannot restore files to that volume unless you also delete the SnapMirror relationship or reverse the SnapMirror direction.
- 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 ONTAP volumes, Microsoft SQL Server instances and databases, VMware workloads and Kubernetes workloads (Preview).



Some of this documentation is provided as a technology preview. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

### What you can do with BlueXP backup and recovery

Use BlueXP backup and recovery to accomplish the following goals:

- ONTAP volume workloads:
  - Create local snapshots, replicate to secondary storage, and back up ONTAP volumes from on-premises ONTAP or Cloud Volumes ONTAP systems to object storage in your public or private cloud account.
  - Create block-level, incremental forever backups that are stored on another ONTAP cluster and in object storage in the cloud.
- Microsoft SQL Server workloads:
  - Back up Microsoft SQL Server instances and databases from on-premises ONTAP, Cloud Volumes ONTAP, or Amazon FSx for NetApp ONTAP.
  - Restore Microsoft SQL Server databases.
  - Clone Microsoft SQL Server databases.
- VM workloads:
  - 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.
  - Restore virtual machines data from the cloud back to the on-premises vCenter with BlueXP backup and recovery.
- Kubernetes workloads (Preview):
  - Manage and protect your Kubernetes applications and resources all in one place.
  - Use protection policies to structure your incremental backups.
  - Restore applications and resources to the same or different clusters and namespaces.

### Benefits of using BlueXP backup and recovery

BlueXP backup and recovery provides the following benefits:

- **Efficient:** BlueXP backup and recovery performs block-level, incremental-forever replication, which significantly reduces the amount of data that's replicated and stored. This helps to minimize network traffic and storage costs.

- **Secure:** BlueXP backup and recovery encrypts data in transit and at rest, and it uses secure communication protocols to protect your data.
- **Cost-effective:** BlueXP backup and recovery uses the lowest-cost storage tiers available in your cloud account, which helps to reduce costs.
- **Automated:** BlueXP backup and recovery automatically generates backups based on a predefined schedule, which helps to ensure that your data is protected.
- **Flexible:** BlueXP backup and recovery enables you to restore data to the same or different working environment, which provides flexibility in data recovery.

## Cost

NetApp doesn't charge you for using the trial version. However, you are responsible for the costs associated with the cloud resources that you use, such as storage and data transfer costs.

There are two types of costs associated with using the backup-to-object feature of BlueXP backup and recovery with ONTAP systems:

- Resource charges
- Service charges

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.

Because 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 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 that are backed up to object storage. This capacity is also known as Front-End Terabytes (FETB).



For Microsoft SQL Server, charges apply when you initiate the replication of snapshots to a secondary ONTAP target or object storage.

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 as a free trial. You can use the service without a license key for a limited time.

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

- **Bring your own license (BYOL):** A license purchased from NetApp that can be used with any cloud provider.
- **Pay as you go (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 organization or account.

[Learn how to set up 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.

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

### Annual contract

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

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

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

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

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

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

## Supported data sources, working environments, and backup targets

### Workload data sources supported

The service protects the following workloads:

- ONTAP volumes
- Microsoft SQL Server instances and databases for physical, VMware Virtual Machine File System (VMFS), and VMware Virtual Machine Disk (VMDK) NFS
- VMware datastores
- Kubernetes workloads (Preview)

### Working environments supported

- On-premises ONTAP SAN (iSCSI protocol) and NAS (using NFS and CIFS protocols) with ONTAP version 9.8 and greater
- Cloud Volumes ONTAP 9.8 or greater for AWS (using SAN and NAS)
- Cloud Volumes ONTAP 9.8 or greater for Microsoft Azure (using SAN and NAS)
- Amazon FSx for NetApp ONTAP

### Backup targets supported

- Amazon Web Services (AWS) S3
- Microsoft Azure Blob
- StorageGRID
- ONTAP S3

## BlueXP backup and recovery uses the Plug-in for Microsoft SQL Server

BlueXP backup and recovery installs the Plug-in for Microsoft SQL Server on the server that hosts Microsoft SQL Server. The Plug-in is a host-side component that enables application-aware data protection management of Microsoft SQL Server databases and instances.

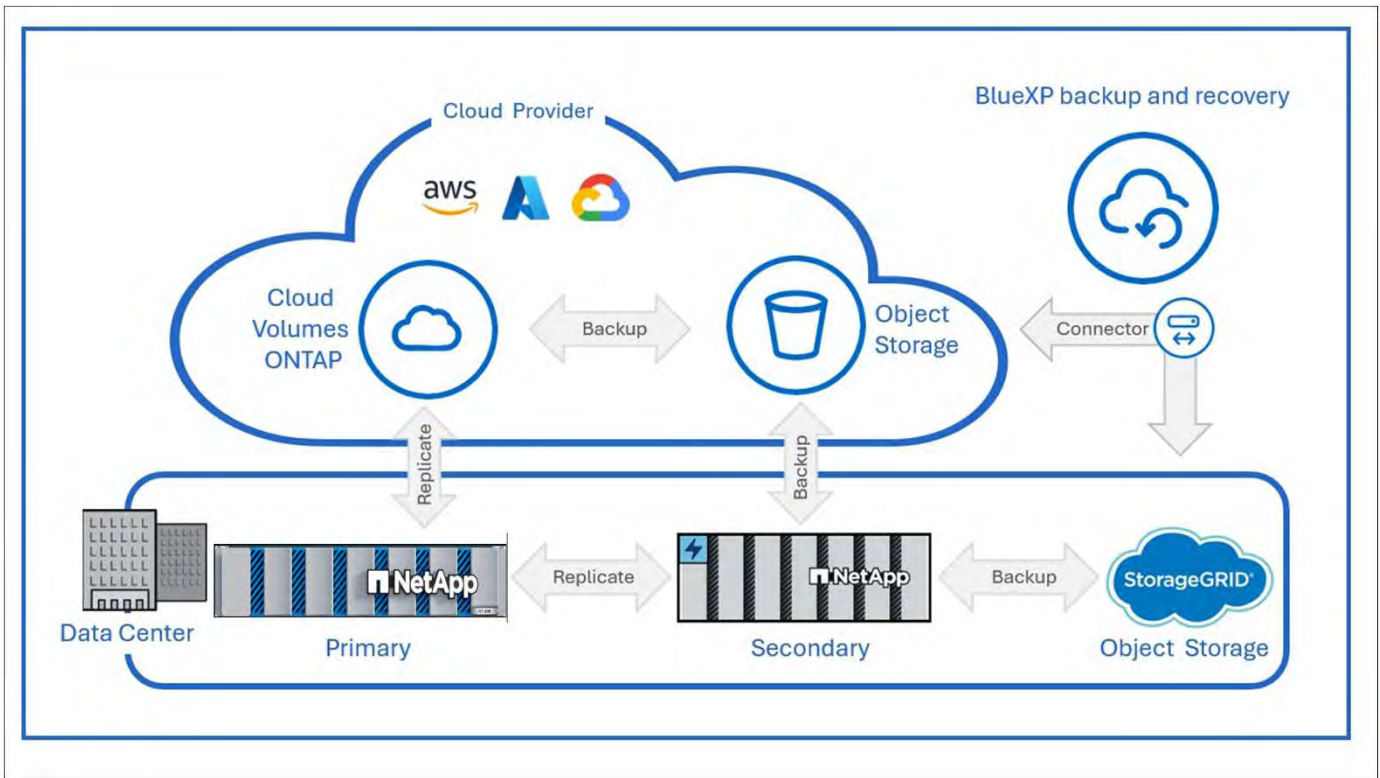
## How BlueXP backup and recovery works

When you enable BlueXP backup and recovery, the service performs a full backup of your data. After the initial



backup, all additional backups are incremental. This keeps network traffic to a minimum.

The following image shows the relationship among components.



Primary to object storage is also supported, not just from secondary storage to object storage.

### Where backups reside in object store locations

Backup copies are stored in an object store that BlueXP creates in your cloud account. There's one object store per cluster or 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 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.

### Backup copies are associated with your BlueXP organization

Backup copies are associated with the BlueXP organization in which the BlueXP Connector resides. [Learn about BlueXP identity and access management.](#)

If you have multiple Connectors in the same BlueXP organization, each Connector displays the same list of backups.

## Terms that might help you with BlueXP backup and recovery

You might benefit by understanding some terminology related to protection.

- **Protection:** Protection in BlueXP backup and recovery means ensuring that snapshots and immutable backups occur on a regular basis to a different security domain using protection policies.
- **Workload:** A workload in BlueXP backup and recovery can include Microsoft SQL Server instances and databases, VMware datastores, or ONTAP volumes.

## BlueXP backup and recovery prerequisites

Get started with BlueXP backup and recovery by verifying the readiness of your operational environment, BlueXP Connector, and BlueXP account. To use BlueXP backup and recovery, you'll need these prerequisites.

### For ONTAP 9.8 and later

An ONTAP One license must be enabled on the on-premises ONTAP instance.

### Prerequisites for backups to object storage


To use object storage as backup targets, you need an account with AWS S3, Microsoft Azure Blob, StorageGRID, or ONTAP and the appropriate access permissions configured.

See [Set up backup destinations before you use BlueXP backup and recovery](#) for details about how to set up the backup destinations.

### Microsoft SQL Server workload requirements

To use BlueXP backup and recovery for SQL Server workloads, you need the following host system, space, and sizing prerequisites.

Item	Requirements
Operating systems	Microsoft Windows For the latest information about supported versions, see the <a href="#">NetApp Interoperability Matrix Tool</a> .
Microsoft SQL Server versions	Version 2012 and later are supported for VMware Virtual Machine File System (VMFS) and VMware Virtual Machine Disk (VMDK) NFS.

Item	Requirements
SnapCenter Server version	<p>SnapCenter Server version 5.0 or greater is required if you are going to import your existing data from SnapCenter into BlueXP backup and recovery.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>If you already have SnapCenter, first check to be sure you've met the prerequisites before importing from SnapCenter. See <a href="#">Prerequisites for importing resources from SnapCenter</a>.</p> </div>
Minimum RAM for the plug-in on the SQL Server host	1 GB
Minimum install and log space for the plug-in on the SQL Server host	<p>5 GB</p> <p>Allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of backups performed and the frequency of data protection operations. If there is not sufficient space, the logs will not be created for the operations.</p>
Required software packages	<ul style="list-style-type: none"> <li>• ASP.NET Core Runtime 8.0.12 Hosting Bundle (and all subsequent 8.0.x patches)</li> <li>• PowerShell Core 7.4.2</li> </ul> <p>For the latest information about supported versions, see the <a href="#">NetApp Interoperability Matrix Tool</a>.</p>

## Kubernetes workload requirements

You need specific requirements to discover Kubernetes resources and protect your Kubernetes applications.

For BlueXP requirements, refer to [In BlueXP](#).

- A primary ONTAP system (ONTAP 9.16.1 or later)
- A Kubernetes cluster - Supported Kubernetes distributions and versions include:
  - Anthos On-Prem (VMware) and Anthos on bare metal 1.16
  - Kubernetes 1.27 - 1.33
  - OpenShift 4.10 - 4.18
  - Rancher Kubernetes Engine 2 (RKE2) v1.26.7+rke2r1, v1.28.5+rke2r1
- NetApp Trident 24.10 or later
- NetApp Trident protect 25.07 or later (installed during Kubernetes workload discovery)
- NetApp Trident protect Connector 25.07 or later (installed during Kubernetes workload discovery)
  - Make sure that TCP port 443 is unfiltered in the outbound direction between the Kubernetes cluster, the Trident protect Connector, and the Trident protect proxy.

## In BlueXP

- A BlueXP user should have the required role and privileges to perform operations on Microsoft SQL Server and Kubernetes workloads. To discover the resources, you must have the BlueXP backup and recovery role of Super admin. See [BlueXP backup and recovery role-based access to features](#) for details about the roles and permissions required to perform operations in BlueXP backup and recovery.
- A BlueXP organization with at least one active BlueXP Connector that connects to on-premises ONTAP clusters or Cloud Volumes ONTAP. Refer to the **Initial Preview setup process** below.
- At least one BlueXP working environment with a NetApp on-premises ONTAP or Cloud Volumes ONTAP cluster.
- A BlueXP Connector

Refer to [Learn how to configure a BlueXP Connector](#) and [standard BlueXP requirements](#).

- The Preview version requires the Ubuntu 22.04 LTS operating system for the Connector.

## Set up BlueXP

The next step is to set up BlueXP and the BlueXP backup and recovery service.

Review [standard BlueXP requirements](#).

### Create a BlueXP Connector

You should reach out to your NetApp Product Team to try out this service. Then, when you use the BlueXP Connector, it will include the appropriate capabilities for the service.

To create a Connector in BlueXP before using the service, refer to the BlueXP documentation that describes [how to create a BlueXP Connector](#).

### Where to install the BlueXP Connector

To complete a restore operation, the Connector can be installed in the following locations:

- For Amazon S3, the Connector can be deployed on your premises.
- For Azure Blob, the Connector can be deployed on your premises.
- 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



References to "on-premises ONTAP systems" includes FAS and AFF systems.

## 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 to **NetApp Intelligent Services** 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 organization or 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 to **NetApp Intelligent Services**. 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.

## End the free trial

If you want to continue using BlueXP backup and recovery after the free trial ends, you must set up a paid subscription. You can do this from the BlueXP interface by navigating to the billing section and selecting a subscription plan that fits your needs. If you don't want to continue using BlueXP backup and recovery, you can end the free trial.

When you end the free trial without subscribing to a paid plan, your data is automatically deleted 60 days after the free trial ends. You can optionally have the system delete your data immediately.

## Steps

1. From the BlueXP backup and recovery landing page, select **View free trial**.

**QUESTION TO REVIEWERS:** How do users get to the Landing page if they're on other BR pages?

2. Select **End free trial**.
3. Select **Delete data immediately after ending my free trial** to delete your data immediately.
4. Type **end trial** in the box.
5. Select **End** to confirm.

## 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 to **NetApp Intelligent Services** in the

Marketplace 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.

[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 Marketplace offering for NetApp Intelligent Services for pricing details.](#)
- Azure: [Go to the Marketplace offering for NetApp Intelligent Services for pricing details.](#)
- Google Cloud: [Go to the Marketplace offering for NetApp Intelligent Services 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 you use 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 you use 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 you use 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 organization or 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.



[Learn about adding licenses with digital wallet.](#)

## Set up backup destinations before you use BlueXP backup and recovery

Before you use BlueXP backup and recovery, perform a few steps to set up backup destinations.

Before you begin, review [prerequisites](#) to ensure that your environment is ready.

### Prepare the backup destination

Prepare one or more of the following backup destinations:

- NetApp StorageGRID.

Refer to [Discover StorageGRID](#).

Refer to [StorageGRID documentation](#) for details about StorageGRID.

- Amazon Web Services. Refer to [Amazon S3 documentation](#).

Do the following to prepare AWS as a backup destination:

- Set up an account in AWS.
- Configure S3 permissions in AWS, listed in the next section.
- For details about managing your AWS storage in BlueXP, refer to [Manage your Amazon S3 buckets](#).

- Microsoft Azure.

- Refer to [Azure NetApp Files documentation](#).
- Set up an account in Azure.
- Configure [Azure permissions](#) in Azure.
- For details about managing your Azure storage in BlueXP, refer to [Manage your Azure storage accounts](#).

After you configure options in the backup destination itself, you will later configure it as a backup destination in the BlueXP backup and recovery service. For details about how to configure the backup destination in BlueXP backup and recovery, refer to [Discover backup targets](#).

### Set up S3 permissions

You'll need to configure two sets of AWS S3 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. Ensure that the Connector has the required permissions. For details, see [BlueXP policy permissions](#).





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"
    }
  ]
}

```

# Log in to BlueXP backup and recovery

You use NetApp BlueXP to log in to the BlueXP backup and recovery service.

BlueXP backup and recovery uses role-based access control (RBAC) to govern the access that each user has to specific actions.

For details about the actions that each role can perform, see [BlueXP backup and recovery user roles](#).

To log in to BlueXP, you can use your NetApp Support Site credentials or you can sign up for a NetApp cloud login using your email and a password. [Learn more about logging in](#).

## Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery restore admin role. [Learn about BlueXP access roles for all services](#).

If this is your first time accessing BlueXP backup and recovery and to add a Connector, you must have the Organization admin or the Backup and Recovery super admin role.

## Steps

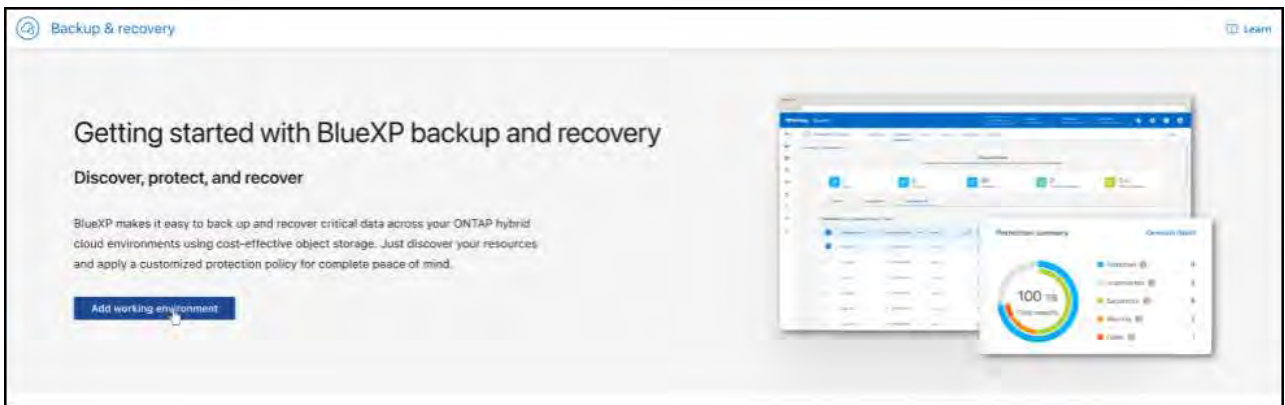
1. Open a web browser and go to the [BlueXP console](#).

The NetApp BlueXP login page appears.

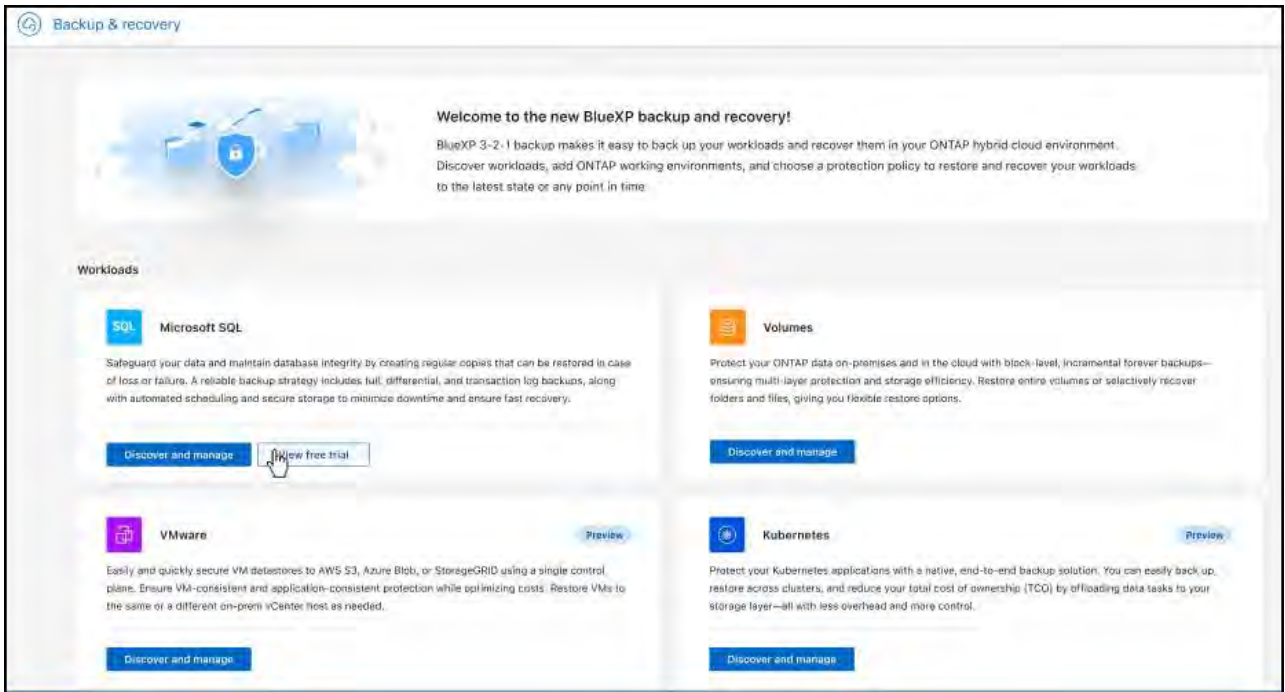
2. Log in to BlueXP.

3. From the BlueXP left navigation, select **Protection > Backup and recovery**.

- If this is your first time logging in to this service and you don't yet have a working environment, the "Welcome to the new BlueXP backup and recovery" landing page appears and shows an option to **Add working environment**. For details about adding a working environment to BlueXP, see [Getting started with BlueXP standard mode](#).



- If this is your first time logging in to this service, you already have a working environment in BlueXP, but you haven't started the free trial, the "Welcome to the new BlueXP backup and recovery" landing page appears and shows an option to **View free trial**.



- If this is your first time logging in to this service and you already have a working environment in BlueXP, but haven't discovered any resources, the "Welcome to the new BlueXP backup and recovery" landing page appears and shows an option to **Discover and manage**.

4. If you haven't done so already, select the **Discover and manage** option.

For Microsoft SQL Server workloads, refer to [Discover Microsoft SQL Server workloads](#).

## Discover offsite backup targets in BlueXP backup and recovery

Complete a few steps to discover or manually add offsite backup targets in BlueXP backup and recovery.

### Discover a backup target

Before you use BlueXP backup and recovery, you should configure your backup targets of Amazon Web Services (AWS) S3, Microsoft Azure Blob Storage, Google Cloud Storage, or StorageGRID.

You can discover these targets automatically or manually add them.

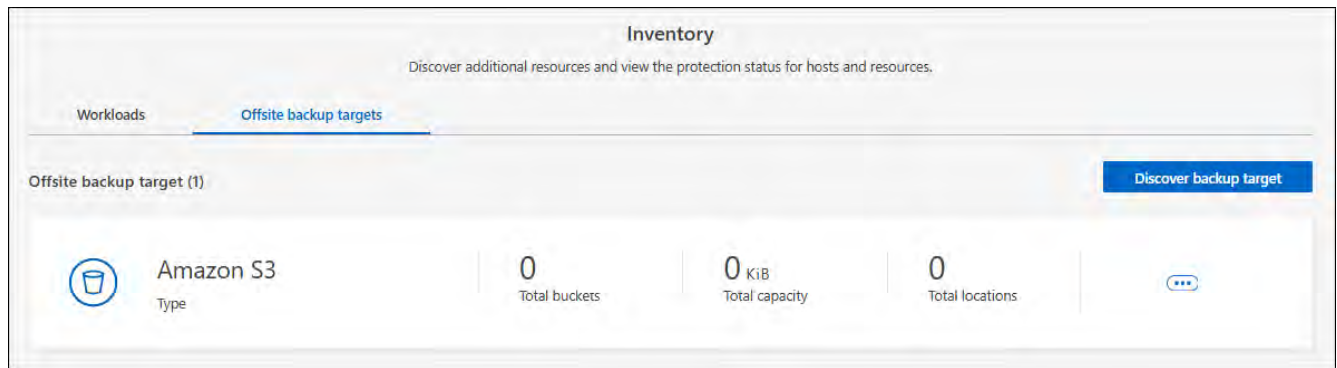
Provide the credentials needed to access the storage account system. These credentials are used to discover the workloads that you want to back up.

### Before you begin

To add an offsite backup target, at least one workload has to be discovered.

### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the **Offsite backup targets** tab.




3. Select **Discover backup target**.
4. Select one of the backup target types: **Amazon Web Services (AWS) S3**, **Microsoft Azure Blob Storage**, **StorageGRID** or **ONTAP S3**.
5. In the **Choose credentials location** section, choose the location where the credentials reside and then choose how to associate the credentials.
6. Select **Next**.
7. Enter the credentials information. The information differs depending on the type of backup target you selected and the credentials location that you chose.
  - For AWS:
    - **Credential name:** Enter the AWS credential name.
    - **Access key:** Enter the AWS secret.
    - **Secret key:** Enter the AWS secret key.
  - For Azure:
    - **Credential name:** Enter the Azure Blob Storage credential name.
    - **Client secret:** Enter the Azure Blob Storage client secret.
    - **Application (client) ID:** Select the Azure Blob Storage application ID.
    - **Directory tenant ID:** Enter the Azure Blob Storage tenant ID.
  - For StorageGRID:
    - **Credential name:** Enter the StorageGRID credential name.
    - **Gateway Node FQDN:** Enter a FQDN name for StorageGRID.
    - **Port:** Enter the port number for StorageGRID.
    - **Access key:** Enter the StorageGRID S3 access key.
    - **Secret key:** Enter the StorageGRID S3 secret key.
  - For ONTAP S3:
    - **Credential name:** Enter the ONTAP S3 credential name.
    - **Gateway Node FQDN:** Enter a FQDN name for ONTAP S3.
    - **Port:** Enter the port number for ONTAP S3.
    - **Access key:** Enter the ONTAP S3 access key.
    - **Secret key:** Enter the ONTAP S3 secret key.
8. Select **Discover**.

## Add a bucket for a backup target

Rather than have BlueXP backup and recovery discover buckets automatically, you can manually add a bucket to an offsite backup target.

### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select **Offsite backup targets**.
3. Select the target and on the right, select the **Actions**  icon and select **Add bucket**.
4. Enter the bucket information. The information differs depending on the type of backup target you selected.
  - For AWS:
    - **Bucket name**: Enter the name of the S3 bucket. The prefix of "netapp-backup" is a required prefix and is automatically added to the name you provide.
    - **AWS account**: Enter the AWS account name.
    - **Bucket region**: Enter the AWS region for the bucket.
    - **Enable S3 Object Lock**: Select this option to enable S3 Object Lock for the bucket. S3 Object Lock prevents objects from being deleted or overwritten for a specified retention period, providing an additional layer of data protection. You can enable this only when you are creating a bucket and you cannot turn it off later.
      - **Governance mode**: Select this option to enable governance mode for the S3 Object Lock bucket. Governance mode enables you to protect objects from being deleted or overwritten by most users, but allows certain users to alter the retention settings.
      - **Compliance mode**: Select this option to enable compliance mode for the S3 Object Lock bucket. Compliance mode prevents any user, including the root user, from altering the retention settings or deleting objects until the retention period expires.
    - **Versioning**: Select this option to enable versioning for the S3 bucket. Versioning enables you to keep multiple versions of objects in the bucket, which can be useful for backup and recovery purposes.
    - **Tags**: Select tags for the S3 bucket. Tags are key-value pairs that can be used to organize and manage your S3 resources.
    - **Encryption**: Select the type of encryption for the S3 bucket. The options are either AWS S3-managed keys or AWS Key Management Service key. If you select AWS Key Management Service keys, you must provide the key ID.
  - For Azure:
    - **Subscription**: Select the name of the Azure Blob Storage container.
    - **Resource group**: Select the name of the Azure resource group.
    - **Instance details**:
      - **Storage account name**: Enter the name of the Azure Blob Storage container.
      - **Azure region**: Enter the Azure region for the container.
      - **Performance type**: Select the performance type of either standard or premium for the Azure Blob Storage container indicating the level of performance required.
      - **Encryption**: Select the type of encryption for the Azure Blob Storage container. The options are either Microsoft-managed keys or customer-managed keys. If you select customer-managed keys, you must provide the key vault name and key name.


- For StorageGRID:
  - **Backup target name:** Select the name of the StorageGRID bucket.
  - **Bucket name:** Enter the name of the StorageGRID bucket.
  - **Region:** Enter the StorageGRID region for the bucket.
  - **Enable versioning:** Select this option to enable versioning for the StorageGRID bucket. Versioning enables you to keep multiple versions of objects in the bucket, which can be useful for backup and recovery purposes.
  - **Object locking:** Select this option to enable object locking for the StorageGRID bucket. Object locking prevents objects from being deleted or overwritten for a specified retention period, providing an additional layer of data protection. You can enable this only when you are creating a bucket and you cannot turn it off later.
  - **Capacity:** Enter the capacity for the StorageGRID bucket. This is the maximum amount of data that can be stored in the bucket.
- For ONTAP S3:
  - **Backup target name:** Select the name of the ONTAP S3 bucket.
  - **Bucket target name:** Enter the name of the ONTAP S3 bucket.
  - **Capacity:** Enter the capacity for the ONTAP S3 bucket. This is the maximum amount of data that can be stored in the bucket.
  - **Enable versioning:** Select this option to enable versioning for the ONTAP S3 bucket. Versioning enables you to keep multiple versions of objects in the bucket, which can be useful for backup and recovery purposes.
  - **Object locking:** Select this option to enable object locking for the ONTAP S3 bucket. Object locking prevents objects from being deleted or overwritten for a specified retention period, providing an additional layer of data protection. You can enable this only when you are creating a bucket and you cannot turn it off later.

5. Select **Add**.

## Change credentials for a backup target

Enter the credentials needed to access the backup target.

### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select **Offsite backup targets**.
3. Select the target and on the right, select the **Actions**  icon and select **Change credentials**.
4. Enter the new credentials for the backup target. The information differs depending on the type of backup target you selected.
5. Select **Done**.

## Switch to different BlueXP backup and recovery workloads


You can switch among the different BlueXP backup and recovery workloads. Some workloads use a different UI.

**How do you know which UI you are using?**

The taskbar for Microsoft SQL Server and Kubernetes (Preview) workloads looks like this:



The menu bar for ONTAP volumes and VMware workloads looks like this:




## Switch to a different workload

You can switch to a different workload in the BlueXP backup and recovery UI.

### Steps

1. From the BlueXP left navigation, select **Protection > Backup and recovery**.
2. From the top right corner of the page, select the **Switch workload** drop-down list.
3. Select the workload that you want to switch to.



The page refreshes and shows the selected workload in the appropriate UI.

## Configure BlueXP backup and recovery settings

After you set up BlueXP, configure the backup and recovery settings, which include adding credentials for host resources, importing SnapCenter resources, configuring log directories, and configuring VMware vCenter settings. You should do this before you actively start backing up and recovering your data.

- [Add credentials for host resources](#) for the Windows and SQL Server hosts that you imported from SnapCenter and add credentials. (Microsoft SQL Server workloads only)
- [Maintain VMware vCenter settings](#).
- [Import and manage SnapCenter host resources](#). (Microsoft SQL Server workloads only)
- [Configure log directories in snapshots for Windows hosts](#).

### Required BlueXP role

Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin. Learn about [Backup and recovery roles and privileges](#). Learn about [BlueXP access roles for all services](#).

## Add credentials for host resources

Add credentials for the host resources that you want to import from SnapCenter. Host credentials are used to discover new workloads and apply backup policies.



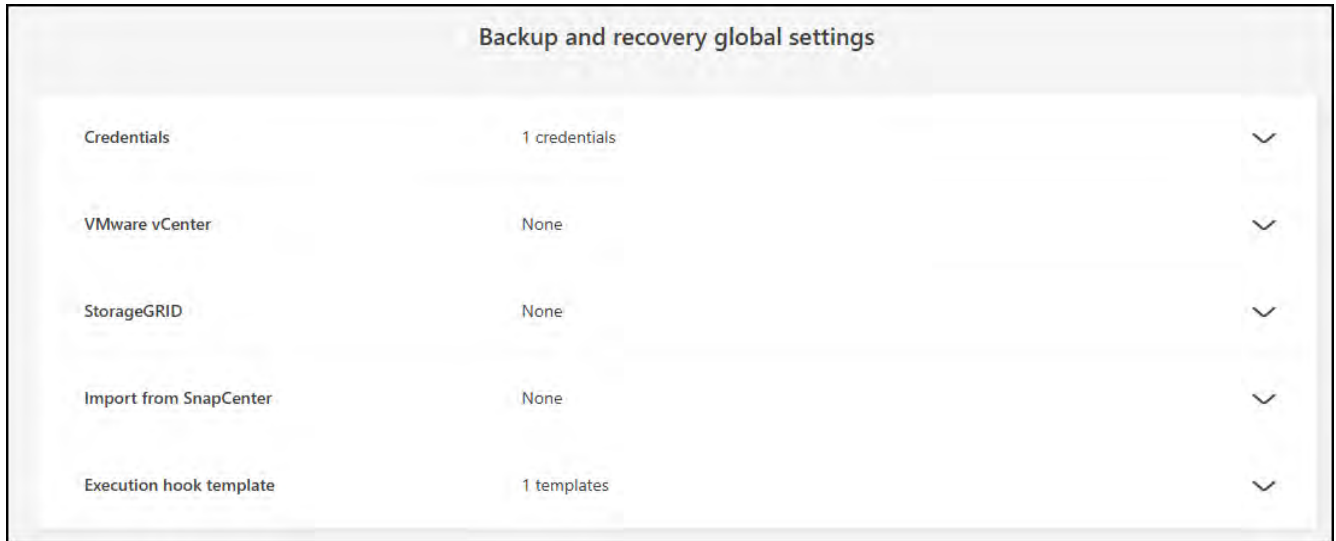
If you don't already have credentials, you can create them. These credentials must have required permissions to access and manage the host workloads.

You need to configure the following types of credentials:

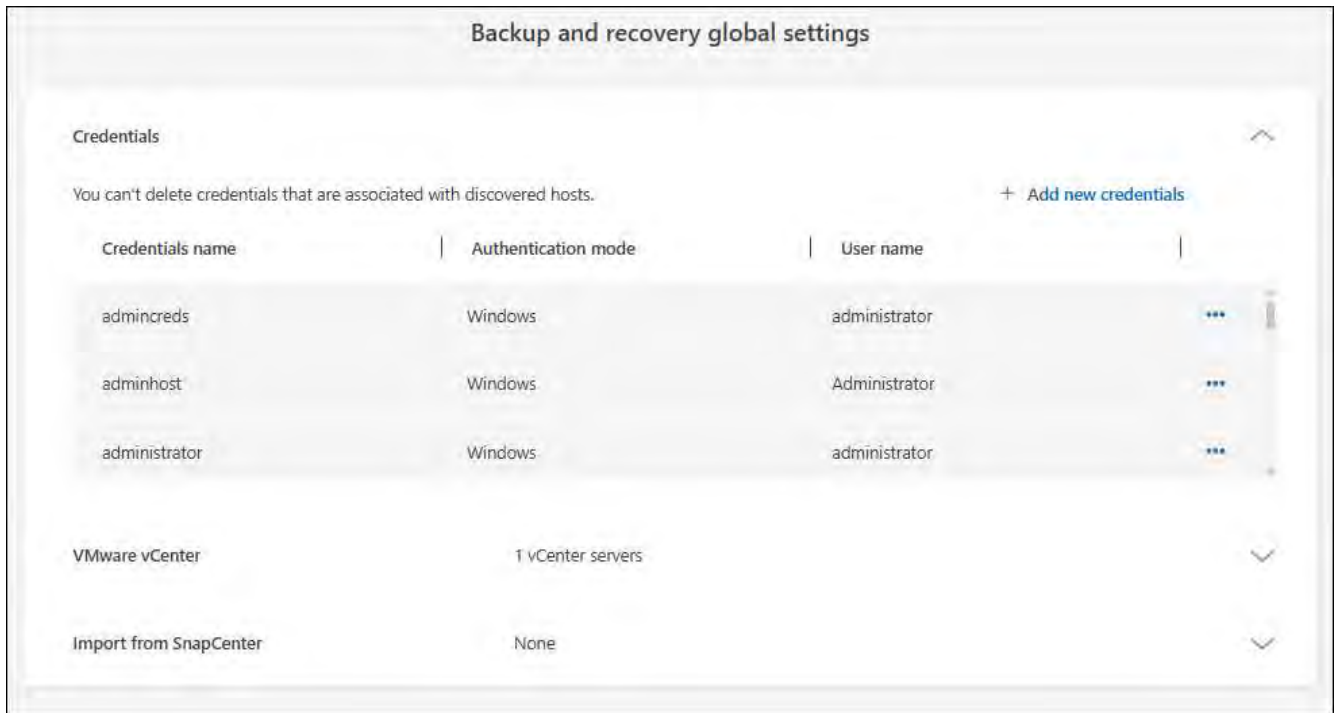
- Microsoft SQL Server credentials
- SnapCenter Windows host credentials

### Steps

1. From the BlueXP backup and recovery menu, select **Settings**.



2. Select the down arrow for **Credentials**.



3. Select **Add new credentials**.

Settings > Add credentials

### Add credentials

Add credentials to discover the hosts you want to access for backup purposes.

**Credentials name** ⓘ

**Authentication mode**

**Connectors**

**Domain and user name** ⓘ

**Password**

Add
Close

4. Enter information for the credentials. Different fields appear depending on the Authentication mode you select. Select the Information **i** for more information about the fields.

- **Credentials name:** Enter a name for the credentials.
- **Authentication mode:** Select **Windows** or **Microsoft SQL**.



You need to enter credentials for both Windows and Microsoft SQL Server, so you'll need to add two sets of credentials.

5. If you selected **Windows**:

- **Connector:** Enter the BlueXP Connector IP address.
- **Domain and user name:** Enter the NetBIOS or domain FQDN and user name for the credentials.
- **Password:** Enter the password for the credentials.

6. If you selected **Microsoft SQL**:

- **Host:** Select a discovered SQL Server host address.
- **SQL Server instance:** Select a discovered SQL Server instance.

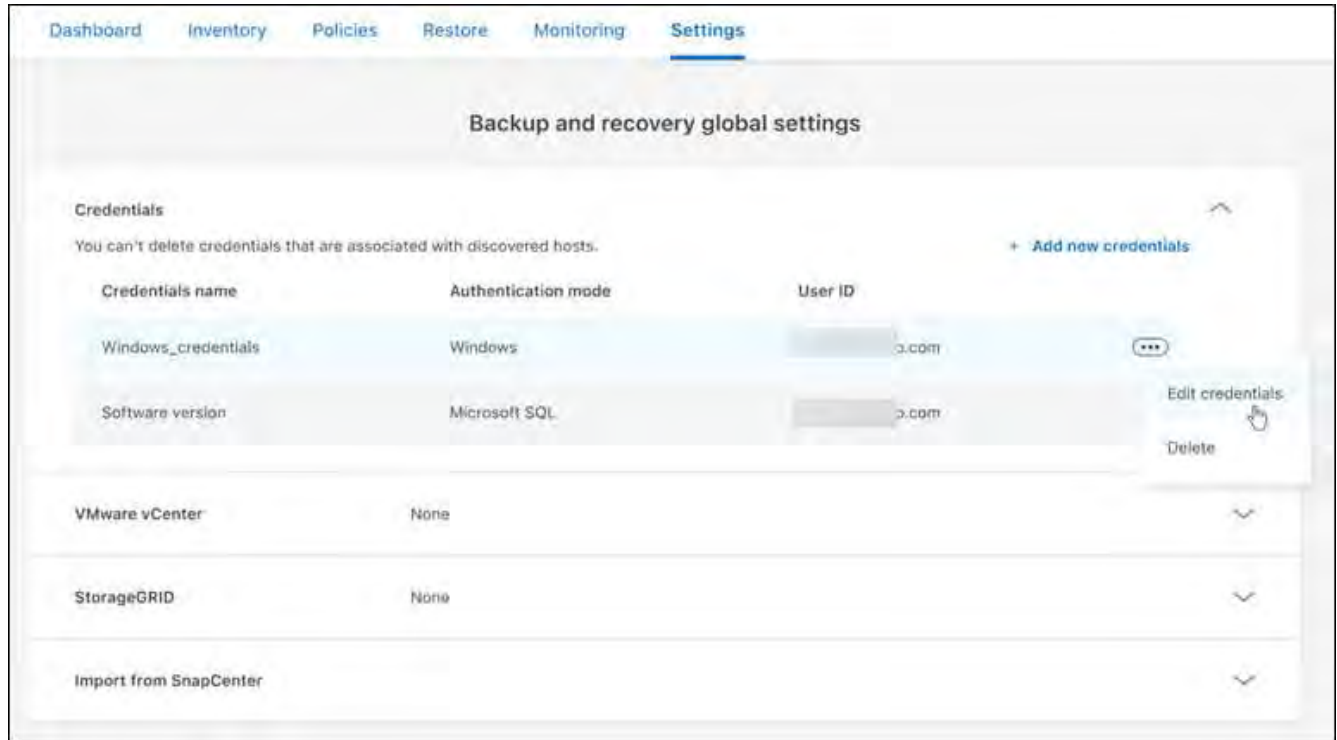
7. Select **Add**.

## Edit credentials for host resources

You can later edit the password for the host resources that you imported from SnapCenter.

### Steps

1. From the BlueXP backup and recovery menu, select **Settings**.
2. Select the down arrow to expand the **Credentials** section.



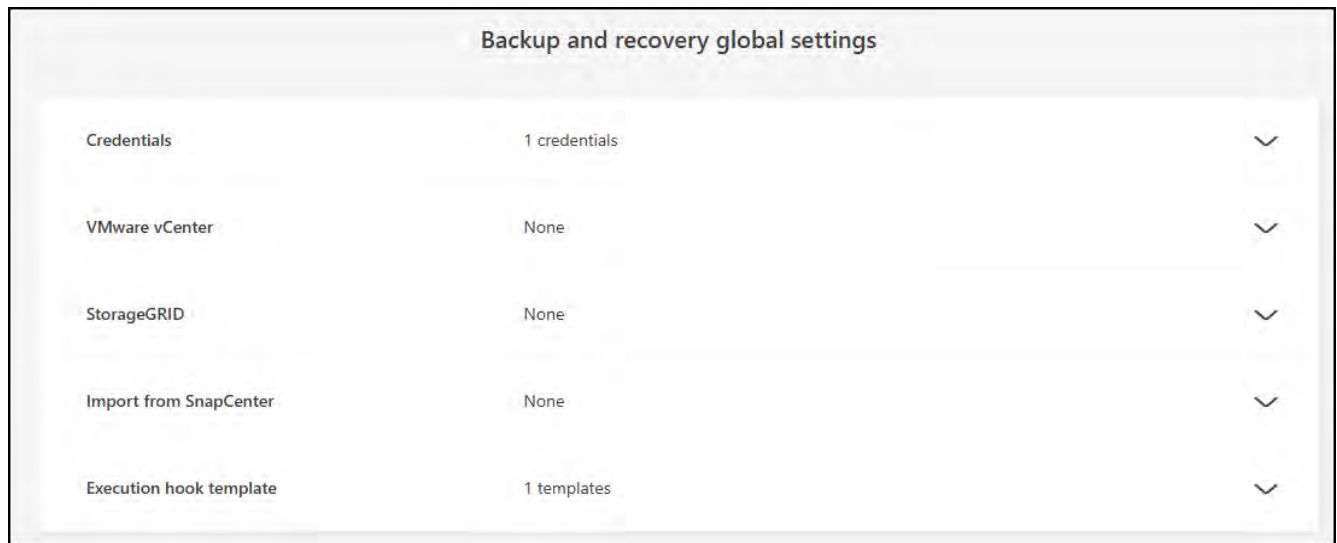
3. Select the Actions icon **⋮** > **Edit credentials**.
  - **Password**: Enter the password for the credentials.
4. Select **Save**.

## Maintain VMware vCenter settings

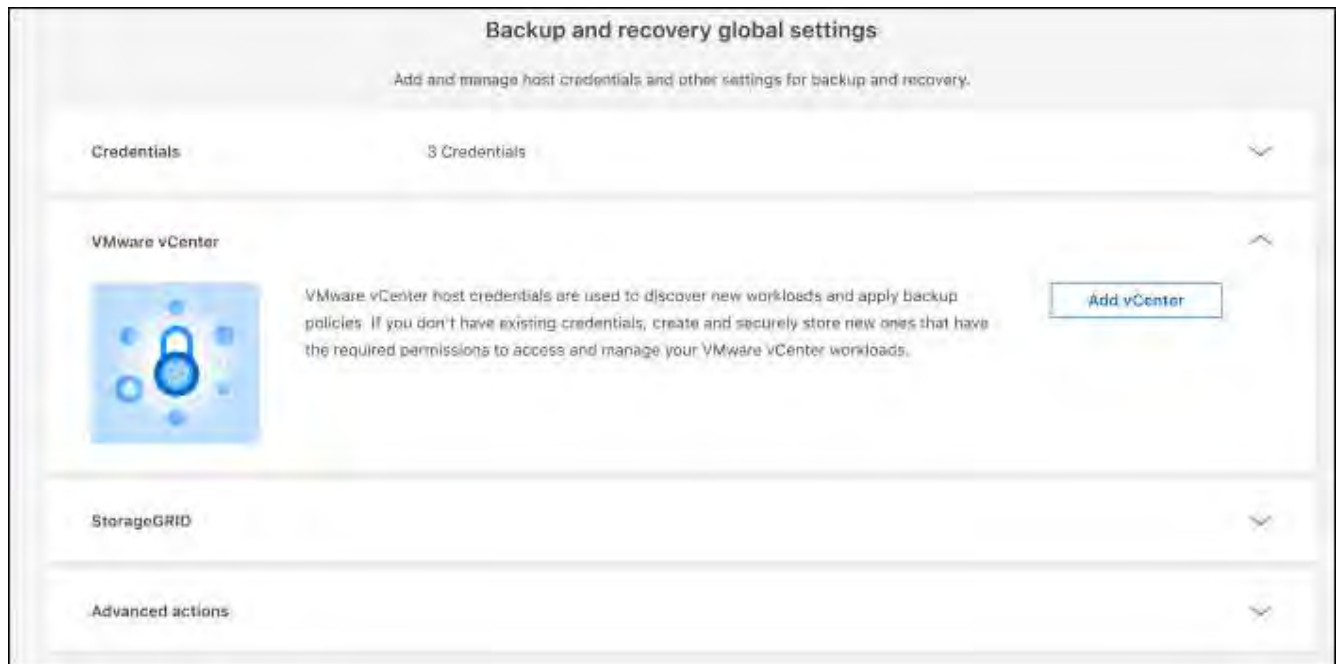
Provide the VMware vCenter credentials to discover the VMware vCenter Server workloads that you want to back up. If you don't have existing credentials, you can create them with the required permissions to access and manage the VMware vCenter Server workloads.

### Steps

1. From the BlueXP backup and recovery menu, select **Settings**.



2. Select the down arrow to expand the **VMware vCenter** section.



3. Select **Add vCenter**.

4. Enter the VMware vCenter Server information.

- **vCenter FQDN or IP address:** Enter a FQDN name or the IP address for the VMware vCenter Server.
- **Username and Password:** Enter the username and password for the VMware vCenter Server.
- **Port:** Enter the port number for the VMware vCenter Server.
- **Protocol:** Select **HTTP** or **HTTPS**.

5. Select **Add**.

## Import and manage SnapCenter host resources

If you previously used SnapCenter to back up your resources, you can import and manage those resources in BlueXP backup and recovery. With this option, you can import SnapCenter Server information to register

multiple Snapcenter servers and discover the database workloads.

This is a two-part process:

- Import SnapCenter Server application and host resources
- Manage selected SnapCenter host resources

### Import SnapCenter Server application and host resources

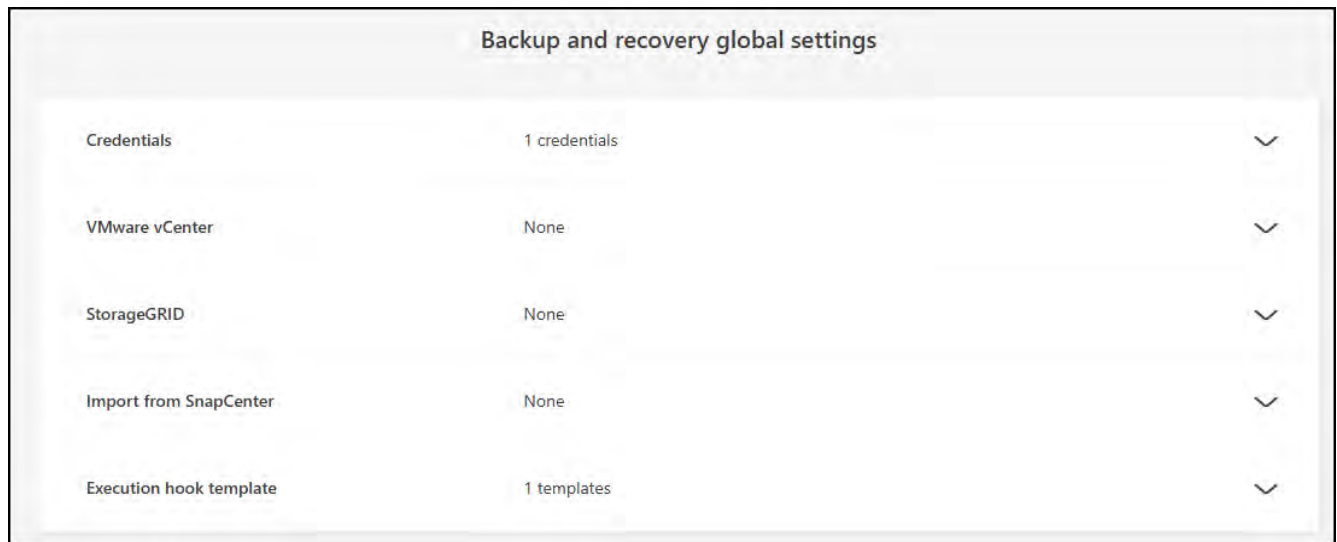
This first step imports host resources from SnapCenter and displays those resources in the BlueXP backup and recovery Inventory page. At that point, the resources are not yet managed by BlueXP backup and recovery.



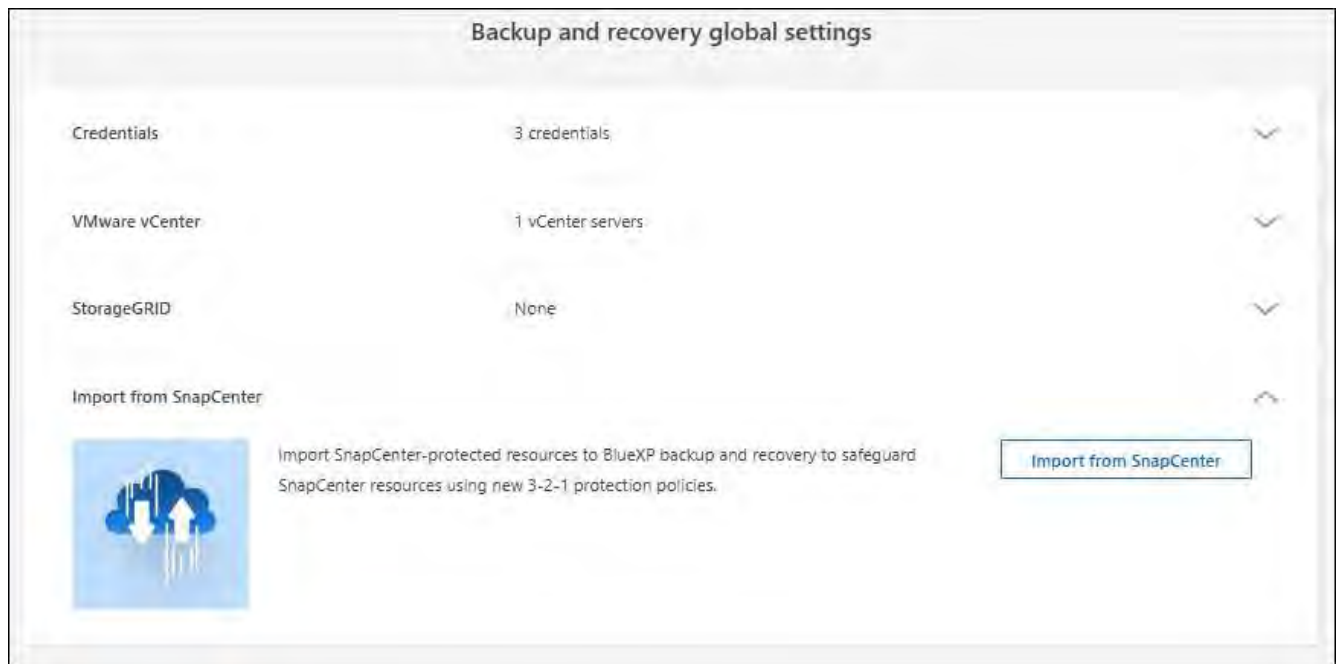
After you import SnapCenter host resources, BlueXP backup and recovery does not take over protection management. To do so, you must explicitly select to manage these resources in BlueXP backup and recovery.

### Steps

1. From the BlueXP backup and recovery menu, select **Settings**.



2. Select the down arrow to expand the **Import from SnapCenter** section.



3. Select **Import from SnapCenter** to import the SnapCenter resources.

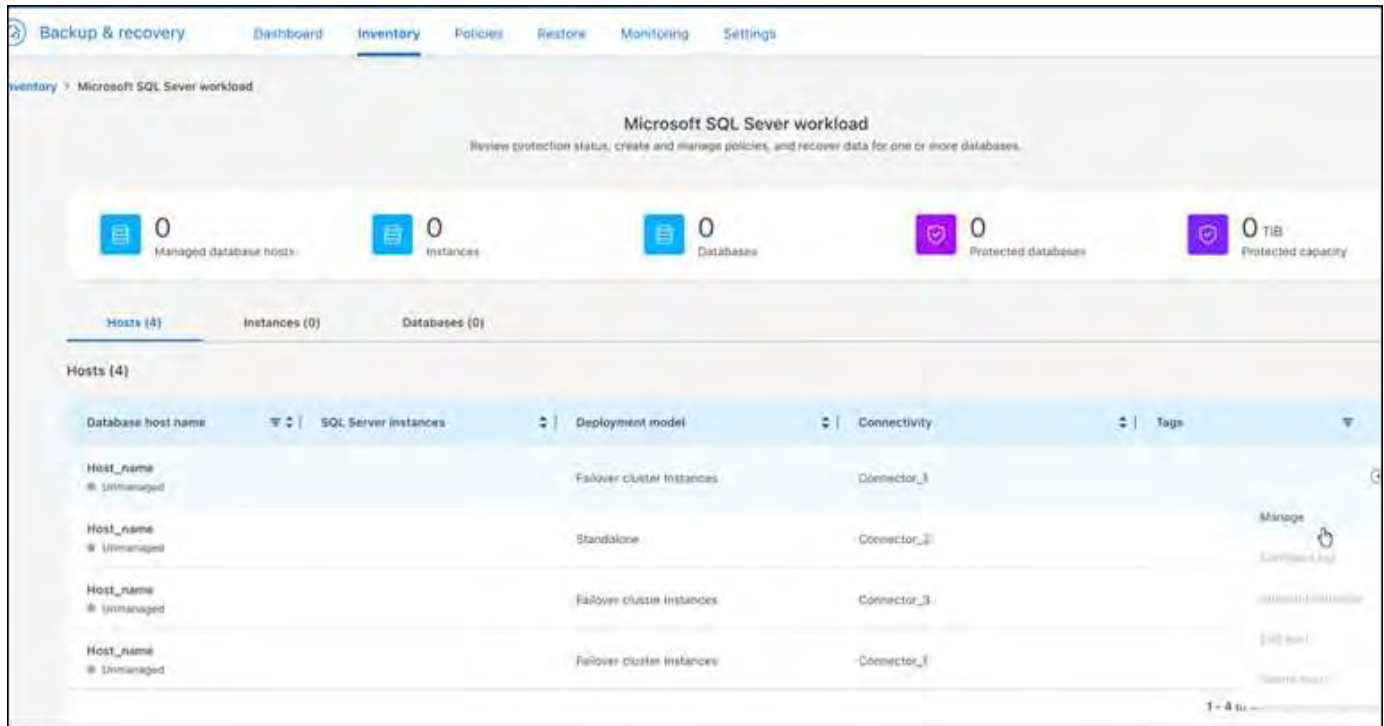
4. Enter **SnapCenter application credentials**:
  - a. **SnapCenter FQDN or IP address**: Enter the FQDN or IP address of the SnapCenter application itself.
  - b. **Port**: Enter the port number for the SnapCenter Server.
  - c. **Username** and **Password**: Enter the username and password for the SnapCenter Server.
  - d. **Connector**: Select the BlueXP Connector for SnapCenter.
5. Enter **SnapCenter server host credentials**:
  - a. **Existing credentials**: If you select this option, you can use the existing credentials that you have already added. Enter the credentials name.
  - b. **Add new credentials**: If you don't have existing SnapCenter host credentials, you can add new credentials. Enter the credentials name, authentication mode, user name, and password.
6. Select **Import** to validate your entries and register the SnapCenter Server.



If the SnapCenter Server is already registered, you can update the existing registration details.

## Result

The Inventory page shows the imported SnapCenter resources.



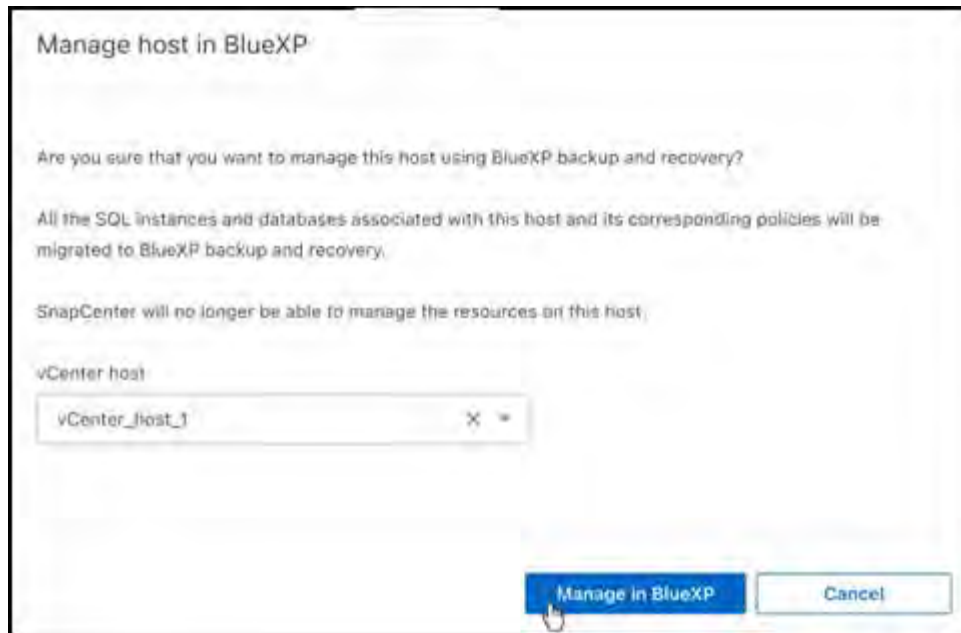
## Manage SnapCenter host resources

After you import the SnapCenter resources, manage those host resources in BlueXP backup and recovery. After you select to manage those imported resources, BlueXP backup and recovery can back up and recover the resources that you are importing from SnapCenter. You no longer need to manage those resources in SnapCenter Server.

## Steps

1. After you import the SnapCenter resources, on the Inventory page that appears, select the SnapCenter resources that you imported that you want to have BlueXP backup and recovery manage from now on.
2. Select the Actions icon **...** > **Manage** to manage the resources.





### 3. Select **Manage in BlueXP**.

The Inventory page shows **Managed** under the host name to indicate that the selected host resources are now managed by BlueXP backup and recovery.

## Edit imported SnapCenter resources

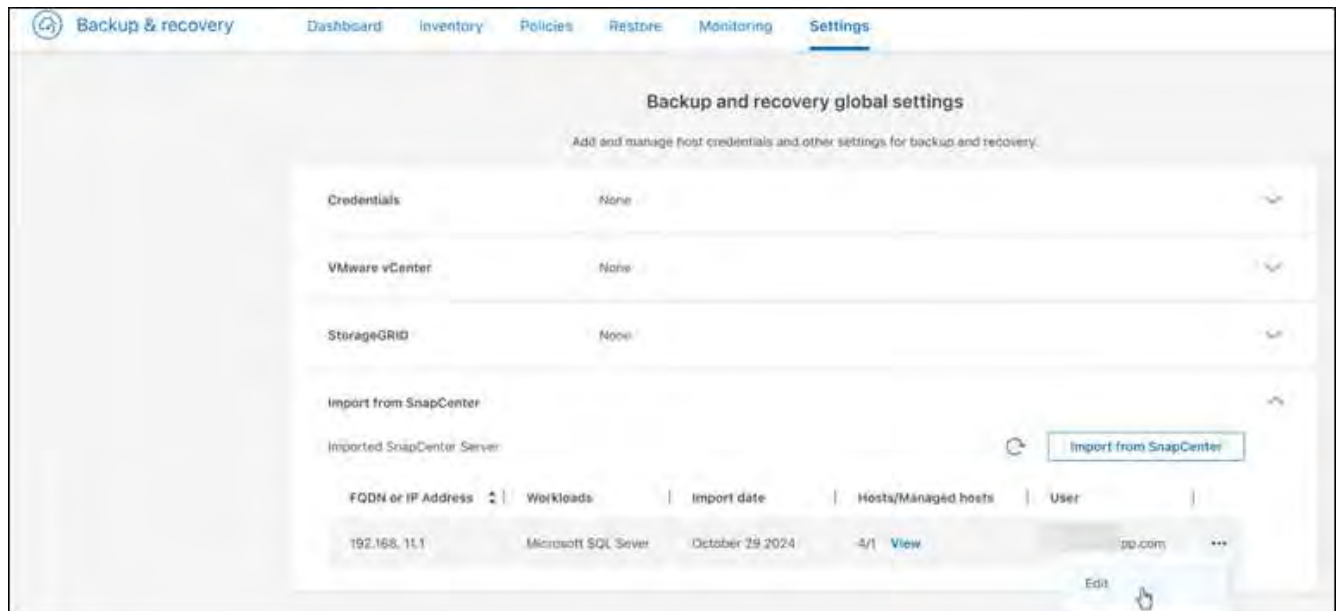
You can later re-import SnapCenter resources or edit the imported SnapCenter resources to update the registration details.

You can change only the port and password details for the SnapCenter Server.

### Steps

1. From the BlueXP backup and recovery menu, select **Settings**.
2. Select the down arrow for **Import from SnapCenter**.

The Import from SnapCenter page shows all previous imports.



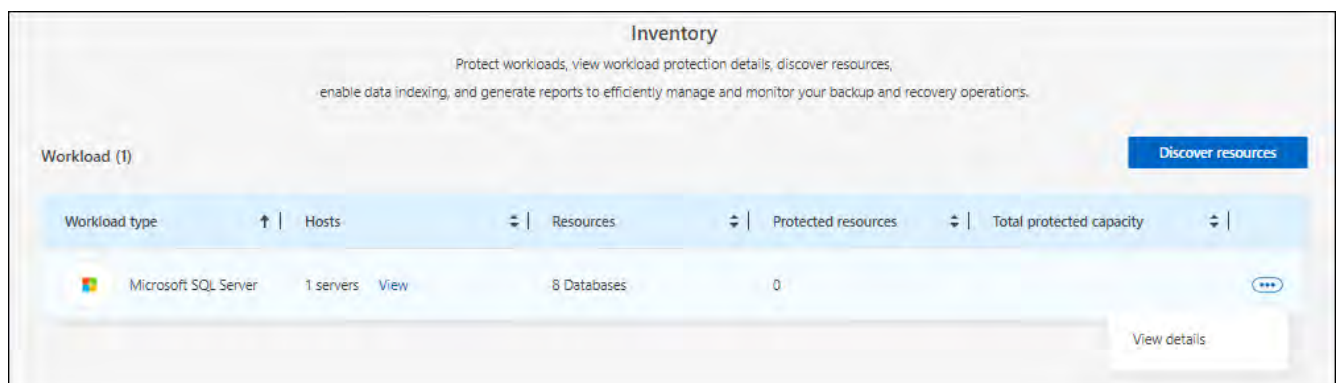
3. Select the Actions icon **...** > **Edit** to update the resources.
4. Update the SnapCenter password and port details, as needed.
5. Select **Import**.

## Configure log directories in snapshots for Windows hosts

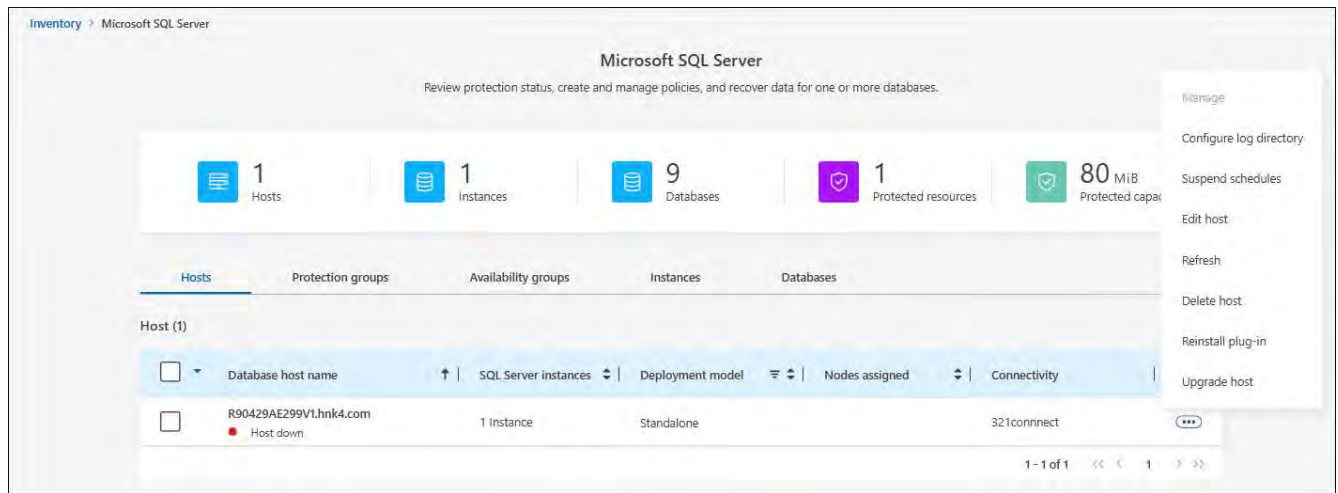
Before you create policies for Windows hosts, you should configure log directories in snapshots for Windows hosts. Log directories are used to store the logs that are generated during the backup process.

### Steps

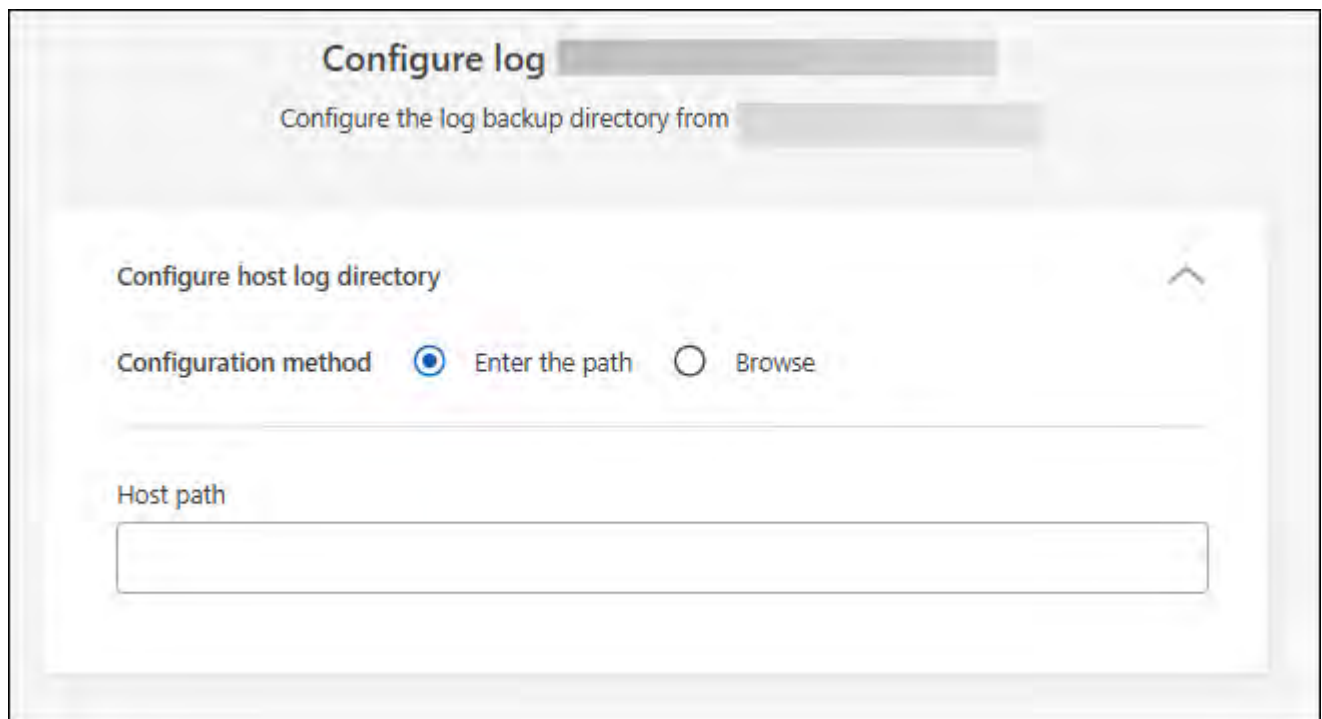
1. From the BlueXP backup and recovery menu, select **Inventory**.



2. From the Inventory page, select a workload and then select the Actions icon **...** > **View details** to display the workload details.
3. From the Inventory details page showing Microsoft SQL Server, select the Hosts tab.



- From the Inventory details page, select a host and select the Actions icon **...** > **Configure log directory**.



- Either browse or enter the path for the log directory.
- Select **Save**.

# Use BlueXP backup and recovery

## View protection health on the BlueXP backup and recovery Dashboard

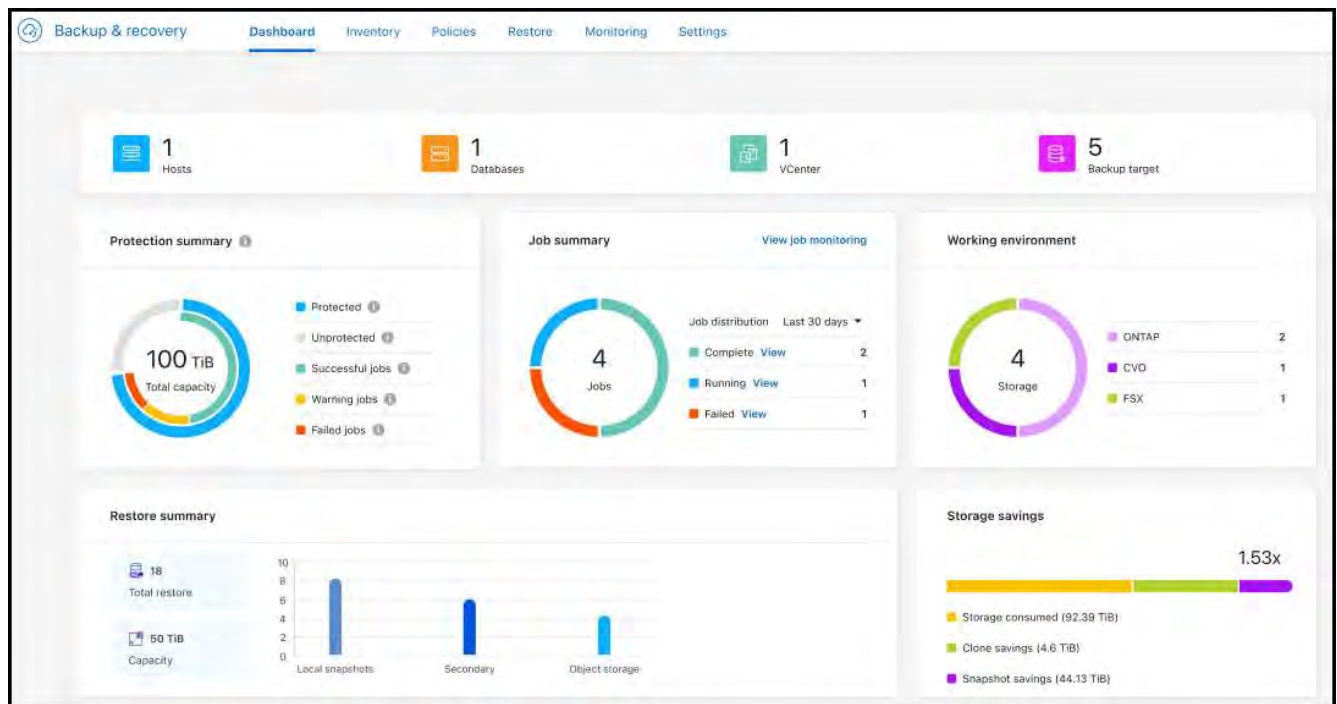
Monitoring the health of your workloads ensures that you are aware of issues with workload protection and can take steps to resolve them. View the status of your backups and restores on the BlueXP backup and recovery Dashboard. You can review the system summary, Protection summary, Job summary, Restore summary, and more.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, Backup and Recovery clone admin, or Backup and Recovery viewer role. Learn about [Backup and recovery roles and privileges](#). Learn about [BlueXP access roles for all services](#).

### Steps

1. From the BlueXP backup and recovery menu, select **Dashboard**.



## View the overall system summary

The System summary provides the following information:

- Number of hosts or VMs discovered
- Number of Kubernetes clusters discovered
- Number of backup targets on object storage
- Number of vCenters

- Number of storage clusters in ONTAP

## View the Protection summary

Review the following information in the Protection summary:

- The total number of protected and unprotected databases.



A protected database is one that has a backup policy assigned. An unprotected database is one that doesn't have a backup policy assigned to it.

- The number of backups that were successful, have a warning, or have failed.
- The total capacity discovered by the backup service and the capacity that is protected versus unprotected. Hover over the "i" icon to see the details.

## View the Job summary

Review the total jobs completed, running or failed in the Job summary.

### Steps

1. For each job distribution, change a filter to show the summary of failed, running and complete based on the number of days, for example, the last 30 days, last 7 days, last 24 hours, or last 1 year.
2. View details of the failed, running and complete jobs by selecting **View job monitoring**.

## View the Restore summary

Review the following information on the Restore summary:

- The total number of restore jobs performed
- The total amount of capacity that has been restored
- The number of restore jobs performed on local, secondary, and object storage. Hover over the chart to see the details.

# Create and manage policies to govern backups in BlueXP backup and recovery

In BlueXP backup and recovery, create your own policies that govern the backup frequency, the time the backup is taken, and the number of backup files that are retained.



Some of these options and configuration sections are not available for all workloads.

If you import resources from SnapCenter, you might encounter some differences with policies used in SnapCenter and those used in BlueXP backup and recovery. See [Policy differences between SnapCenter and BlueXP backup and recovery](#).

You can accomplish the following goals related to policies:

- Create a local snapshot policy
- Create a policy for replication to secondary storage

- Create a policy for object storage settings
- Configure advanced policy settings
- Edit policies
- Delete policies

## View policies

1. From the BlueXP backup and recovery menu, select **Policies**.

Name	Workload	Backup type	Architecture	Resources protected	Ransomware protection
azure321	Microsoft SQL Server	Full backup	3-2-1 fan-out	0 View	Snapshot locking on secondary storage, DataLock locking on object storage
azure321new	Microsoft SQL Server	Full backup	3-2-1 fan-out	1 View	Snapshot locking on secondary storage, DataLock locking on object storage
test_	Microsoft SQL Server	Full backup	Local snapshots	0 View	Snapshot locking on local snapshots
test_	Microsoft SQL Server	Full backup	Local snapshots	0 View	Snapshot locking on local snapshots
test_policy	Microsoft SQL Server	Full backup	Disk to disk	0 View	Snapshot locking on local snapshots

2. Review these policy details.

- **Workload:** Examples include Microsoft SQL Server, Volumes, VMware, or Kubernetes.
- **Backup type:** Examples include full backup and log backup.
- **Architecture:** Examples include local snapshot, fan-out, cascading, disk to disk, and disk to object store.
- **Resources protected:** Shows how many resources out of the total resources on that workload are protected.
- **Ransomware protection:** Shows if the policy includes snapshot locking on the local snapshot, snapshot locking on secondary storage, or DataLock locking on object storage.

## Create a policy

You can create policies that govern your local snapshots, replications to secondary storage, and backups to object storage. Part of your 3-2-1 strategy involves creating a snapshot copy of the Microsoft SQL Server instances or databases on the **primary** storage system.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin. Learn about [Backup and recovery roles and privileges](#). Learn about [BlueXP access roles for all services](#).

### Before you begin

If you plan on replicating to secondary storage and want to use snapshot locking on local snapshots or on

remote ONTAP secondary storage, you first need to initialize the ONTAP compliance clock on the cluster level. This is a requirement for enabling snapshot locking in the policy.

For instructions on how to do this, refer to [Initialize the compliance clock in ONTAP](#).

For information about snapshot locking in general, refer to [Snapshot locking in ONTAP](#).

## Steps

1. From the BlueXP backup and recovery menu, select **Policies**.
2. From the Policies page, select **Create new policy**.

The screenshot shows the 'New policy' configuration interface. At the top, it says 'New policy' and 'Create backup and recovery policy to protect your data'. There is an 'Expand all' link in the top right. The main content area is divided into several sections, each with a down arrow to expand it: 'Details' (Name), 'Backup architecture' (Action required), 'Local snapshot settings', 'Secondary settings' (Action required), and 'Object store settings' (Action required). Below these is an 'Advanced settings' section with a 'Select' dropdown menu. The dropdown menu is open, showing four options: 'Select All', 'Maximum transfer r...', 'Ransomware scan', and 'Backup retries', all of which are checked. At the bottom of the page, there are two buttons: 'Create' and 'Cancel'.

3. In the Policies page, provide the following information.
  - **Details** section:
    - Workload type of "Microsoft SQL Server" is selected by default for this version.
    - Enter a policy name.
  - **Backup architecture** section: Select the down arrow and choose the architecture for the backup, such as fan-out, cascading, and disk to disk.
    - **Local snapshot:** Local snapshot on the selected volume. Local snapshots are a key component of data protection strategies, capturing the state of your data at specific points in time. This creates read-only, point-in-time copies of production volumes where your workloads are running. The snapshot consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last snapshot. You can use local snapshots to recover from data loss or corruption, as well as to create backups for disaster recovery purposes.



- **3-2-1 fanout:** (Not available for Kubernetes workloads) Primary storage (disk) to secondary storage (disk) to cloud (object store). Creates multiple copies of data across different storage systems, such as ONTAP to ONTAP and ONTAP to object-store configurations. This can be a cloud hyperscaler object store or a private object store — StorageGRID. These configurations help in achieving optimal data protection and disaster recovery.



This option is not available for Amazon FSx for NetApp ONTAP.

- **3-2-1 cascaded:** (Not available for Kubernetes workloads) Primary storage (disk) to secondary storage (disk) and primary storage (disk) to cloud storage (object store). This can be a cloud hyperscaler object store or a private object store — StorageGRID. This creates a chain of data replication across multiple systems to ensure redundancy and reliability.



This option is not available for Amazon FSx for NetApp ONTAP.

- **Disk to disk:** (Not available for Kubernetes workloads) Primary storage (disk) to secondary storage (disk). The ONTAP to ONTAP data protection strategy replicates data between two ONTAP systems to ensure high availability and disaster recovery. This is typically achieved using SnapMirror, which supports both synchronous and asynchronous replication. This method ensures that your data is continuously updated and available across multiple locations, providing robust protection against data loss.
- **Disk-to-object store:** Primary storage (disk) to cloud (object store). This replicates data from an ONTAP system to an object storage system, such as AWS S3, Azure Blob Storage or StorageGRID. This is typically achieved using SnapMirror Cloud, which provides incremental forever backups by transferring only changed data blocks after the initial baseline transfer. This can be a cloud hyperscaler object store or a private object store — StorageGRID. This method is ideal for long-term data retention and archiving, offering a cost-effective and scalable solution for data protection.
- **Disk-to-disk fanout:** (Not available for Kubernetes workloads) Primary storage (disk) to secondary storage (disk) and primary storage (disk) to secondary storage (disk).



You can configure multiple secondary settings for the disk-to-disk fanout option.

## Create a local snapshot policy

Provide information for the local snapshot.

- Select the **Add schedule** option to select the snapshot schedule or schedules. You can have a maximum of 5 schedules.
- **Snapshot frequency:** Select the frequency of hourly, daily, weekly, monthly, or yearly. The yearly frequency is not available for Kubernetes workloads.
- **Snapshot retention:** Enter the number of snapshots to keep.
- **Enable log backup:** (Not available for Kubernetes workloads) Check the option to back up logs and set the frequency and retention of the log backups. To do this, you must have already configured a log backup. See [Configure log directories](#).
- **Provider:** (Kubernetes workloads only) Select the storage provider that hosts the Kubernetes application resources.
- **Backup target:** (Kubernetes workloads only) Select the storage bucket that hosts the Kubernetes application resources. The application resource definitions at the time of the snapshot are stored in this



bucket. Ensure that the bucket is accessible within your backup environment.

- Optionally, select **Advanced** at the right of the schedule to set the SnapMirror label and enable snapshot locking (not available for Kubernetes workloads).
  - **SnapMirror label:** The label serves as a marker for transferring a specified snapshot according to the retention rules of the relationship. Adding a label to a snapshot marks it as a target for SnapMirror replication.
  - **Offset from an hour:** Enter the number of minutes to offset the snapshot from the start of the hour. For example, if you enter **15**, the snapshot will be taken at 15 minutes past the hour.
  - **Enable silent hours:** Select whether you want to enable silent hours. Silent hours are a period during which no snapshots are taken, allowing for maintenance or other operations without interference from backup processes. This is useful for reducing the load on the system during peak usage times or maintenance windows.
  - **Enable snapshot locking:** Select whether you want to enable tamper-proof snapshots. Enabling this option ensures that the snapshots cannot be deleted or altered until the specified retention period has expired. This feature is crucial for protecting your data against ransomware attacks and ensuring data integrity.
  - **Snapshot locking period:** Enter the number of days, months, or years that you want to lock the snapshot.

### Create a policy for secondary settings (replication to secondary storage)

Provide information for the replication to secondary storage. Schedule information from the local snapshot settings appears for you in the secondary settings. These settings are not available for Kubernetes workloads.

- **Backup:** Select the frequency of hourly, daily, weekly, monthly, or yearly.
- **Backup target:** Select the target system on secondary storage for the backup.
- **Retention:** Enter the number of snapshots to keep.
- **Enable snapshot locking:** Select whether you want to enable tamper-proof snapshots.
- **Snapshot locking period:** Enter the number of days, months, or years that you want to lock the snapshot.
- **Transfer to secondary:**
  - The **ONTAP transfer schedule - Inline** option is selected by default and that indicates that snapshots are transferred to the secondary storage system immediately. You don't need to schedule the backup.
  - Other options: If you choose a deferred transfer, the transfers are not immediate and you can set a schedule.
- **SnapMirror and SnapVault SMAS secondary relationship:** Use SnapMirror and SnapVault SMAS secondary relationships for SQL Server workloads.

### Create a policy for object storage settings

Provide information for the backup to object storage. These settings are called "Backup settings" for Kubernetes workloads.



The fields that appear differ depending on the provider and architecture selected.

### Create a policy for AWS object storage

Enter information in these fields:

- **Provider:** Select **AWS**.
- **AWS account:** Select the AWS account.
- **Backup target:** Select a registered S3 object storage target. Ensure that the target is accessible within your backup environment.
- **IPspace:** Select the IPspace to use for the backup operations. This is useful if you have multiple IPspaces and want to control which one is used for backups.
- **Schedule settings:** Select the schedule that was set for the local snapshots. You can remove a schedule, but you cannot add one because the schedules are set according to the local snapshot schedules.
- **Retention copies:** Enter the number of snapshots to keep.
- **Run at:** Choose the ONTAP transfer schedule to back up data to object storage.
- **Tier your backups from object store to archival storage:** If you choose to tier backups to archive storage (for example, AWS Glacier), select the tier option and the number of days to archive.

#### Create a policy for Microsoft Azure object storage

Enter information in these fields:

- **Provider:** Select **Azure**.
- **Azure subscription:** Select the Azure subscription from those discovered.
- **Azure resource group:** Select the Azure resource group from those discovered.
- **Backup target:** Select a registered object storage target. Ensure that the target is accessible within your backup environment.
- **IPspace:** Select the IPspace to use for the backup operations. This is useful if you have multiple IPspaces and want to control which one is used for backups.
- **Schedule settings:** Select the schedule that was set for the local snapshots. You can remove a schedule, but you cannot add one because the schedules are set according to the local snapshot schedules.
- **Retention copies:** Enter the number of snapshots to keep.
- **Run at:** Choose the ONTAP transfer schedule to back up data to object storage.
- **Tier your backups from object store to archival storage:** If you choose to tier backups to archive storage, select the tier option and the number of days to archive.

#### Create a policy for StorageGRID object storage

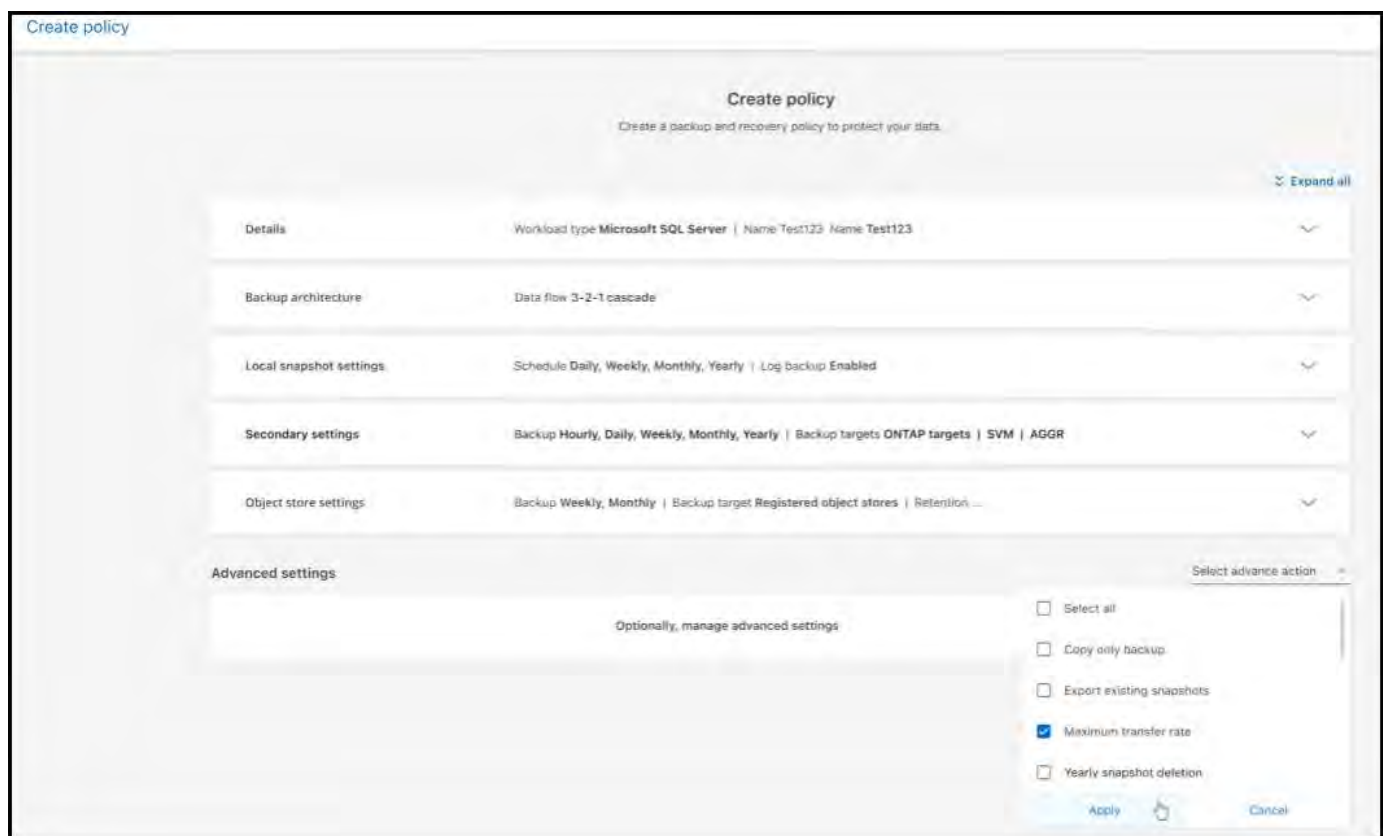
Enter information in these fields:

- **Provider:** Select **StorageGRID**.
- **StorageGRID credentials:** Select the StorageGRID credentials from those discovered. These credentials are used to access the StorageGRID object storage system and were entered in the Settings option.
- **Backup target:** Select a registered S3 object storage target. Ensure that the target is accessible within your backup environment.
- **IPspace:** Select the IPspace to use for the backup operations. This is useful if you have multiple IPspaces and want to control which one is used for backups.
- **Schedule settings:** Select the schedule that was set for the local snapshots. You can remove a schedule, but you cannot add one because the schedules are set according to the local snapshot schedules.
- **Retention copies:** Enter the number of snapshots to keep for each frequency.

- **Transfer schedule for object storage:** (Not available for Kubernetes workloads) Choose the ONTAP transfer schedule to back up data to object storage.
- **Enable integrity scan:** (Not available for Kubernetes workloads) Select whether you want to enable integrity scans (snapshot locking) on the object storage. This ensures that the backups are valid and can be restored successfully. The integrity scan frequency is set to 7 days by default. To protect your backups from being modified or deleted, select the **Integrity scan** option. The scan occurs only on the latest snapshot. You can enable or disable integrity scans on the latest snapshot.
- **Tier your backups from object store to archival storage:** (Not available for Kubernetes workloads) If you choose to tier backups to archive storage, select the tier option and the number of days to archive.

## Configure advanced settings in the policy

Optionally, you can configure advanced settings in the policy. These settings are available for all backup architectures, including local snapshots, replication to secondary storage, and backups to object storage. These settings are not available for Kubernetes workloads.



## Steps

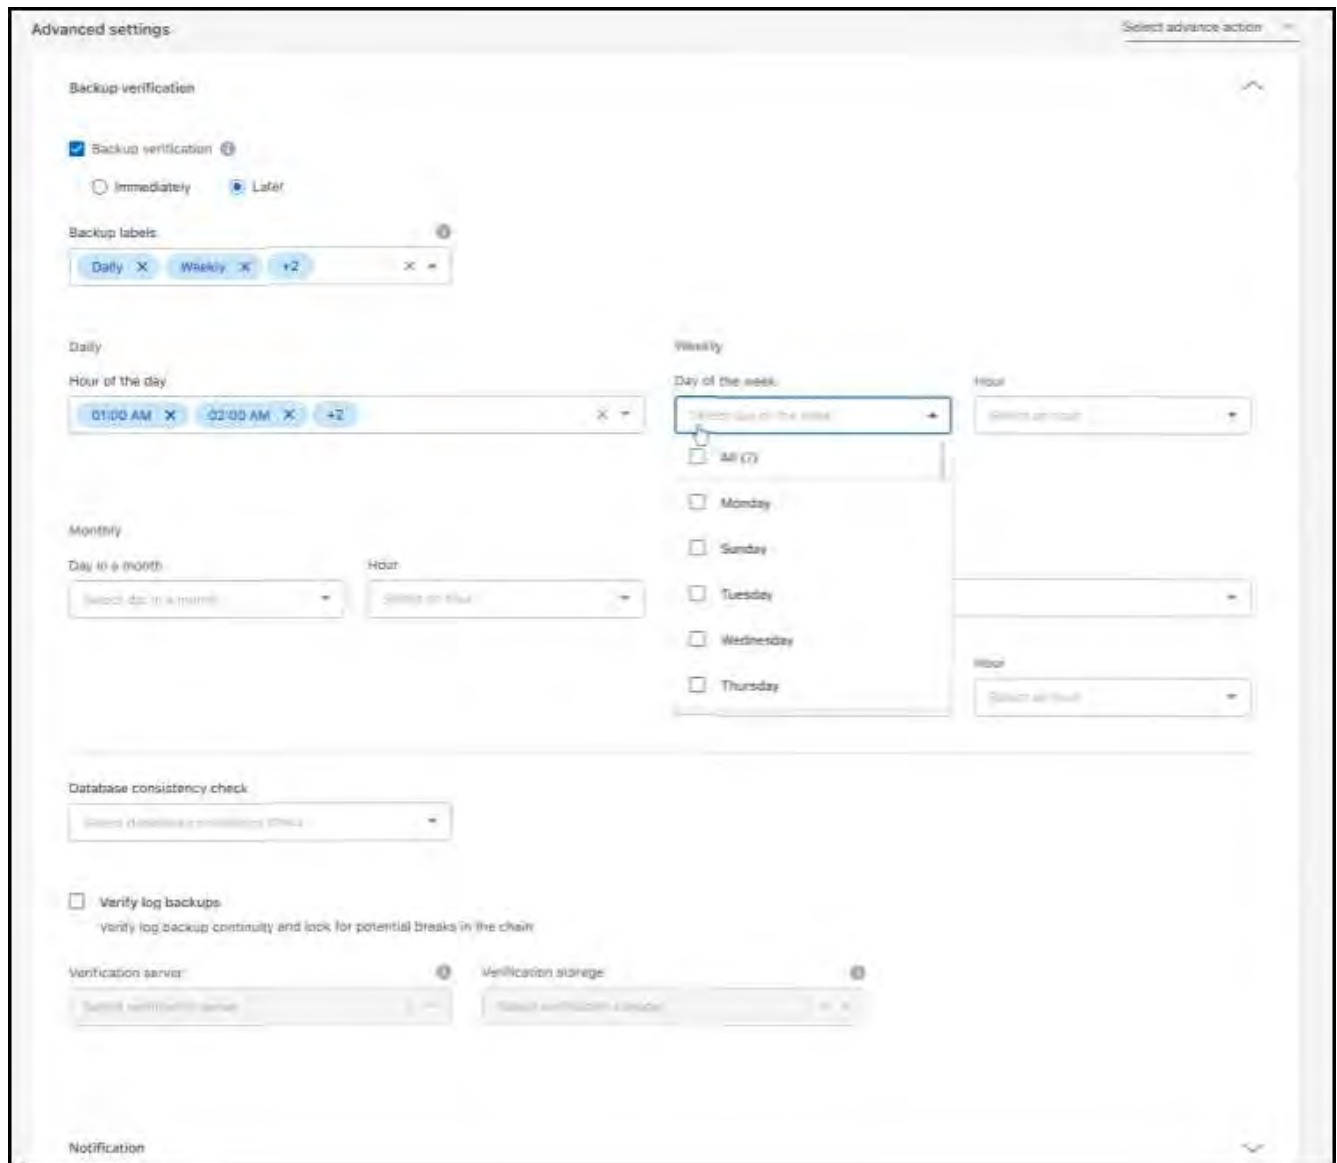
1. From the BlueXP backup and recovery menu, select **Policies**.
2. From the Policies page, select **Create new policy**.
3. In the **Policy > Advanced** settings section, select the down arrow and select the option.
4. Provide the following information:
  - **Copy only backup:** Choose copy-only backup (a type of Microsoft SQL Server backup) that lets you back up your resources by using another backup application.
  - **Availability group settings:** Select preferred backup replicas or specify a particular replica. This setting is useful if you have a SQL Server availability group and want to control which replica is used for backups.

- **Maximum transfer rate:** To not set a limit on bandwidth usage, select **Unlimited**. If you want to limit the transfer rate, select **Limited** and select the network bandwidth between 1 and 1,000 Mbps allocated to upload backups to object storage. 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 backup traffic is affecting normal user workloads, consider decreasing the amount of network bandwidth that is used during the transfer.
- **Backup retries:** To retry the job in case of a failure or interruption, select **Enable job retries during failure**. Enter the maximum number of snapshot and backup job retries and the retry time interval. The recount must be less than 10. This setting is useful if you want to ensure that the backup job is retried in case of a failure or interruption.



If the snapshot frequency is set to 1 hour, the maximum delay along with the retry count shouldn't exceed 45 minutes.

- **Ransomware scan:** Select whether you want to enable ransomware scanning on each bucket. This requires DataLock locking on object storage. Enter the frequency of the scan in days. This option applies to AWS and Microsoft Azure object storage. Note that this option might incur additional charges, depending on the cloud provider.
- **Backup verification:** Select whether you want to enable backup verification and whether you want it immediately or later. This feature ensures that the backups are valid and can be restored successfully. We recommend that you enable this option to ensure the integrity of your backups. By default, backup verification runs from secondary storage if secondary storage is configured. If secondary storage isn't configured, backup verification runs from primary storage.



Additionally, configure the following options:

- **Daily, Weekly, Monthly, or Yearly** verification: If you chose **Later** as the backup verification, select the frequency of backup verification. This ensures that backups are regularly checked for integrity and can be restored successfully.
- **Backup labels**: Enter a label for the backup. This label is used to identify the backup in the system and can be useful for tracking and managing backups.
- **Database consistency check**: Select whether you want to enable database consistency checks. This option ensures that the databases are in a consistent state before the backup is taken, which is crucial for ensuring data integrity.
- **Verify log backups**: Select whether you want to verify log backups. Select the verification server. If you chose disk-to-disk or 3-2-1, also select the verification storage location. This option ensures that the log backups are valid and can be restored successfully, which is important for maintaining the integrity of your databases.
- **Networking**: Select the network interface to use for the backup operations. This is useful if you have multiple network interfaces and want to control which one is used for backups.
  - **IPspace**: Select the IPspace to use for the backup operations. This is useful if you have multiple

IPspaces and want to control which one is used for backups.

- **Private endpoint configuration:** If you are using a private endpoint for your object storage, select the private endpoint configuration to use for the backup operations. This is useful if you want to ensure that the backups are transferred securely over a private network connection.
- **Notification:** Select whether you want to enable email notifications for backup operations. This is useful if you want to be notified when a backup operation starts, completes, or fails.
- **SnapMirror and snapshot format:** Optionally, enter your own snapshot name in a policy that governs the backups for Microsoft SQL Server workloads. Enter the format and custom text. If you chose to backup to secondary storage, you can also add a SnapMirror volume prefix and suffix.

**Create policy**  
Create a backup and recovery policy to protect your data.

[Expand all](#)

**Details** Workload type: Microsoft SQL Server | Name: Test123 | Name: Test123

**Backup architecture** Data flow: 3-2-1 cascade

**Local snapshot settings** Schedule: Daily, Weekly, Monthly, Yearly | Log backup: Enabled

**Secondary settings** Backup: Hourly, Daily, Weekly, Monthly, Yearly | Backup targets: ONTAP targets | SVM | AGGR

**Object store settings** Backup: Weekly, Monthly | Backup target: Registered object stores | Retention: ...

**Advanced settings** [Select advanced action](#)

**SnapMirror volume and snapshot format**

Use custom name format for snapshot copy

Snapshot name format: Protection group X \$Policy X -5 | Custom text: Test\_text

Provide SnapMirror volume format (ONTAP Secondary)

Prefix: Vol\_ | Suffix: \_Dest

## Edit a policy

You can edit backup architecture, backup frequency, retention policy, and other settings for a policy.

You can add another protection level when you edit a policy, but you cannot remove a protection level. For example, if the policy is only protecting local snapshots, you can add replication to secondary storage or backups to object storage. If you have local snapshots and replication, you can add object storage. However, if you have local snapshots, replication, and object storage, you cannot remove one of these levels.

If you are editing a policy that backs up to object storage, you can enable archival.


If you imported resources from SnapCenter, you might encounter some differences policies used in SnapCenter and those used in BlueXP backup and recovery. See [Policy differences between SnapCenter and](#)

[BlueXP backup and recovery](#).

### Required BlueXP role

Organization admin or Folder or project admin. [Learn about BlueXP access roles for all services](#).

### Steps

1. In BlueXP, got to **Protection** > **Backup and recovery**.
2. Select the **Policies** tab.
3. Select the policy that you want to edit.
4. Select the **Actions**  icon, and select **Edit**.


## Delete a policy

You can delete a policy if you no longer need it.



You cannot delete a policy that is associated with a workload.

### Steps

1. In BlueXP, got to **Protection** > **Backup and recovery**.
2. Select the **Policies** tab.
3. Select the policy that you want to delete.
4. Select the **Actions**  icon, and select **Delete**.
5. Review the information in the confirmation dialog box, and select **Delete**.

# Protect ONTAP volume workloads

## 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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

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.

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

### Backup-to-object features:

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



## Restore features:

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

## Supported working environments for backup and restore operations

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

### Supported regions

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

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

### Supported backup destinations

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

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

### Supported restore destinations

You can restore ONTAP data from a backup file that resides in a secondary working environment (a replicated

volume) or in object storage (a backup file) to the following working environments. Snapshot copies reside on the source working environment and can be restored only to that same system.

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

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

### Supported volumes

BlueXP backup and recovery supports the following types of volumes:

- FlexVol read-write volumes
- FlexGroup volumes (requires ONTAP 9.12.1 or later)
- SnapLock Enterprise volumes (requires ONTAP 9.11.1 or later)
- SnapLock Compliance for on-premises volumes (requires ONTAP 9.14 or later)
- SnapMirror data protection (DP) destination volumes



BlueXP backup and recovery does not support backups of FlexCache volumes.

See the sections on [Backup and restore limitations for ONTAP volumes](#) 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.

## 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 organization or account.

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

### **Pay-as-you-go subscription**

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

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

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

### **Annual contract**

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

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

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

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

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

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

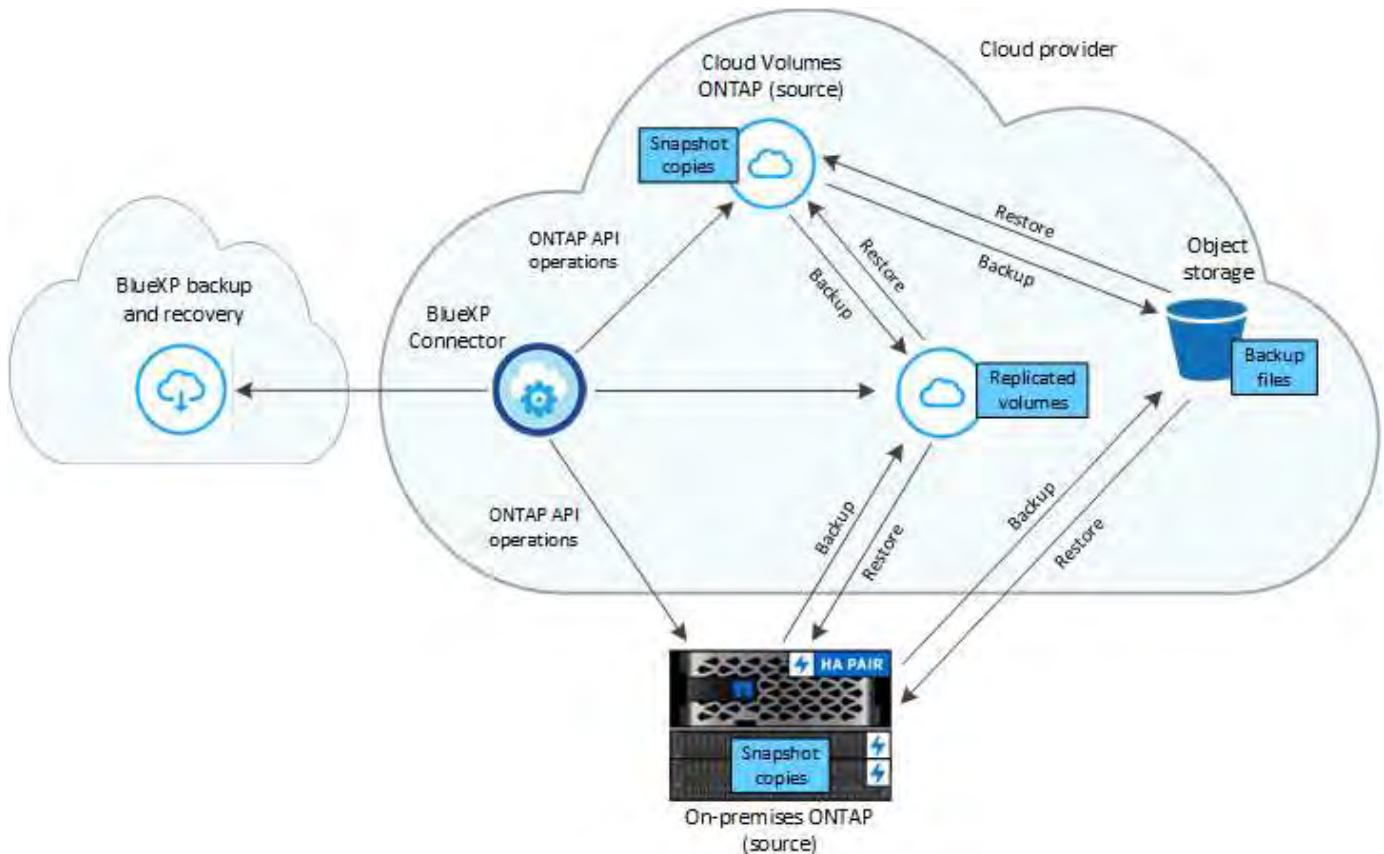
### **How BlueXP backup and recovery works**

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



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

The following image shows the relationship between each component:



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

### Where backups reside

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

- *Snapshot copies* reside on the source volume in the source working environment.
- *Replicated volumes* reside on the secondary storage system - a Cloud Volumes ONTAP or on-premises ONTAP system.
- *Backup copies* are stored in an object store that BlueXP creates in your cloud account. There's one object store per cluster/working environment, and BlueXP names the object store as follows: "netapp-backup-clusteruid". 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 reach 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).



The retention period for backups of data protection volumes is the same as defined in the source SnapMirror relationship. You can change this if you want by using the API.

## Backup file protection settings

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

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

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

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

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

For more details, refer to this information:

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





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

### Archival storage for older backup files

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

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

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

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

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

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

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

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

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

See [xref:./prev-ontap-policy-object-options.html](#)] for 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 with BlueXP backup and recovery

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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

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)

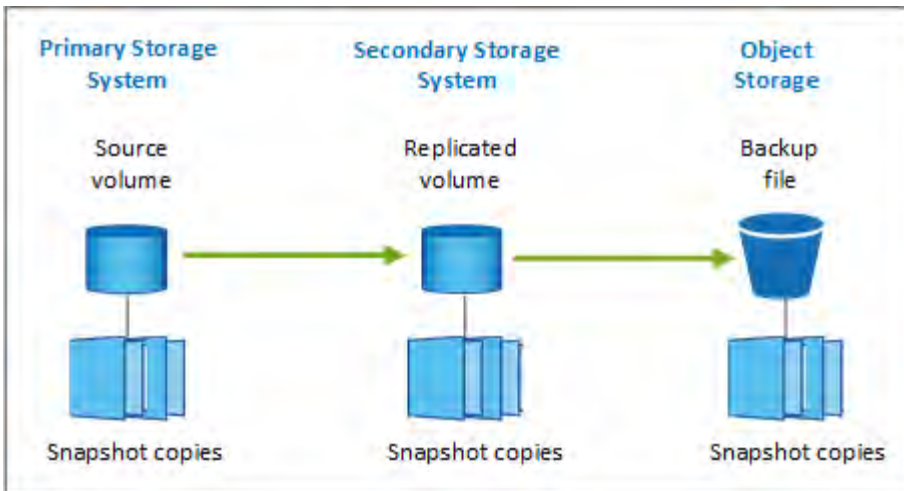
### Which protection features will you use

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.](#)

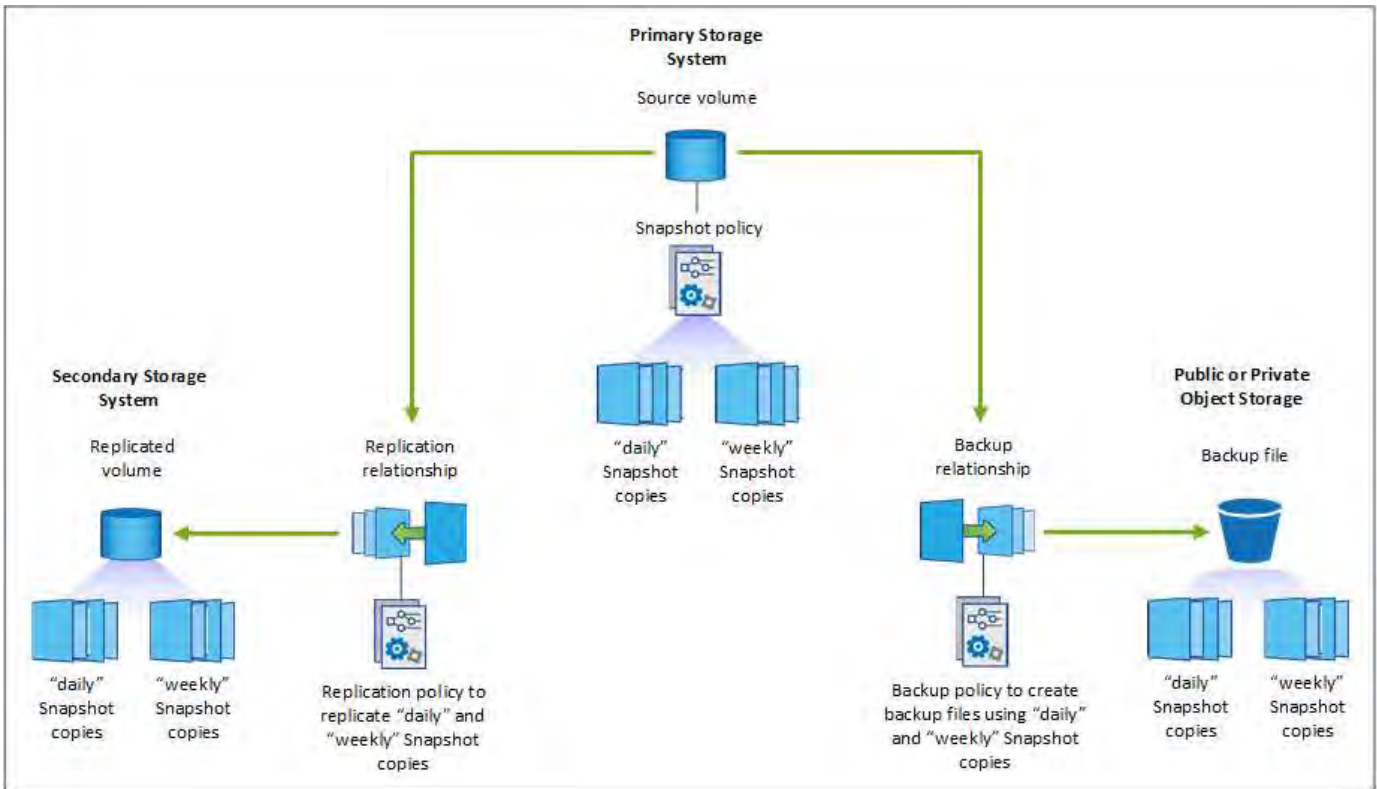
Backup type	Backup speed	Backup cost	Restore speed	Restore cost
Snapshot	High	Low (disk space)	High	Low
Replication	Medium	Medium (disk space)	Medium	Medium (network)
Cloud backup	Low	Low (object space)	Low	High (provider fees)

### Which backup architecture will you use

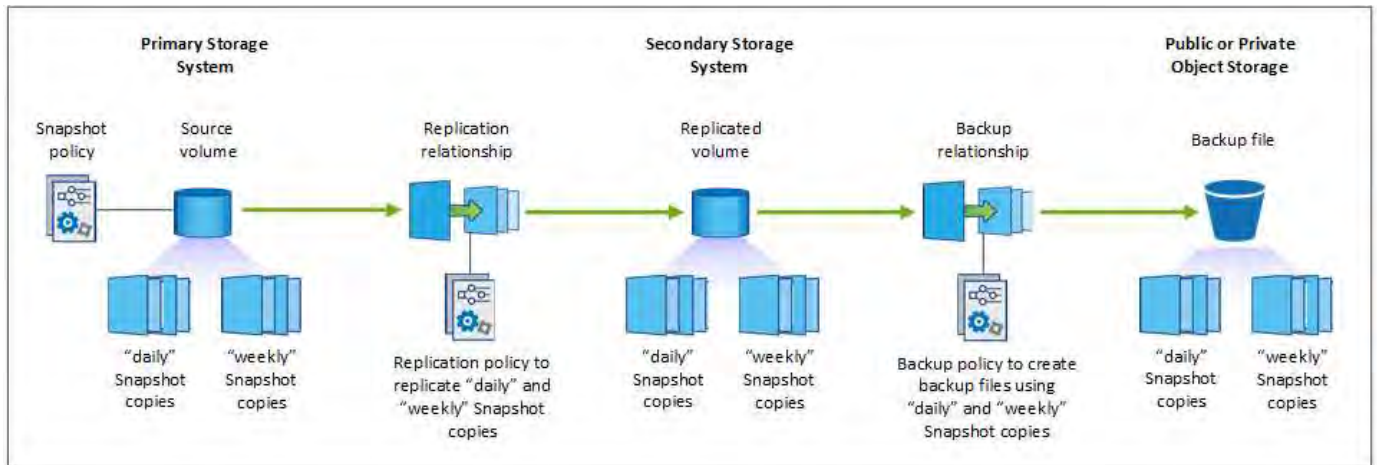
When creating both replicated volumes and backup files, you can choose a fan-out or cascade architecture to back up your volumes.

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

Fan-out	Cascade
Easier to set up because all policies, networking, and ONTAP configurations are done on the source system	Requires some networking and ONTAP configuration to be done from the secondary system as well.

## Will you use the default policies for snapshots, replications, and backups

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

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

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

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

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

- [Create a snapshot policy using System Manager or the ONTAP CLI](#)
- [Create a replication policy using System Manager or the ONTAP CLI](#)

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

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

Policy Description	Command
Simple snapshot policy	<code>snapshot policy create -policy WeeklySnapshotPolicy -enabled true -schedule1 weekly -count1 10 -vservers ClusterA -snapmirror-label1 weekly</code>
Simple backup to cloud	<code>snapmirror policy create -policy &lt;policy_name&gt; -transfer -priority normal -vservers &lt;vservers_name&gt; -create -snapshot-on-source false -type vault snapmirror policy add-rule -policy &lt;policy_name&gt; -vservers &lt;vservers_name&gt; -snapmirror-label &lt;snapmirror_label&gt; -keep</code>

Policy Description	Command
Backup to cloud with DataLock and Ransomware protection	<pre> snapmirror policy create -policy CloudBackupService-Enterprise -snapshot-lock-mode enterprise -vserver &lt;vserver_name&gt; snapmirror policy add-rule -policy CloudBackupService-Enterprise -retention-period 30days </pre>
Backup to cloud with archival storage class	<pre> snapmirror policy create -vserver &lt;vserver_name&gt; -policy &lt;policy_name&gt; -archive-after-days &lt;days&gt; -create -snapshot-on-source false -type vault snapmirror policy add-rule -policy &lt;policy_name&gt; -vserver &lt;vserver_name&gt; -snapmirror-label &lt;snapmirror_label&gt; -keep </pre>
Simple replication to another storage system	<pre> snapmirror policy create -policy &lt;policy_name&gt; -type async-mirror -vserver &lt;vserver_name&gt; snapmirror policy add-rule -policy &lt;policy_name&gt; -vserver &lt;vserver_name&gt; -snapmirror-label &lt;snapmirror_label&gt; -keep </pre>



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. BlueXP backup and recovery can use any bucket and share buckets. 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](#)

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.

### 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 are supported in AWS storage when using ONTAP 9.11.1 or greater software on your clusters, and Azure storage when using ONTAP 9.12.1 or greater software.
  - For AWS, you must enable Object Locking on the bucket using a 30-day retention period.
  - For Azure, you need to create the Storage Class with version-level immutability support.

### Which BlueXP Connector deployment mode are you using

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

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

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

### Support for sites with full internet connectivity

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

supported cloud providers. [See the full list of supported backup destinations.](#)

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

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

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

#### **Support for sites with no internet connectivity**

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

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

## **Manage backup policies for ONTAP volumes with BlueXP backup and recovery**

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

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads.](#)

When you use the activation wizard to enable the backup and recovery service for your volumes, you can



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

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

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

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


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

- BlueXP UI
- System Manager UI
- ONTAP CLI



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

### View policies for a working environment

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

The Policies management page appears.

The screenshot shows the 'Policies Management' page in the BlueXP UI. At the top, there are navigation tabs: Backup and recovery, Volumes (selected), Restore, Applications, Virtual Machines, Kubernetes, Job Monitoring, and Reports. Below the navigation, the breadcrumb path is 'Volumes > Backup Settings > Policies Management'. The main content area shows a 'Working Environment: PrimaryClusterA' with four summary cards: '31 Total Policies', '4 Snapshot Policies', '20 Replication Policies', and '7 Backup Policies'. Below these cards are three tabs: 'Snapshot Policies (4)', 'Replication Policies (20)', and 'Backup Policies (7)'. The 'Snapshot Policies (4)' tab is active, displaying a table with the following data:

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

Snapshot policies are displayed by default.

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

## Create policies

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


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

### Create a snapshot policy before initiating the snapshot

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

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

### Steps

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

The Policies management page appears.

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



- a. Select the frequency of hourly, daily, weekly, monthly, or yearly.
- b. Specify the snapshot labels denoting the schedule and retention.
- c. Enter when and how often the snapshot will be taken.
- d. Retention: Enter the number of snapshots to keep.

7. Select **Create**.

### Snapshot policy example using cascading architecture

This example creates a snapshot policy with two clusters:

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

### Create a replication policy before initiating the replication

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

#### Steps

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

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

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

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

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

#### Steps

1. In the Policy management page, select **Create policy > Create backup policy**.
2. In the Policy Details section, specify the policy name.
3. Specify the SnapMirror labels (maximum of 5) denoting the retention for each label.
4. Specify the settings, including the transfer schedule and when to archive backups.

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

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

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

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

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


7. Select **Create**.

## Edit a policy

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

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

### Steps

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



The process is the same for replication and backup policies.


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

## Delete a policy

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

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

### Steps

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

## Find more information

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

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

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

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

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

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

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

## Backup-to-object policy options in BlueXP backup and recovery

BlueXP backup and recovery enables you to create backup policies with a variety of settings for your on-premises 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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### Backup schedule options

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

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

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)

Yearly backups are deleted automatically from the source system after being transferred to object storage. This default behavior can be changed in the Advanced Settings page for the Working Environment.

### DataLock and Ransomware protection options

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

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

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

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



- If you plan to use DataLock and Ransomware protection, you can enable it when creating your first backup policy and activating BlueXP backup and recovery for that cluster. You can later enable or disable ransomware scanning using BlueXP backup and recovery Advanced Settings.
- 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

With this feature, you can lock the cloud snapshots replicated via SnapMirror to Cloud and also enable the feature to detect a ransomware attack and recover a consistent copy of the snapshot on the object store. This feature is supported on AWS, Azure, and StorageGRID.

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."

Cloud providers use a Retention Until Date (RUD), which is calculated based on the Snapshot Retention Period. The Snapshot Retention Period is calculated based on the label and the retention count defined in the

backup policy.

The minimum Snapshot Retention Period is 30 days. Let's look at some examples of how this works:

- If you choose the **Daily** label with Retention Count 20, the Snapshot Retention Period is 20 days, which defaults to the minimum 30 days.
- If you choose the **Weekly** label with Retention Count 4, the Snapshot Retention Period is 28 days, which defaults to the minimum of 30 days.
- If you choose the **Monthly** label with Retention Count 3, the Snapshot Retention Period is 90 days.
- If you choose the **Yearly** label with Retention Count 1, the Snapshot Retention Period is 365 days.

#### What is Retention Until Date (RUD) and how is it calculated?

The Retention Until Date (RUD) is determined based on the Snapshot Retention Period. The Retention Until Date is calculated by summing the Snapshot Retention Period and a Buffer.

- Buffer is the Buffer for Transfer Time (3 days) + Buffer for Cost Optimization (28 days), which totals as 31 days.
- The minimum Retention Until Date is 30 days + 31 days buffer = 61 days.

Here are some examples:

- If you create a Monthly backup schedule with 12 retentions, your backups are locked for 12 months (plus 31 days) before they are deleted (replaced by the next backup file).
- If you create a backup policy that creates 30 daily, 7 weekly, 12 monthly backups, there are three locked retention periods:
  - The "30 daily" backups are retained for 61 days (30 days plus 31 days buffer),
  - The "7 weekly" backups are retained for 11 weeks (7 weeks plus 31 days), and
  - The "12 monthly" backups are retained for 12 months (plus 31 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 61 days (30 days plus 31 days buffer).



Old backups are deleted after the DataLock Retention Period expires, not after the backup policy retention period.

The DataLock retention setting overrides the policy retention setting from your backup policy. This could affect your storage costs as your backup files will be saved in the object store for a longer period of time.

#### Enable DataLock and Ransomware protection

You can enable DataLock and Ransomware protection when you create a policy. You cannot enable, modify, or disable this after the policy is created.

1. When you create a policy, expand the **DataLock and Ransomware Protection** section.
2. Choose one of the following:
  - **None**: DataLock protection and ransomware protection are disabled.
  - **Unlocked**: DataLock protection and ransomware protection are enabled. Users with specific permissions can overwrite or delete protected backup files during the retention period.

- **Locked:** DataLock protection and ransomware protection are enabled. No users can overwrite or delete protected backup files during the retention period. This satisfies full regulatory compliance.

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

### 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.)

Scans occur in these situations:

- Scans on cloud backup objects are initiated soon after they are transferred to the cloud object storage. 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.
- Ransomware scans can be initiated when the backup is selected for the restore process.
- Scans can be performed on-demand at any time.

### How does the recovery process work?

When a ransomware attack is detected, the service uses the Active Data Connector Integrity Checker REST API to start the recovery process. The oldest version of the data objects is the source of truth and is made into the current version as part of the recovery process.

Let's see how this works:

- In the event of a ransomware attack, the service tries to overwrite or delete the object in the bucket.
- Because the cloud storage is versioning-enabled, it automatically creates a new version of the backup object. If an object is deleted with versioning turned on, it is marked as deleted but is still retrievable. If an object is overwritten, previous versions are stored and marked.
- When a ransomware scan is initiated, the checksums are validated for both object versions and compared. If the checksums are inconsistent, potential ransomware has been detected.
- The recovery process involves reverting to the last known good copy.

### 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
On-premises ONTAP system	Amazon S3 Azure Blob NetApp StorageGRID

## Requirements

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

### AWS S3 permissions

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

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

- For Azure:

- Your clusters must running ONTAP 9.12.1 or greater
- The Connector can be deployed in the cloud or on your premises
- For StorageGRID:
  - Your clusters must running ONTAP 9.11.1 or greater
  - Your StorageGRID systems must be running 11.6.0.3 or greater
  - The Connector must be deployed on your premises (it can be installed in a site with or without internet access)
  - The following S3 permissions must be part of the IAM role that provides the Connector with permissions:

#### **StorageGRID S3 permissions**

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



## Restrictions

- The DataLock and Ransomware protection feature is not available if you have configured archival storage in the backup policy.
- The DataLock option you select when activating BlueXP backup and recovery must be used for all backup policies for that cluster.
- You cannot use multiple DataLock modes on a single cluster.
- If you enable DataLock, all volume backups will be locked. You can't mix locked and non-locked volume backups for a single cluster.
- DataLock and Ransomware protection is applicable for new volume backups using a backup policy with DataLock and Ransomware protection enabled. You can later enable or disable these features using the Advanced Settings option.
- FlexGroup volumes can use DataLock and Ransomware protection only when using ONTAP 9.13.1 or greater.

## Tips on how to mitigate DataLock costs

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

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

You can choose different levels of protection:

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



Microsoft Azure uses a Lock and Unlock mode instead.

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

For more tips to mitigate DataLock costs, refer to <https://community.netapp.com/t5/Tech-ONTAP-Blogs/Understanding-BlueXP-Backup-and-Recovery-DataLock-and-Ransomware-Feature-TCO/ba-p/453475>

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

## Archival storage options

When using AWS, Azure, or Google cloud storage, you can move older backup files to a less expensive archival storage class or access tier after a certain number of days. You can also choose to send your backup

files to archival storage immediately without being written to standard cloud storage. Just enter **0** as the "Archive After Days" to send your backup file directly to archival storage. This can be especially helpful for users who rarely need to access data from cloud backups or users who are replacing a backup to tape solution.

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



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

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

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

- 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.](#)

## Manage backup-to-object storage options in BlueXP backup and recovery Advanced Settings

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



These settings are available for backup-to-object storage only. None of these settings affect your Snapshot or replication settings.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads.](#)

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

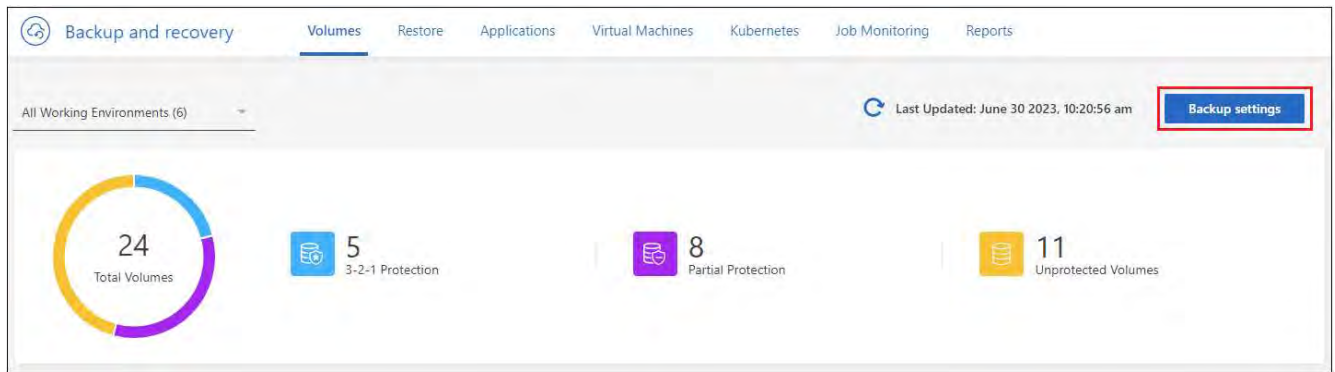
- Changing the network bandwidth allocated to upload backups to object storage using the Max Transfer Rate option
- Changing whether historical Snapshot copies are exported as backup files and included in your initial baseline backup files for future volumes
- Changing whether "yearly" snapshots are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

### View cluster-level backup settings

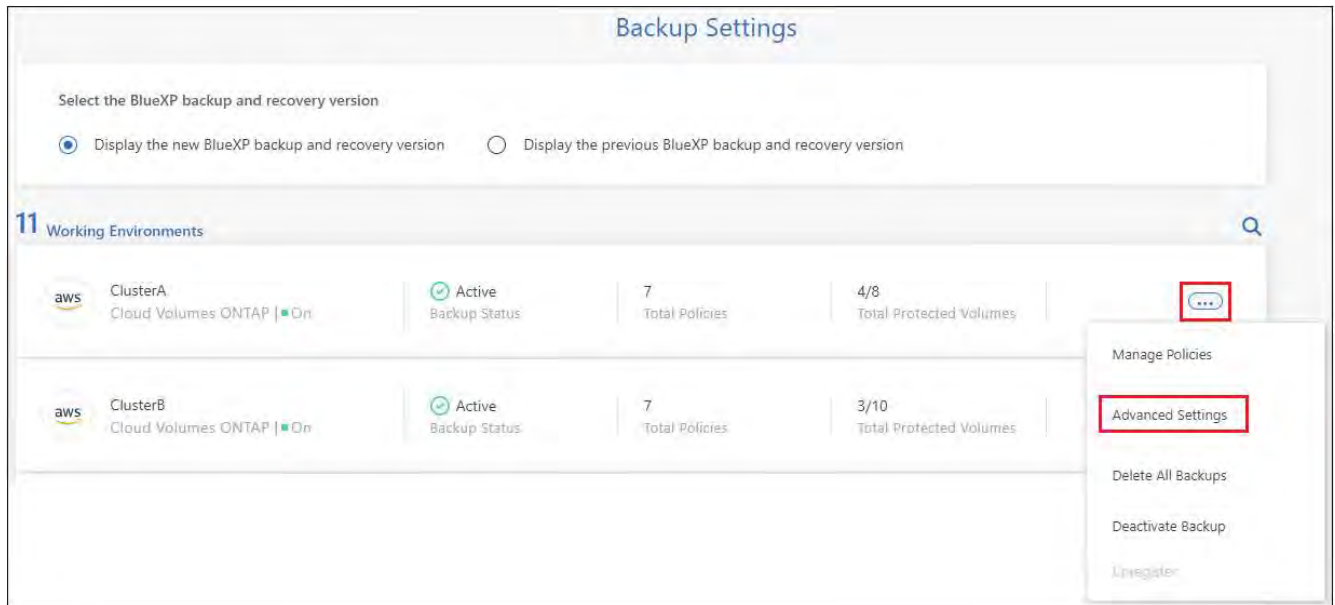
You can view the cluster-level backup settings for each working environment.

#### Steps

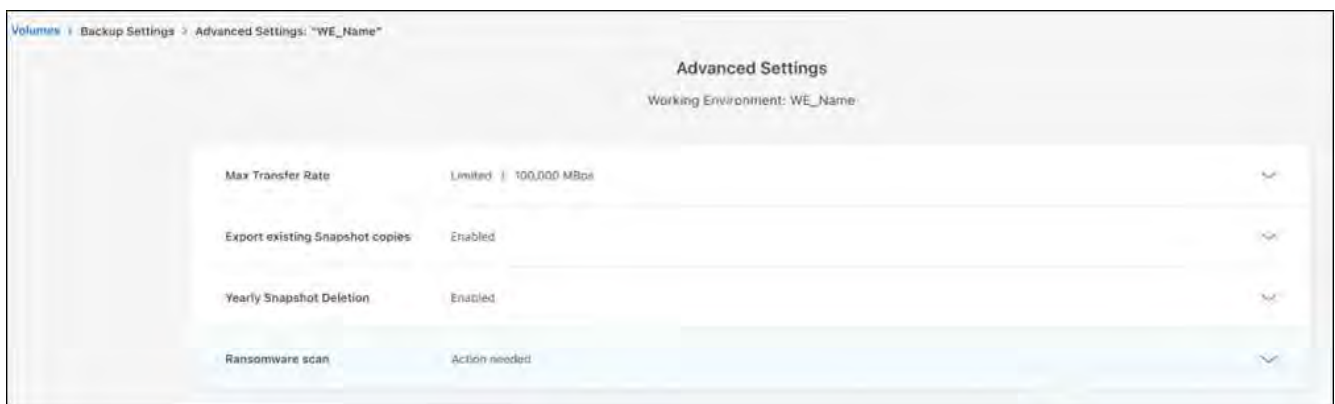
1. From the BlueXP menu, select **Protection > Backup and recovery**.
2. From the **Volumes** tab, select **Backup Settings**.



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



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



4. Expand the option and make the change.

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

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

## Change the network bandwidth available to upload backups to object storage

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

### Steps

1. From the **Volumes** tab, select **Backup Settings**.
2. From the *Backup Settings* page, click **...** for the working environment and select **Advanced Settings**.
3. In the Advanced Settings page, expand the **Max Transfer Rate** section.



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

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

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

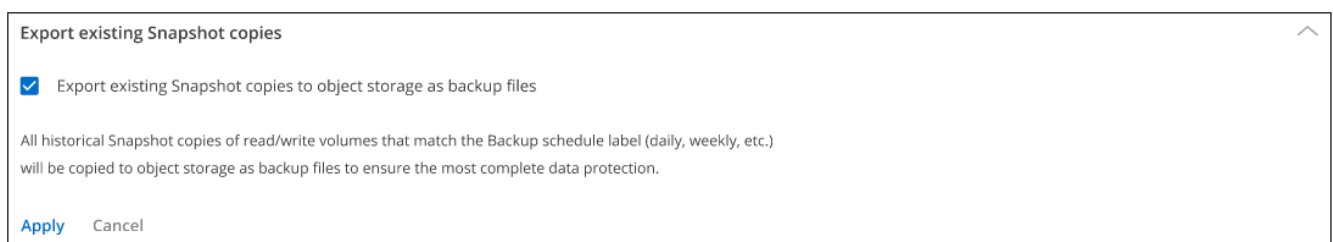
## Change whether historical snapshot copies are exported as backup files

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

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

### Steps

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



The screenshot shows a dialog box titled "Export existing Snapshot copies" with a close button in the top right corner. It contains a checked checkbox with the label "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 of the dialog box are two buttons: "Apply" and "Cancel".

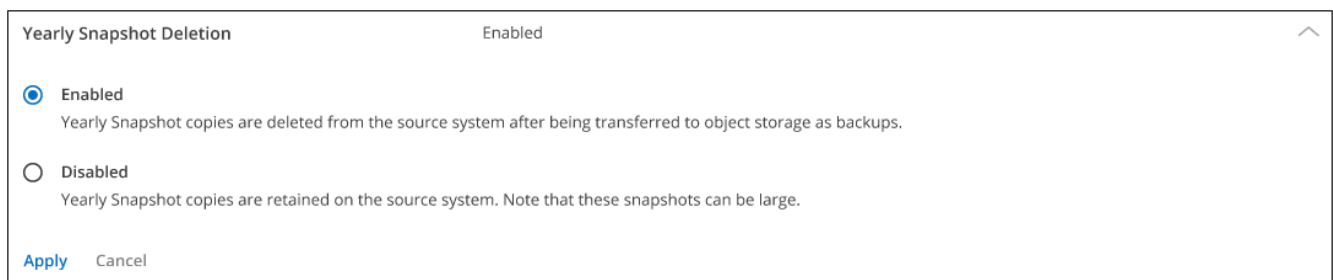
4. Select whether you want existing Snapshot copies to be exported.
5. Select **Apply**.

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

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

#### Steps

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



The screenshot shows a dialog box titled "Yearly Snapshot Deletion" with a status of "Enabled" in the top right corner. There are two radio button options: "Enabled" (selected) and "Disabled". The "Enabled" option has a sub-description: "Yearly Snapshot copies are deleted from the source system after being transferred to object storage as backups." The "Disabled" option has a sub-description: "Yearly Snapshot copies are retained on the source system. Note that these snapshots can be large." At the bottom left, there are two buttons: "Apply" and "Cancel".

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

### Enable or disable ransomware scans

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

For details about DataLock and Ransomware Protection options, refer to [DataLock and Ransomware Protection options](#).

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



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

Scheduled ransomware scans run only on the latest snapshot copy.

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

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

#### Steps

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

3. In the Advanced Settings page, expand the **Ransomware scan** section.
4. Enable or disable **Ransomware scan**.
5. Select **Scheduled ransomware scan**.
6. Optionally, change the every week default scan to days or weeks.
7. Set the how often in days or weeks that the scan should run.
8. Select **Apply**.

## Back up Cloud Volumes ONTAP data to Amazon S3 with BlueXP backup and recovery

Complete a few steps in BlueXP backup and recovery to get started backing up volume data from your Cloud Volumes ONTAP systems to Amazon S3.

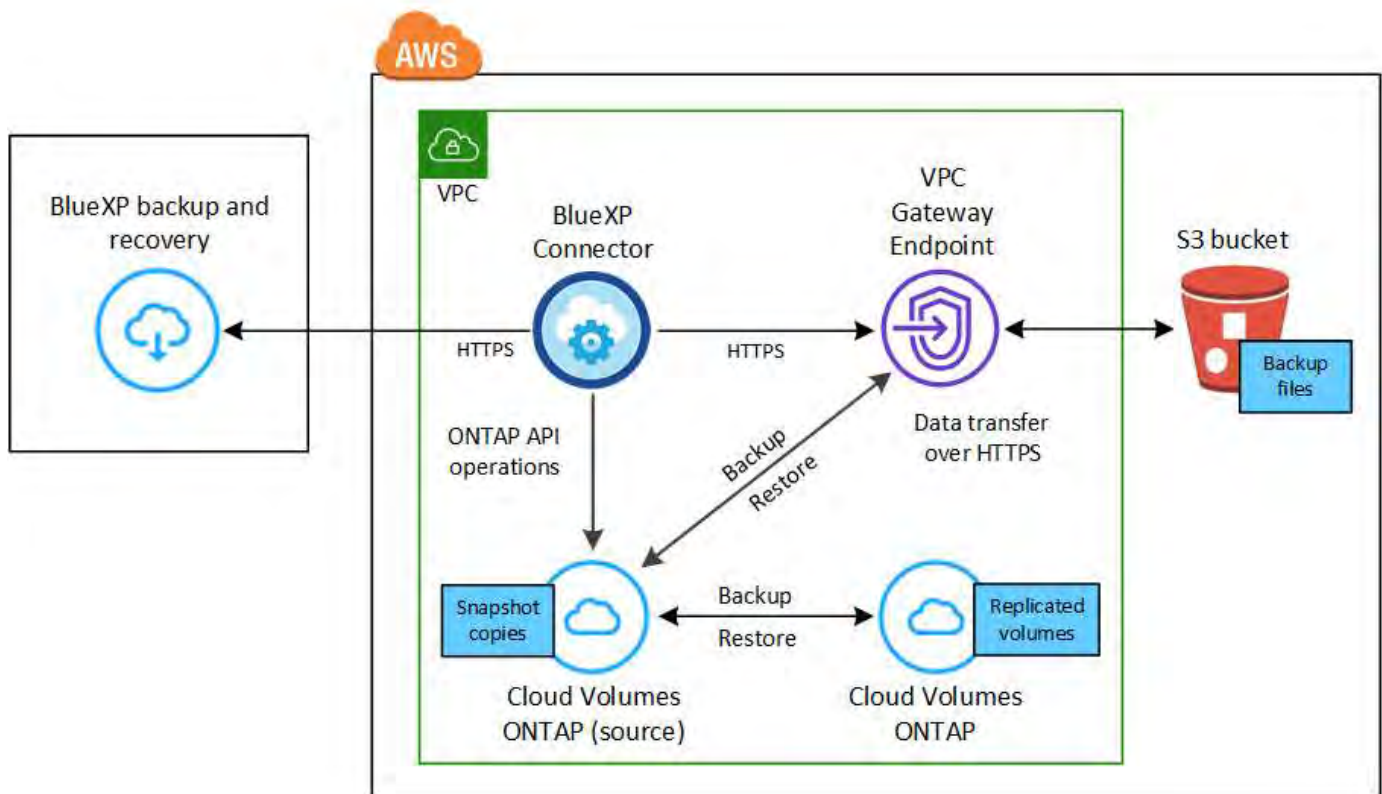
**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### 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, 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 prev-ontap-backup-cvo-aws.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.

## Activate backups on your ONTAP volumes

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

A wizard takes you through the following major steps:

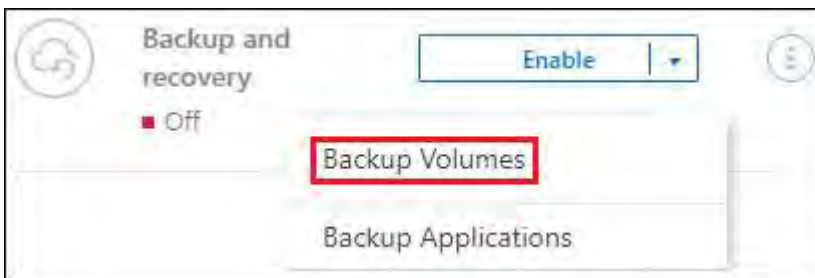
- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

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

### Start the wizard

#### Steps

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



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

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

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

2. Continue with the following options:

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

### Select the volumes that you want to back up

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

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



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

### Steps

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

6. Select **Next**.

### Review your selections

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

### Steps

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

### Result

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

A replicated volume is created in the destination cluster that will be synchronized with the primary 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 page](#).

#### 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.

## Back up Cloud Volumes ONTAP data to Azure Blob storage with BlueXP backup and recovery

Complete a few steps in BlueXP backup and recovery to get started backing up volume data from your Cloud Volumes ONTAP systems to Azure Blob storage.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

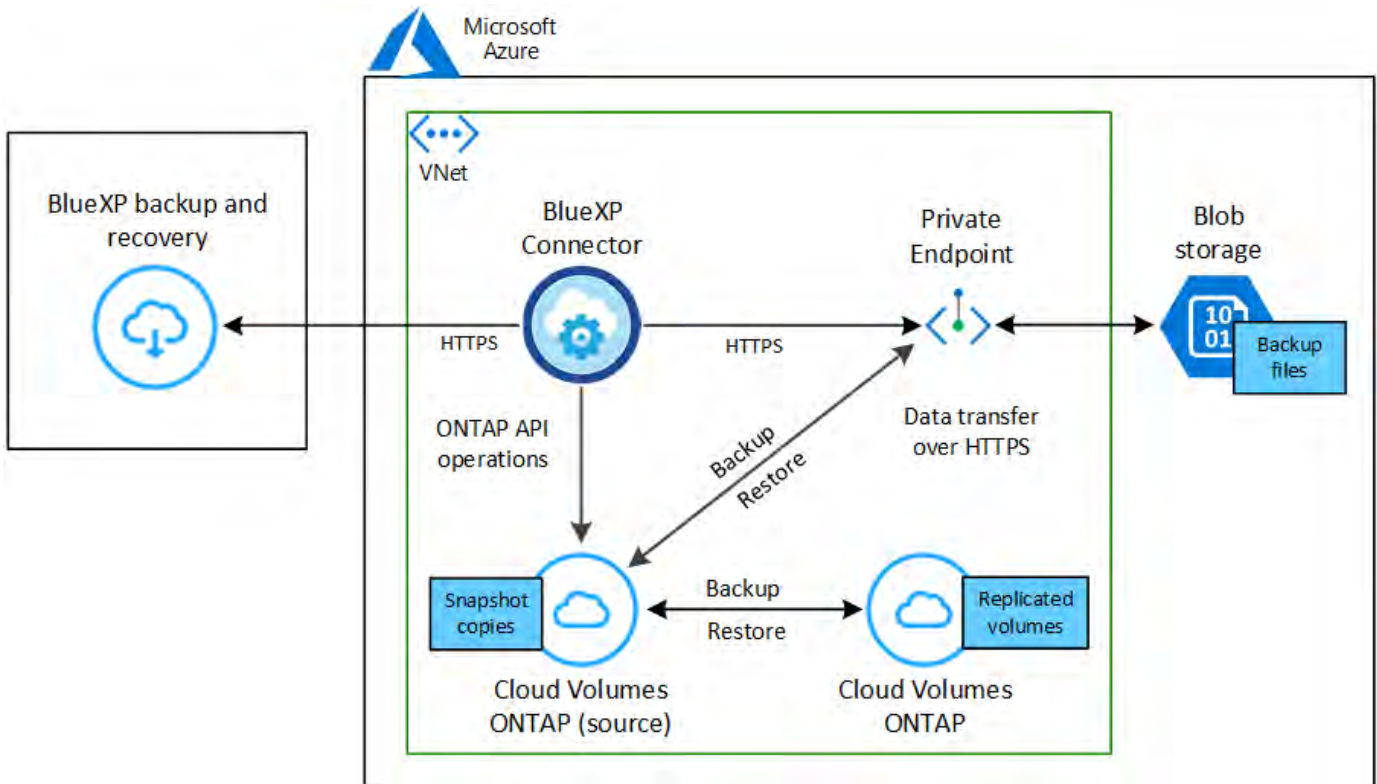
#### Verify support for your configuration

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, 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.

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

## Required information for using customer-managed keys for data encryption

You can use your own customer-managed keys for data encryption in the activation wizard instead of using the default Microsoft-managed encryption keys. In this case, you will need to have the Azure Subscription, Key Vault name, and the Key. [See how to use your own keys.](#)

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

## Create your Azure Blob storage account

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

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

## Verify ONTAP networking requirements for replicating volumes

Unresolved directive in prev-ontap-backup-cvo-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.

## 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 select **Continue**.
4. Fill out the Details & Credentials page and be sure that an Azure Marketplace subscription is in place, and select **Continue**.
5. On the Services page, leave the service enabled and select **Continue**.
6. Complete the pages in the wizard to deploy the system.

## Result

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

## Enable BlueXP backup and recovery on an existing system

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

### Steps

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

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

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

## Activate backups on your ONTAP volumes

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

A wizard takes you through the following major steps:

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

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

### Start the wizard

#### Steps

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

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

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

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

2. Continue with the following options:
  - If you already have a BlueXP Connector, you're all set. Just select **Next**.

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

### Select the volumes that you want to back up

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

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



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

### Steps

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.
  - To back up individual volumes, check the box for each volume.
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 five 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 five 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.

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



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



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

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

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



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

- **Networking:** Choose the IPspace, and whether you'll be using a Private Endpoint. Private Endpoint is disabled by default.
  - a. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.
  - b. Optionally, choose whether you'll use an Azure private endpoint that you have previously configured. [Learn about using an Azure private endpoint.](#)
- **Backup policy:** Select an existing backup-to-object storage policy.



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

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

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

#### 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 Google Cloud Storage with BlueXP backup and recovery

Complete a few steps in BlueXP backup and recovery to get started backing up volume data from your Cloud Volumes ONTAP systems to Google Cloud Storage.

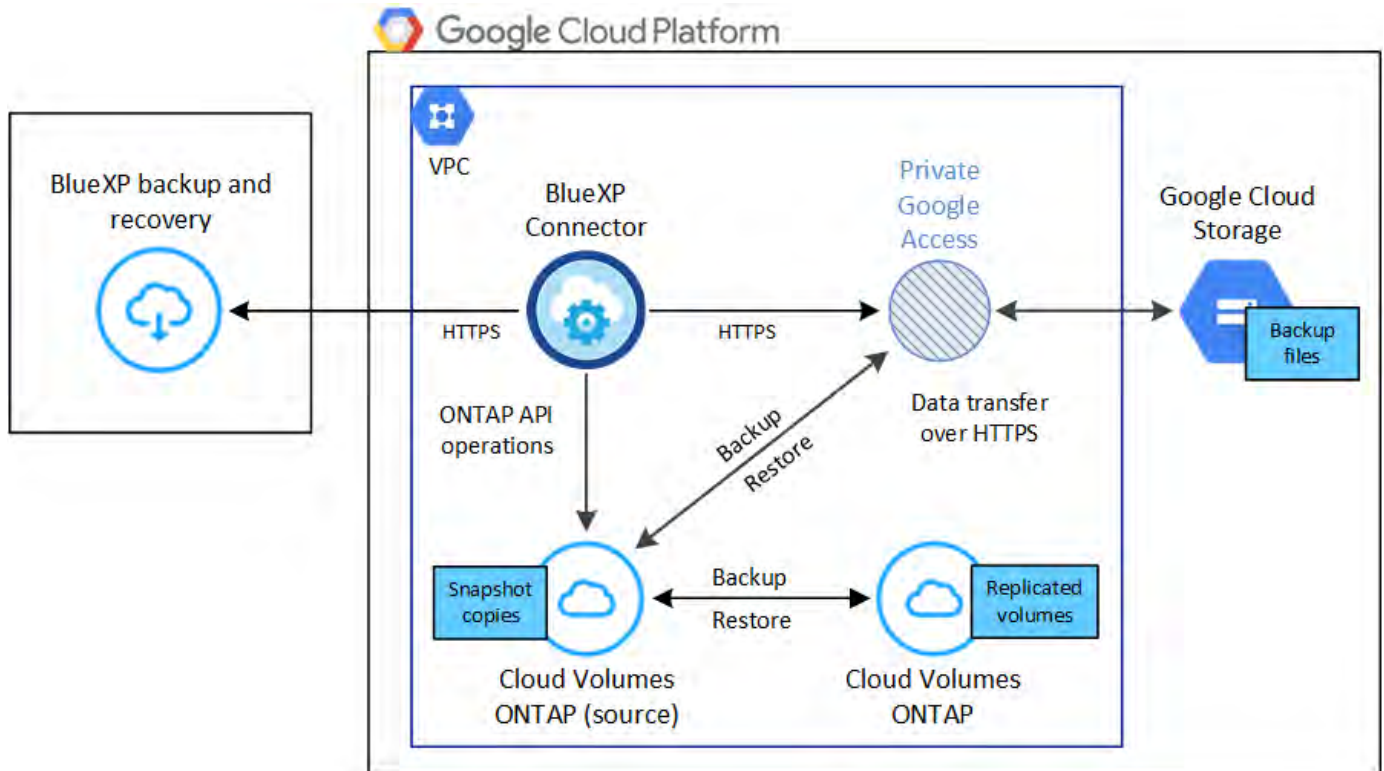
**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

## 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.

## GCP Service Account

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



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

## Verify license requirements

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

For BlueXP backup and recovery BYOL licensing, you need the serial number from NetApp that enables you to use the service for the duration and capacity of the license. [Learn how to manage your BYOL licenses.](#)

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

## Prepare your BlueXP Connector

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

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

## Verify or add permissions to the Connector

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

### Steps

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

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

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

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

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

- You'll need to have the Key Ring and the Key Name so you can add this information in the activation wizard. [Learn more about customer-managed encryption keys.](#)

- You'll need to verify that these required permissions are included in the role for the Connector:

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

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

### **CMEK considerations:**

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

### **Create your own buckets**

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

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

### **Verify ONTAP networking requirements for replicating volumes**

Unresolved directive in prev-ontap-backup-cvo-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 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](#).

### 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.

### Prepare Google Cloud Storage as your backup target

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

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

## Set up permissions

You need to provide storage access keys for a service account that has specific permissions using a custom role. A service account enables BlueXP backup and recovery to authenticate and access Cloud Storage buckets used to store backups. The keys are required so that Google Cloud Storage knows who is making the request.

### Steps

1. In the [Google Cloud Console](#), go to the **Roles** page.
2. [Create a new role](#) with the following permissions:

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

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

You'll need to enter the keys in BlueXP backup and recovery later when you configure the backup service.

## Create your own buckets

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

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

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

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

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

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

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

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

### CMEK considerations:

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

### Activate backups on your ONTAP volumes

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

A wizard takes you through the following major steps:

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

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

### Start the wizard

#### Steps



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

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

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

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

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

2. Continue with the following options:

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

### Select the volumes that you want to back up

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

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



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

### Steps

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

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

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

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 five 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 five 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 five 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 page](#).

#### 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 Amazon S3 with BlueXP backup and recovery

Complete a few steps in BlueXP backup and recovery 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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

#### Identify the connection method

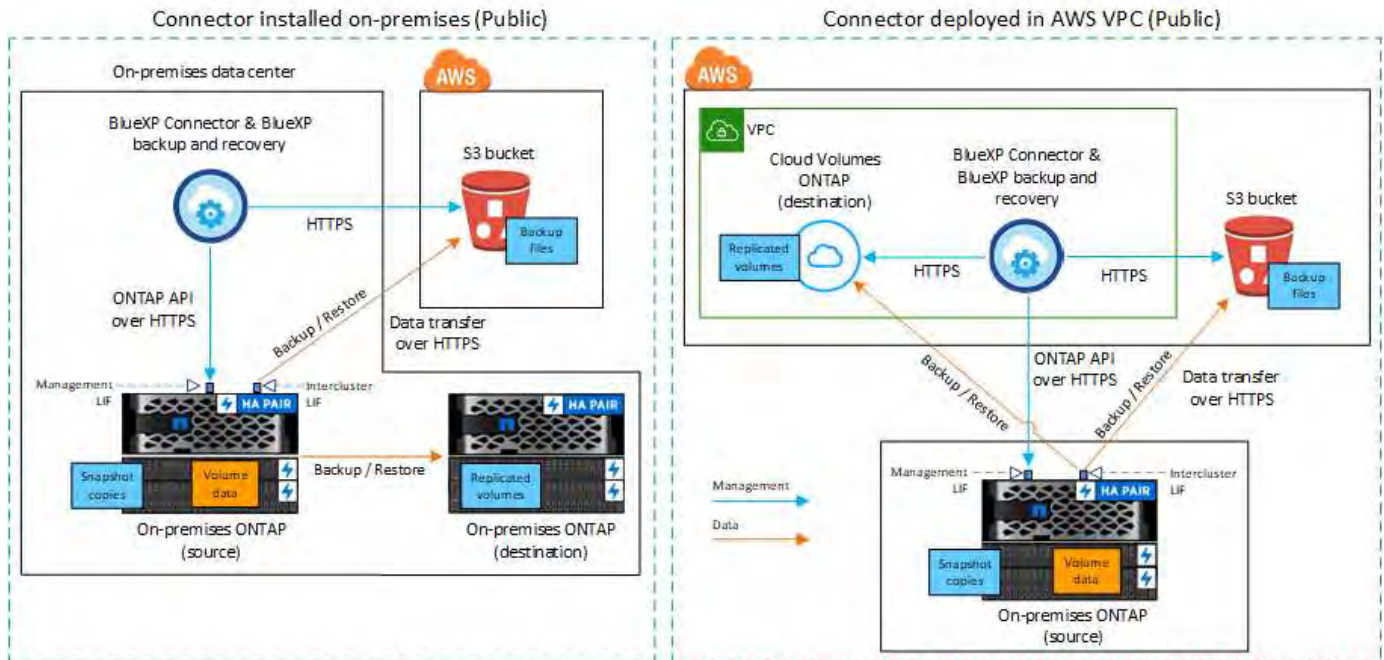
Choose which of the two connection methods you will use when configuring backups from on-premises ONTAP systems to AWS S3.

- **Public connection** - Directly connect the ONTAP system to AWS S3 using a public S3 endpoint.

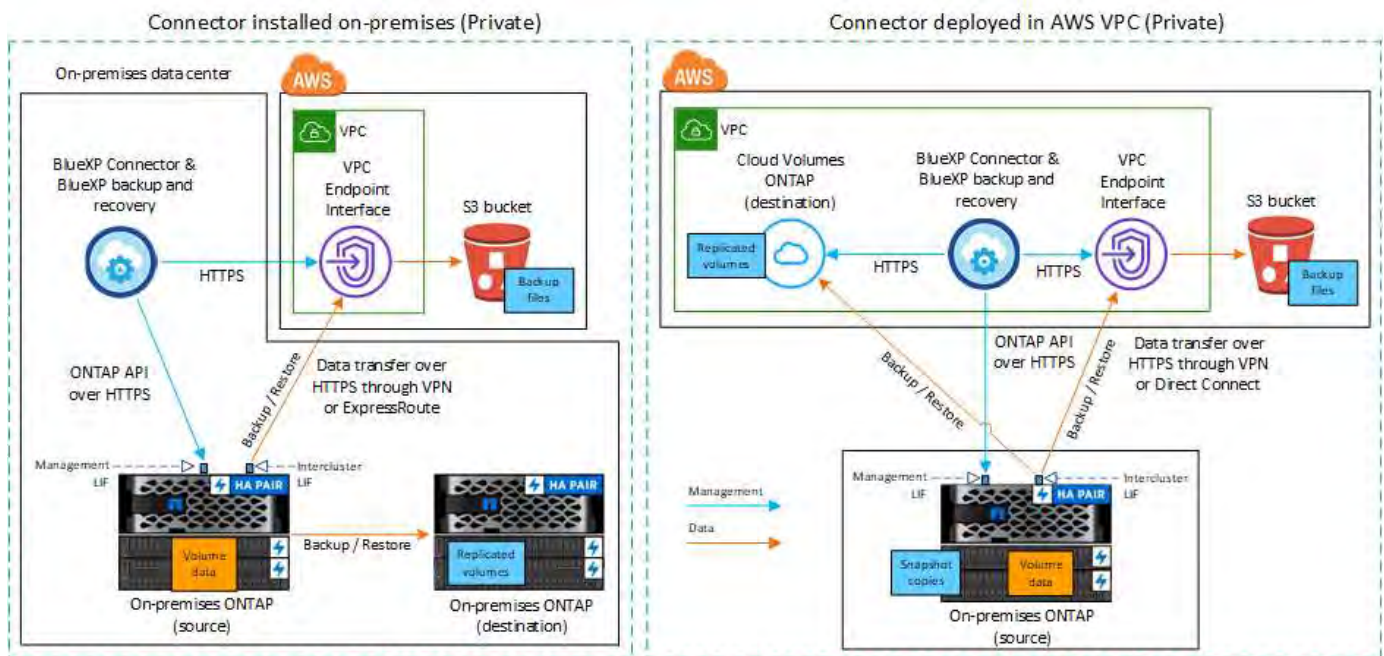
- **Private connection** - Use a VPN or AWS Direct Connect and route traffic through a VPC Endpoint interface that uses a private IP address.

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.
- 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. [Configure your system for a private connection using 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.
- 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, 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 prev-ontap-backup-onprem-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. [Configure your system for a private connection using a VPC endpoint interface.](#)

\*[Ensure that your ONTAP cluster has permissions to access the S3 bucket.]

## Verify ONTAP networking requirements for replicating volumes

Unresolved directive in prev-ontap-backup-onprem-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. Ensure that the Connector has the required permissions. For details, see [BlueXP policy permissions](#).



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 Amazon encryption keys with Cloud Volumes ONTAP.](#)

[Refer to how to use your own Amazon encryption keys with BlueXP backup and recovery.](#)

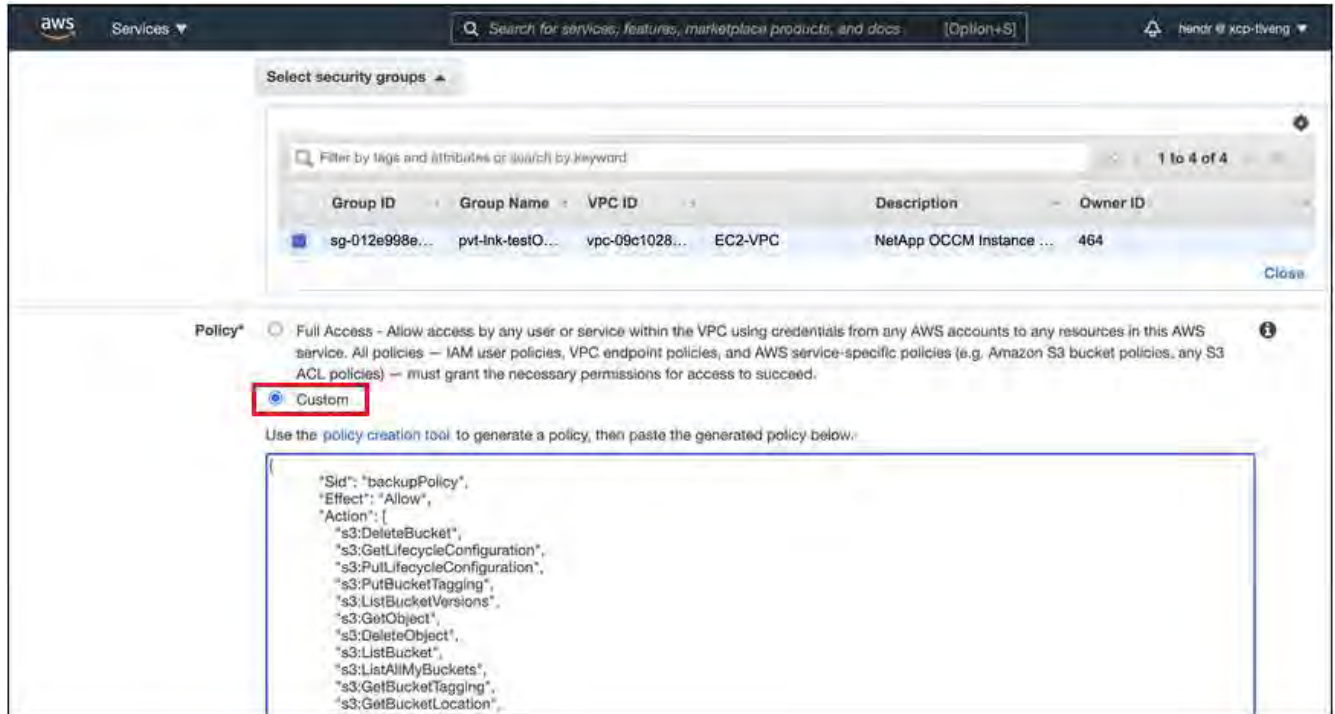
## Configure your system for a private connection using a VPC endpoint interface

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.

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.

## Steps

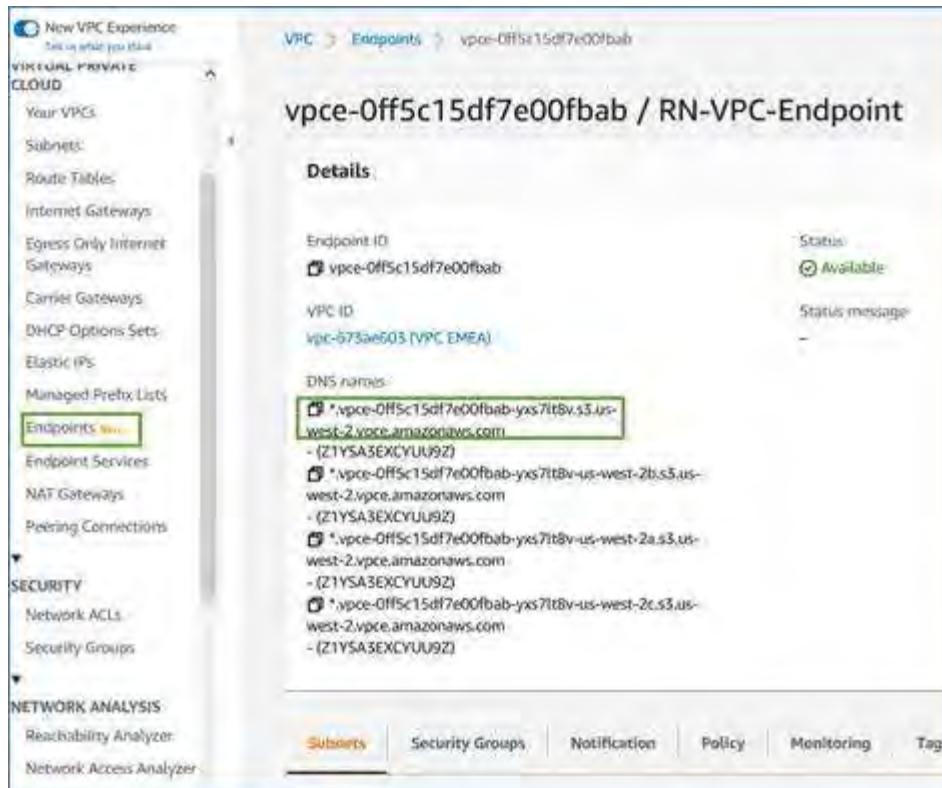
1. Create an Interface endpoint configuration using the Amazon VPC console or the command line. [Refer to details about using AWS PrivateLink for Amazon S3.](#)
2. Modify the security group configuration that's associated with the BlueXP Connector. You must change the policy to "Custom" (from "Full Access"), 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-----
MIIM6zCCC9OgAwIBAgIQA7MGJ4FaD8R8uL0KR3oltTANBgkqhkiG9w0BAQsFADBG
...
...
GqvbOz/oO2NWLlFCqI+xmKlcmiPrZy+/6Af+HH2mLCM4EsI2b+IpBmPkriWnnxo=
-----END CERTIFICATE-----
```

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

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

## Activate backups on your ONTAP volumes

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

A wizard takes you through the following major steps:

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

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

### Start the wizard

#### Steps

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

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

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

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

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

2. Continue with the following options:

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

### Select the volumes that you want to back up

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

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



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

## Steps

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.
  - To back up individual volumes, check the box for each volume.
2. Select **Next**.

## Define the backup strategy

Defining the backup strategy involves setting the following options:

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



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

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

## Steps

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

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

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



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

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

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

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

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



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

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

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

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

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

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

- **Bucket**: Either choose an existing S3 bucket or create a new one. Refer to [Add S3 buckets](#).
- **Encryption key**: If you created a new S3 bucket, enter encryption key information given to you from the provider. Choose whether you'll use the default Amazon S3 encryption keys, or choose your own customer-managed keys from your AWS account, to manage encryption of your data.



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

- **Networking**: Choose the IPspace, and whether you'll be using a Private Endpoint. Private Endpoint is disabled by default.
  - a. The IPspace in the ONTAP cluster where the volumes you want to back up reside. The intercluster LIFs for this IPspace must have outbound internet access.
  - b. Optionally, choose whether you'll use an AWS PrivateLink that you have previously configured. [See details about using AWS PrivateLink for Amazon S3](#).
- **Backup policy**: Select an existing backup policy or create a policy.



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

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

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

### Review your selections

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

### Steps

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

### Result

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

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

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

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

### 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.

## Back up on-premises ONTAP data to Azure Blob storage with BlueXP backup and recovery

Complete a few steps in BlueXP backup and recovery 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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

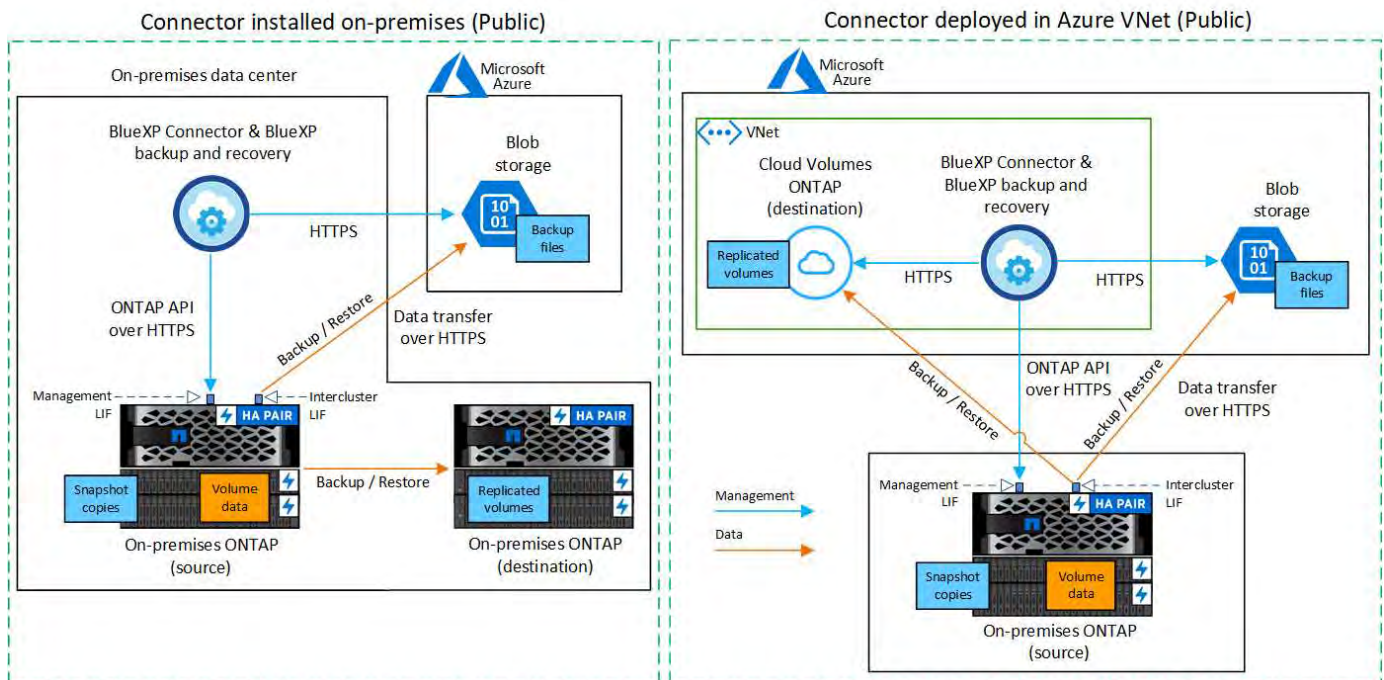
## Identify the connection method

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.

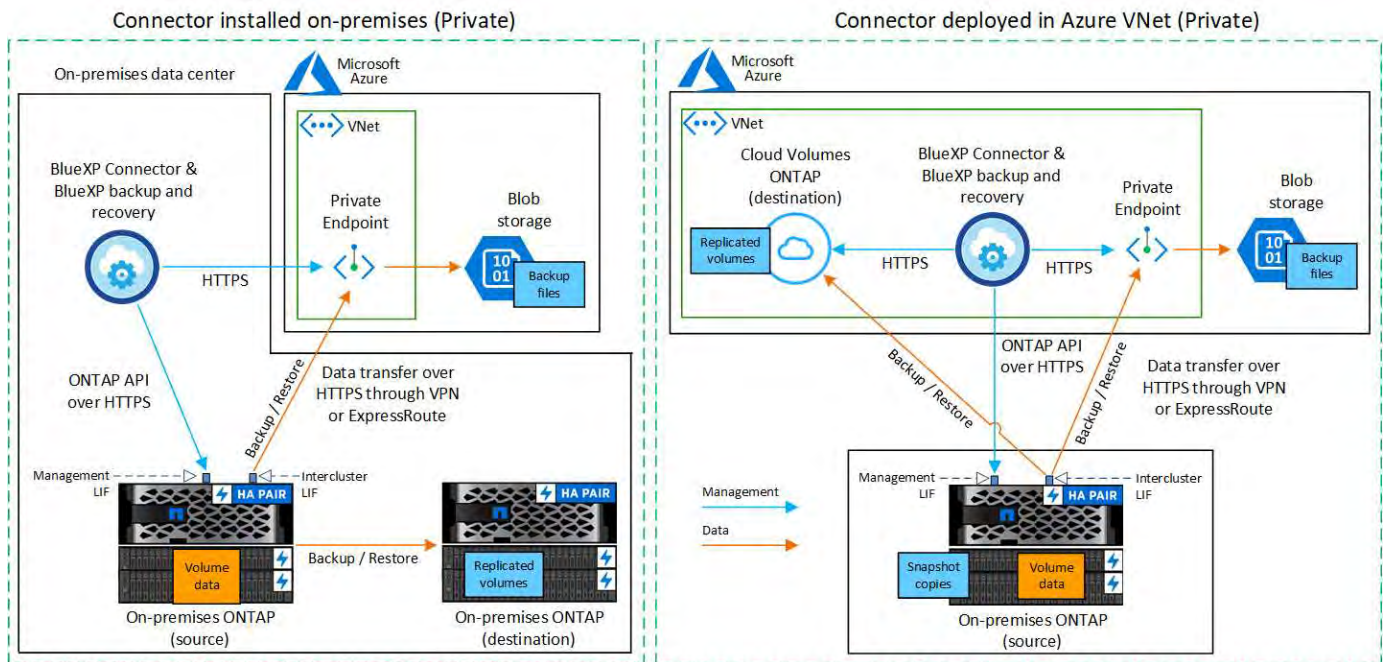
Optionally, you can connect to a secondary ONTAP system for replicated volumes using the public or private connection as well.

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, 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 prev-ontap-backup-onprem-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 prev-ontap-backup-onprem-azure.adoc - include::.../\_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

### Prepare Azure Blob as your backup target

1. You can use your own custom-managed keys for data encryption in the activation wizard instead of using the default Microsoft-managed encryption keys. In this case you will need to have the Azure Subscription, Key Vault name, and the Key. [Learn how to use your own keys](#).

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

2. If you want to have a more secure connection over the public internet from your on-prem data center to the VNet, there is an option to configure an Azure Private Endpoint in the activation wizard. In this case you will need to know the VNet and Subnet for this connection. [Refer to details about using a Private Endpoint](#).

### Create your Azure Blob storage account

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

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

### Activate backups on your ONTAP volumes

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

A wizard takes you through the following major steps:

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

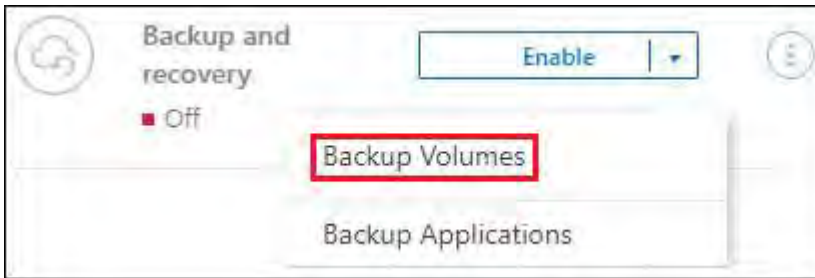
You can also [Show the API commands](#) at the review step, so you can copy the code to automate backup activation for future working environments.

### Start the wizard

#### Steps

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

to the Backup and recovery service in the right-panel.



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

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

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

2. Continue with the following options:

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

### Select the volumes that you want to back up

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

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



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

### Steps

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

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

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

## 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 five 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 five 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 five schedules, typically of different frequencies.
  - For backup-to-object policies, set the DataLock and Ransomware Protection settings. For details on DataLock and Ransomware Protection, refer to [Backup-to-object policy settings](#).
  - Select **Create**.
- **Export existing Snapshot copies to object storage as backup copies:** If there are any local snapshot copies for volumes in this working environment that match the backup schedule label you just



selected for this working environment (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

### Review your selections

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

### Steps

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

### Result

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

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

A Blob storage account is created in the resource group you entered, and the backup files are stored there. The Volume Backup Dashboard is displayed so you can monitor the state of the backups.

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

### 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.

## Back up on-premises ONTAP data to Google Cloud Storage with BlueXP backup and recovery

Complete a few steps in BlueXP backup and recovery 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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

## Identify the connection method

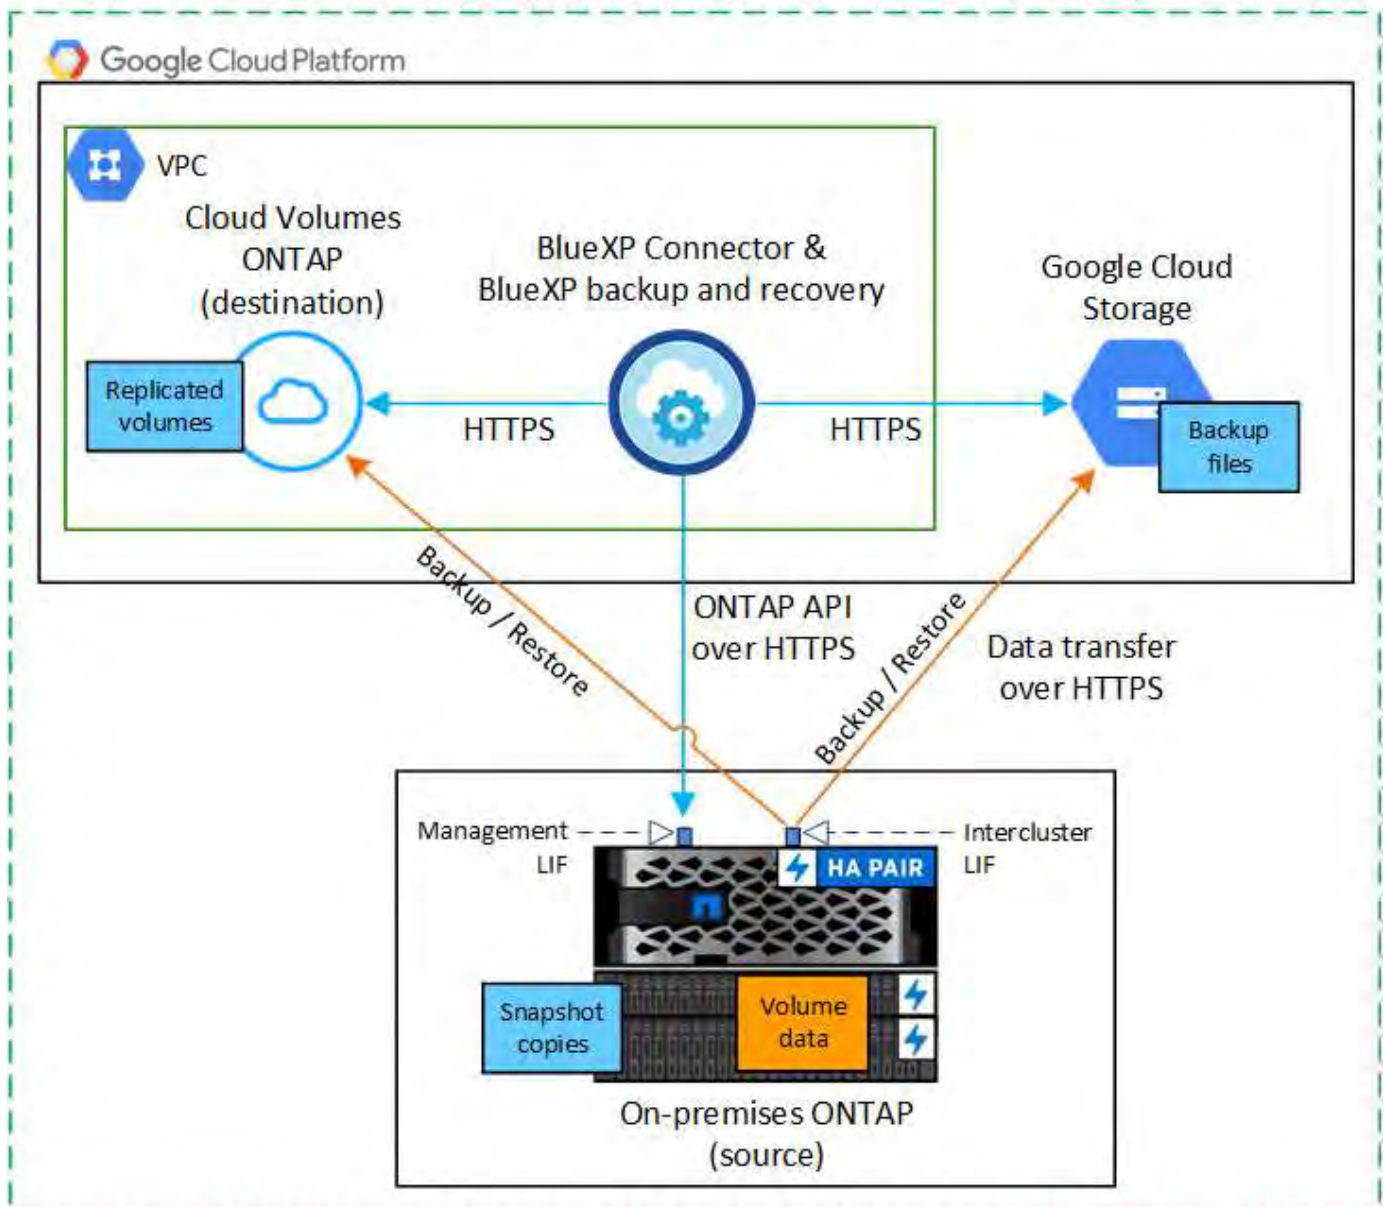
Choose which of the two connection methods you will use when configuring backups from on-premises ONTAP systems to Google Cloud Storage.

- **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.

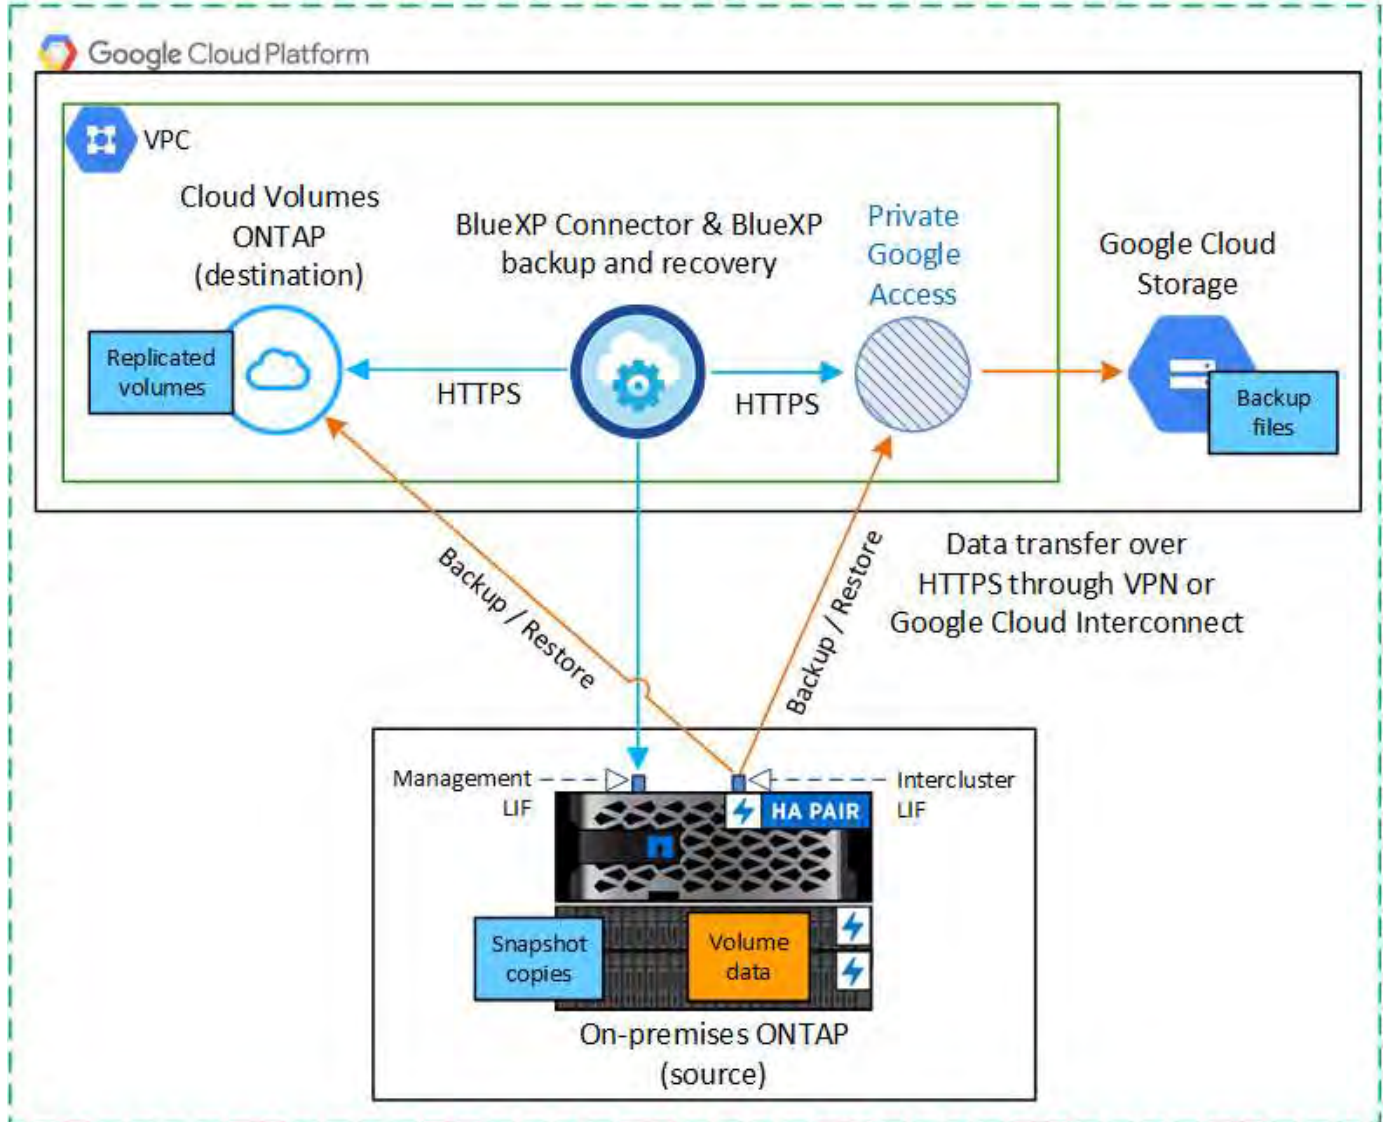
The following diagram shows the **public connection** method and the connections that you need to prepare between the components. The Connector must be deployed in the Google Cloud Platform VPC.

### 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. You specify the region where backups will be stored when you set up the service.

## Prepare your ONTAP clusters

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

### Prepare Google Cloud Storage as your backup target

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

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

#### Set up permissions

You need to provide storage access keys for a service account that has specific permissions using a custom role. A service account enables BlueXP backup and recovery to authenticate and access Cloud Storage buckets used to store backups. The keys are required so that Google Cloud Storage knows who is making the request.

#### Steps

1. In the [Google Cloud Console](#), go to the **Roles** page.
2. [Create a new role](#) with the following permissions:

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

3. In the Google Cloud console, [go to the Service accounts page](#).
4. Select your Cloud project.

5. Select **Create service account** and provide the required information:
  - a. **Service account details**: Enter a name and description.
  - b. **Grant this service account access to project**: Select the custom role that you just created.
  - c. Select **Done**.
6. Go to [GCP Storage Settings](#) and create access keys for the service account:
  - a. Select a project, and select **Interoperability**. If you haven't already done so, select **Enable interoperability access**.
  - b. Under **Access keys for service accounts**, select **Create a key for a service account**, select the service account that you just created, and click **Create Key**.

You'll need to enter the keys in BlueXP backup and recovery later when you configure the backup service.

### Create your own buckets

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

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

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

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

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

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

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

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

### CMEK considerations:

- Both HSM (Hardware-backed) and Software-generated keys are supported.
- Both newly created or imported Cloud KMS keys are supported.

- Only regional keys are supported, global keys are not supported.
- Currently, only the "Symmetric encrypt/decrypt" purpose is supported.
- The service agent associated with the Storage Account is assigned the "CryptoKey Encrypter/Decrypter (roles/cloudkms.cryptoKeyEncrypterDecrypter)" IAM role by BlueXP backup and recovery.

## Activate backups on your ONTAP volumes

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

A wizard takes you through the following major steps:

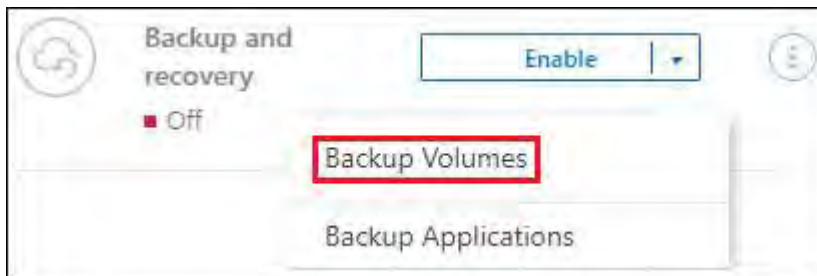
- [Select the volumes that you want to back up](#)
- [Define the backup strategy](#)
- [Review your selections](#)

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

### Start the wizard

#### Steps

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



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

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

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

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



## Select the volumes that you want to back up

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

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



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

## Steps

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.
  - To back up individual volumes, check the box for each volume.
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, 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 five 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, 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 five 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, 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 five 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 page](#).

### 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.

## Back up on-premises ONTAP data to ONTAP S3 with BlueXP backup and recovery

Complete a few steps in BlueXP backup and recovery 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.

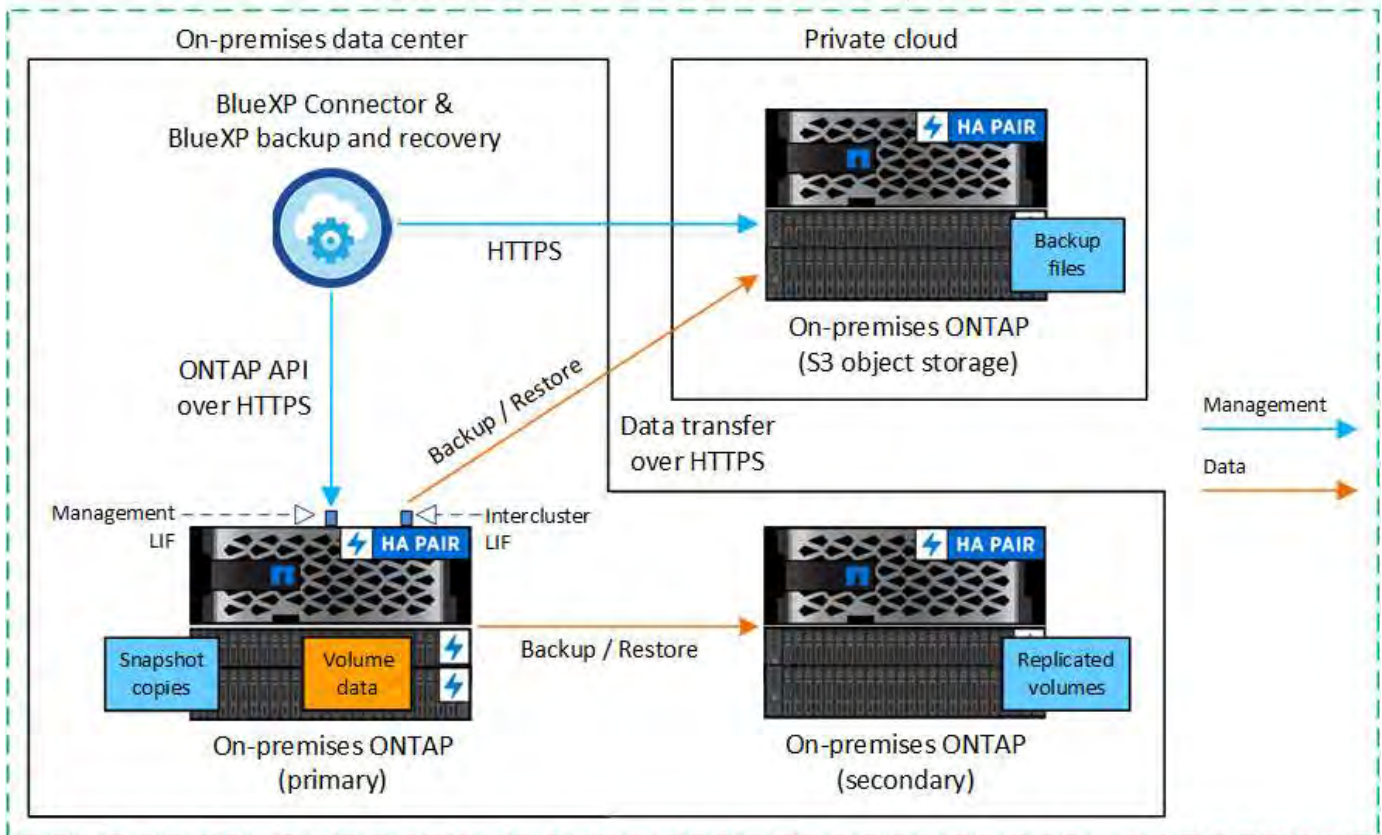
**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### 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.

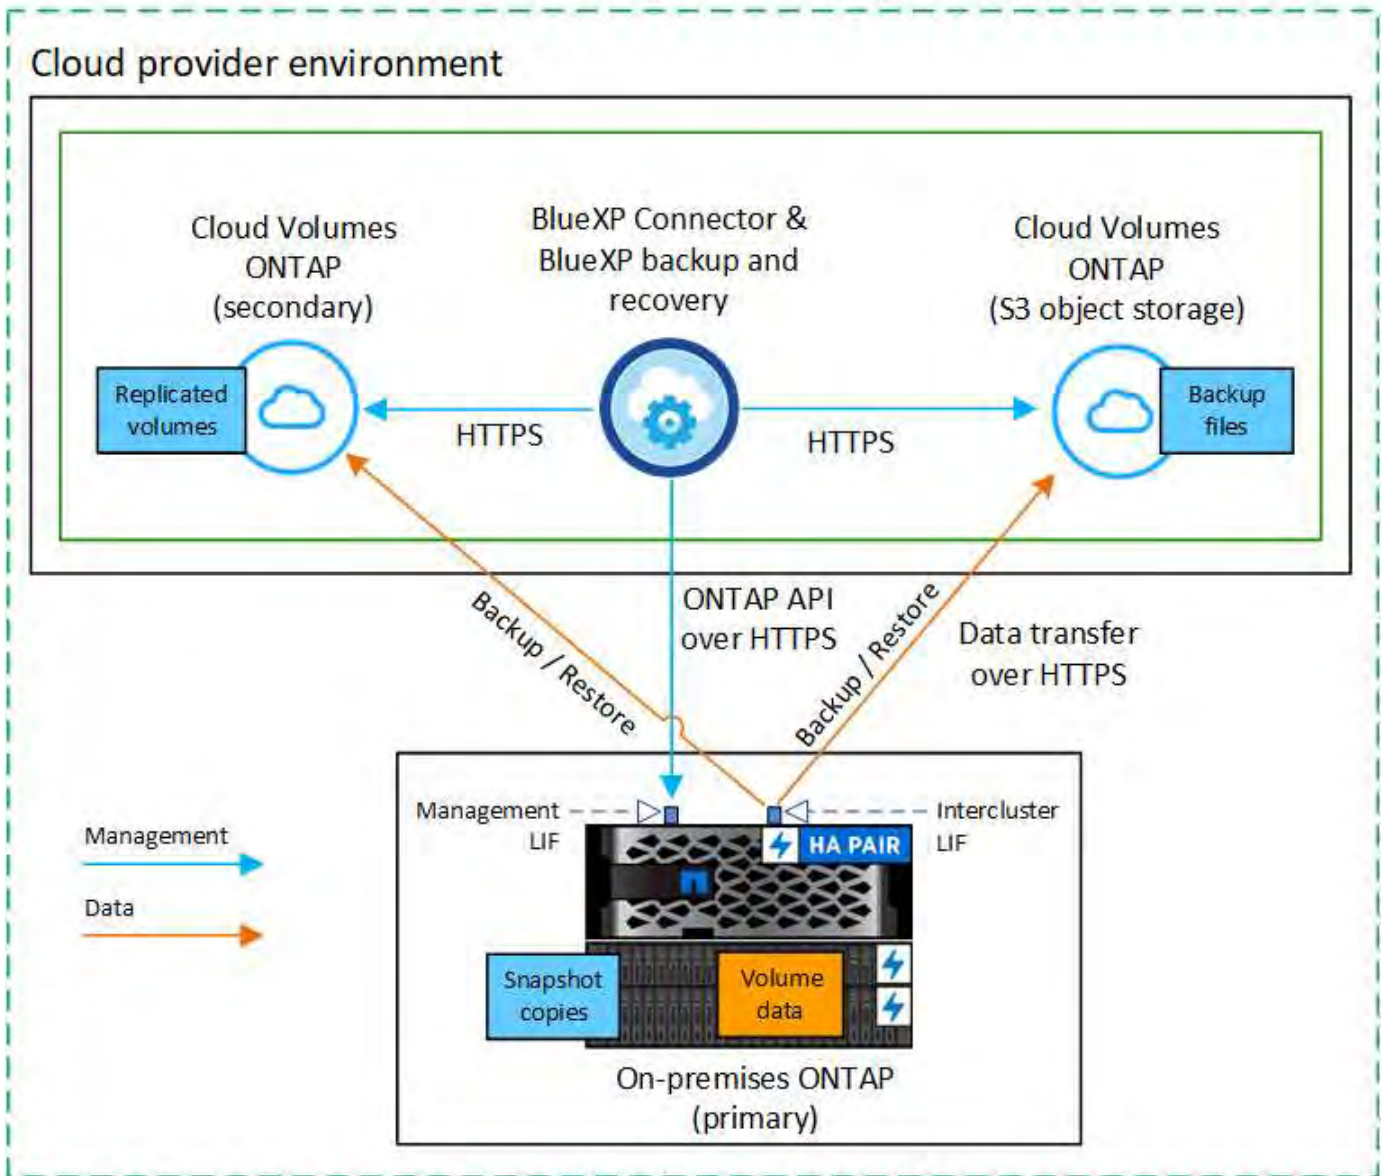
## Connector installed on-premises (Public)



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.

### 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 prev-ontap-backup-onprem-ontaps3.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 use 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 prev-ontap-backup-onprem-ontaps3.adoc - include::.../\_include/backup-onprem-prepare-source-destination-systems-for-repl.adoc[]

#### Prepare ONTAP S3 as your backup target

You must enable a Simple Storage Service (S3) object storage server in the ONTAP cluster that you plan to use for object storage backups. See the [ONTAP S3 documentation](#) for details.

**Note:** You can discover this cluster to the BlueXP Canvas, but it is not identified as being an S3 object storage server, and you can't drag and drop a source working environment onto this S3 working environment to initiate backup activation.

This ONTAP system must meet the following requirements.



## Supported ONTAP versions

ONTAP 9.8 and later is required for on-premises ONTAP systems.

ONTAP 9.9.1 and later is required for Cloud Volumes ONTAP systems.

## S3 credentials

You must have created an S3 user to control access to your ONTAP S3 storage. [See the ONTAP S3 docs for details.](#)

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

These access keys must be associated with a user who has the following permissions:

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

## Activate backups on your ONTAP volumes

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

A wizard takes you through the following major steps:

- Select the volumes that you want to back up
- Define the backup strategy and policies
- Review your selections

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

### Start the wizard

#### Steps

1. Access the Activate backup and recovery wizard using one of the following ways:
  - From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.
  - Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions (...)** option and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage enabled).

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

2. Continue with the following options:

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

### Select the volumes that you want to back up

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

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



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

### Steps

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

1. In the Select Volumes page, select the volume or volumes you want to protect.
  - Optionally, filter the rows to show only volumes with certain volume types, styles, and more to make the selection easier.
  - After you select the first volume, then you can select all FlexVol volumes (FlexGroup volumes can be selected one at a time only). To back up all existing FlexVol volumes, check one volume first and then check the box in the title row.
  - To back up individual volumes, check the box for each volume.
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, 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 five 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 five 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 using this service, 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 five schedules, typically of different frequencies.
- For backup-to-object policies, set the DataLock and Ransomware Protection settings. For details on DataLock and Ransomware Protection, refer to [Backup-to-object policy settings](#).
- Select **Create**.
- **Export existing Snapshot copies to object storage as backup files:** If there are any local snapshot copies for volumes in this working environment that match the backup schedule label you just selected (for example, daily, weekly, etc.), this additional prompt is displayed. Check this box to have all historic Snapshots copied to object storage as backup files to ensure the most complete protection for your volumes.

6. Select **Next**.

### Review your selections

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

### Steps

1. In the Review page, review your selections.
2. Optionally check the box to **Automatically synchronize the Snapshot policy labels with the replication and backup policy labels**. This creates Snapshots with a label that matches the labels in the replication and backup policies. If the policies don't match, backups will not be created.
3. Select **Activate Backup**.

### Result

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

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

An S3 bucket is created in the service account indicated by the S3 access key and secret key you entered, and the backup files are stored there.

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

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

### 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.

## Back up on-premises ONTAP data to StorageGRID with BlueXP backup and recovery

Complete a few steps in BlueXP backup and recovery 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.



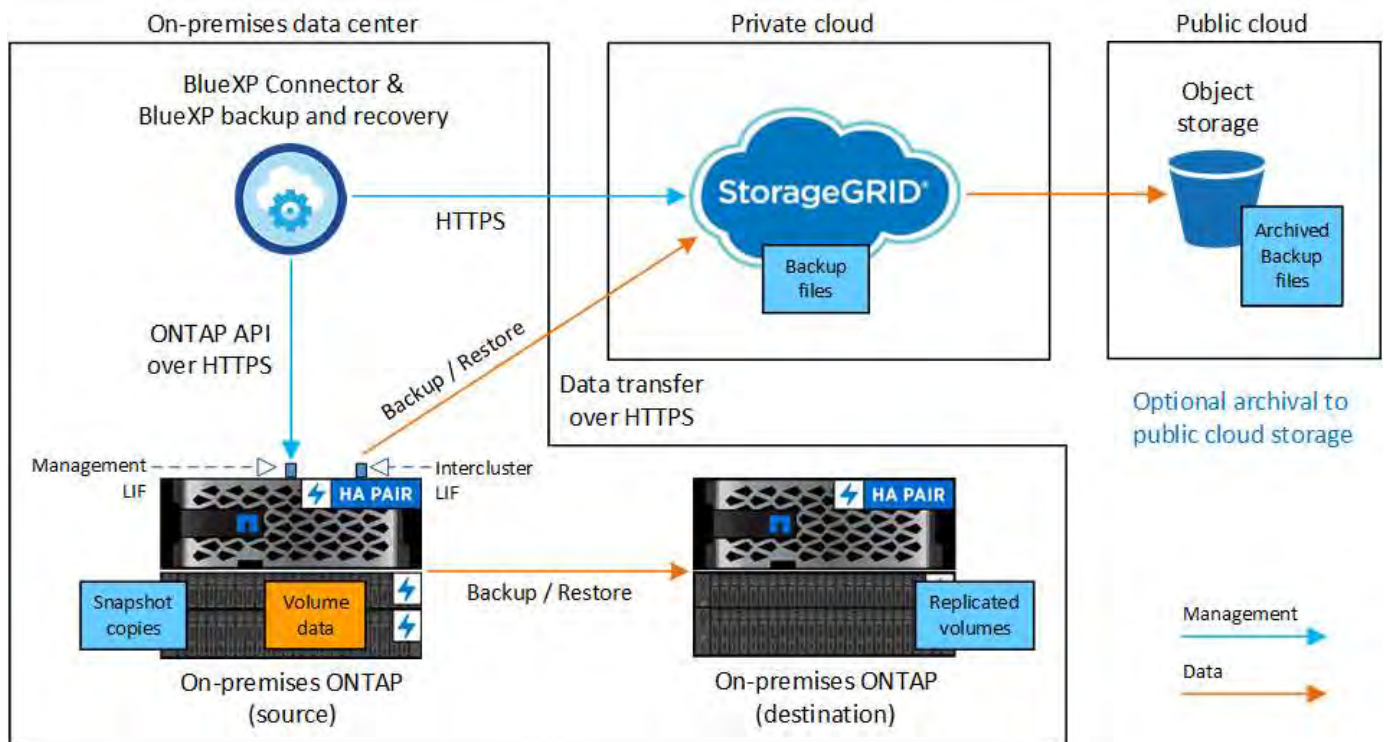
"On-premises ONTAP systems" include FAS, AFF, and ONTAP Select systems.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### 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.

### 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 prev-ontap-backup-onprem-storagegrid.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 use are using a different IPspace than the Default, then you might need to create a static route to get access to the object storage.
- Update firewall rules, if necessary, to allow 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 prev-ontap-backup-onprem-storagegrid.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.



For details about DataLock and Ransomware Protection requirements for StorageGRID, refer to [Backup-to-object policy options](#).

### 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.

To use archival storage, admin node IP access is needed.

Gateway IP access is always needed.

### 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 AWS	Cloud Provider AZURE
Account Select Account	Azure Subscription Select Account
Region Select Region	Region Select Region
AWS Access Key Enter AWS Access Key	Resource Group Type Select an Existing Resource Group
AWS Secret Key Enter AWS Secret Key	Resource Group Select Resource Group
Archive After (Days) (1-999)	Archive After (Days) (1-999)
Storage Class S3 Glacier	Storage Class Azure Archive

The archival policy settings you select will generate an information lifecycle management (ILM) policy in

StorageGRID, and add the settings as "rules."

- If there is an existing active ILM policy, new rules will be added to the ILM policy to move the data to the archive tier.
- If there is an existing ILM policy in the "proposed" state, the creation and activation of a new ILM policy will not be possible. [Learn more about StorageGRID ILM policies and rules.](#)

## Activate backups on your ONTAP volumes

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

A wizard takes you through the following major steps:

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

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

### Start the wizard

#### Steps

1. Access the Activate backup and recovery wizard using one of the following ways:
  - From the BlueXP canvas, select the working environment and select **Enable > Backup Volumes** next to the Backup and recovery service in the right-panel.  
  
If the destination for your backups exists as a working environment on the Canvas, you can drag the ONTAP cluster onto the object storage.
  - Select **Volumes** in the Backup and recovery bar. From the Volumes tab, select the **Actions (...)** option and select **Activate Backup** for a single volume (that does not already have replication or backup to object storage already enabled).

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

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

### Select the volumes that you want to back up

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

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



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

## Steps

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.
  - To back up individual volumes, check the box for each volume.
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, 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 five 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 five 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, 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 five schedules, typically of different frequencies.
- For backup-to-object policies, set the DataLock and Ransomware Protection settings. For details on DataLock and Ransomware Protection, refer to [Backup-to-object policy settings](#).

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

protects your backup files from being modified or deleted, and *Ransomware Protection* scans your backup files to look for evidence of a ransomware attack in your backup files.

- Select **Create**.

If your cluster is using ONTAP 9.12.1 or greater and your StorageGRID system is using version 11.4 or greater, you can choose to tier older backups to public cloud archive tiers after a certain number of days. Current support is for AWS S3 Glacier/S3 Glacier Deep Archive or Azure Archive storage tiers. [See how to configure your systems for this functionality](#).

- **Tier backup to public cloud:** Select the cloud provider that you want to tier backups to and enter the provider details.

Select or create a new StorageGRID cluster. For details about creating a StorageGRID cluster so BlueXP can discover it, refer to [StorageGRID documentation](#).

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

6. Select **Next**.

### Review your selections

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

### Steps

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

### Result

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

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

An S3 bucket is created in the service account indicated by the S3 access key and secret key you entered, and the backup files are stored there.

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

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

### 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.

## Migrate volumes using SnapMirror to Cloud Resync with BlueXP backup and recovery

The SnapMirror to Cloud Resync feature in BlueXP backup and recovery streamlines data protection and continuity during volume migrations in NetApp environments. When a volume is migrated using SnapMirror Logical Replication (LRSE), from one on-premises NetApp deployment to another, or to a cloud-based solution such as Cloud Volumes ONTAP or Cloud Volumes Service, SnapMirror to Cloud Resync ensures that existing cloud backups remain intact and operational.

This feature eliminates the need for a time-consuming and resource-intensive re-baseline operation, enabling backup operations to continue post-migration. This feature is valuable in workload migration scenarios, supporting both FlexVols and FlexGroups, and is available starting with ONTAP version 9.16.1.



This feature is available starting with BlueXP backup and recovery version 4.0.3 released May 2025.

By maintaining backup continuity across environments, SnapMirror to Cloud Resync enhances operational efficiency and reduces the complexity of hybrid and multi-cloud data management.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### Before you begin

Ensure that these prerequisites have been met:

- The destination ONTAP cluster must be running ONTAP version 9.16.1 or later.
- The old source ONTAP cluster must be protected using BlueXP backup and recovery.
- The SnapMirror to Cloud Resync feature is available starting with BlueXP backup and recovery version 4.0.3 released May 2025.
- The latest backup in the object storage must be the common snapshot across the old source, the new source, and the object store. The common snapshot cannot be older than the latest snapshot that is backed up to the object store.
- Both the snapshot and SnapMirror policies, which were used on the older ONTAP must be created on the new ONTAP cluster before starting the resync operation. If any policy is going to be used in the resync process, then that policy must also be created. The Resync operation does not create the policies.
- Ensure that the SnapMirror policy that is applied to the migration volume SnapMirror relationship includes the same label that the cloud relationship uses. To avoid issues, use the policy that governs an exact mirror of the volume and all snapshots.



SnapMirror to Cloud Resync after migrations using SVM-Migrate, SVM-DR, or Head Swap methods are not currently supported.

## How BlueXP backup and recovery SnapMirror to Cloud Resync works

If you complete a technical refresh or migrate volumes from one ONTAP cluster to another, it's important that your backups continue to work without interruption. BlueXP backup and recovery SnapMirror to Cloud Resync helps with this by ensuring that your cloud backups stay consistent even after a volume migration.

Here's an example:

Imagine you have an on-premises volume called Vol1a. This volume has three snapshots: S1, S2, and S3. These snapshots are like restore points. Vol1 is already being backed up to a cloud object store endpoint using SnapMirror to Cloud (SM-C). However, only S1 and S2 have been backed up to object store so far.

Now, you want to migrate Vol1 to another ONTAP cluster. To do this, you create a SnapMirror Logical Replication (LRSE) relationship to a new cloud volume called Vol1b. This transfers all three snapshots—S1, S2, and S3—from Vol1a to Vol1b.

After the migration is complete, you have the following setup:

- The original SM-C relationship (Vol1a → Object store) is deleted.
- The LRSE relationship (Vol1a → Vol1b) is also deleted.
- Vol1b is now your active volume.

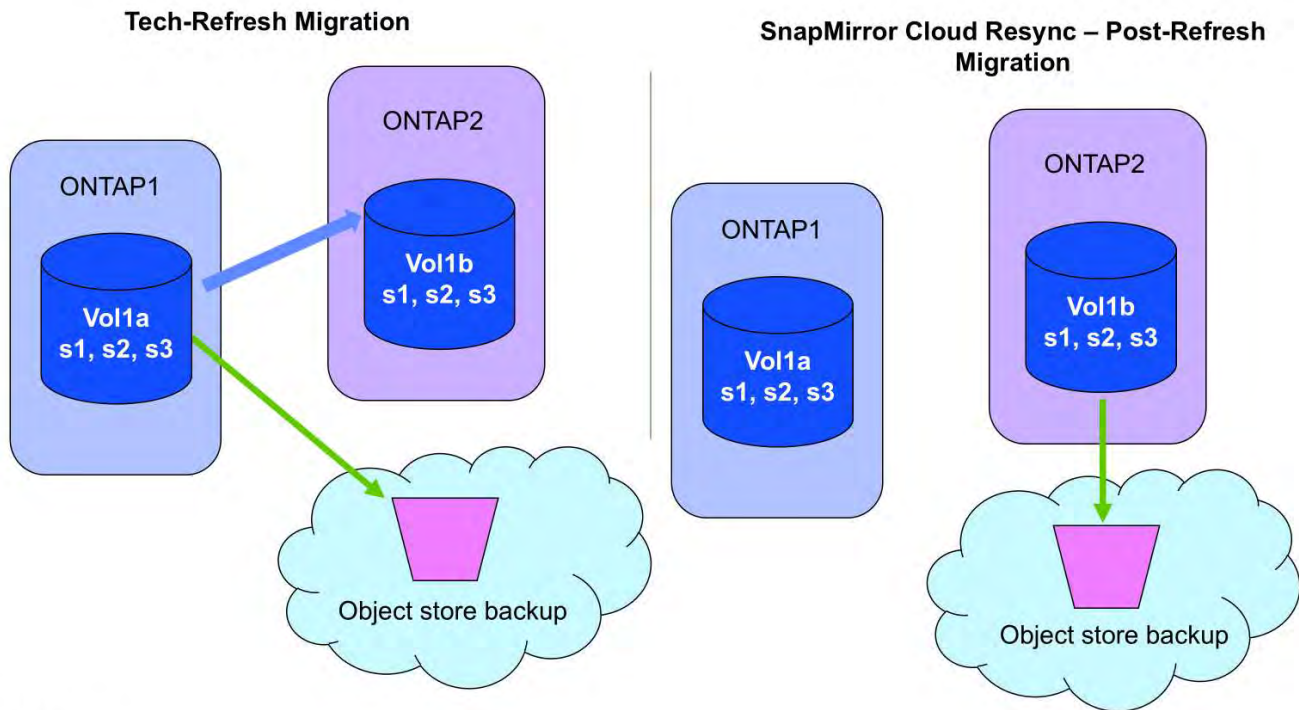
At this point, you want to continue backing up Vol1b to the same cloud endpoint. But instead of starting a full backup from scratch (which would take time and resources), you use SnapMirror to Cloud Resync.

Here's how the resync works:

- The system checks for a common snapshot between Vol1a and Object store. In this case, both have S2.
- Because of this shared snapshot, the system needs to transfer only the incremental changes between S2 and S3.

This means only the new data added after S2 is sent to object store, not the entire volume.

This process avoids re-sending data that's already backed up, saves bandwidth, and ensures that your backup chain continues smoothly after migration.



NetApp

## Procedure notes

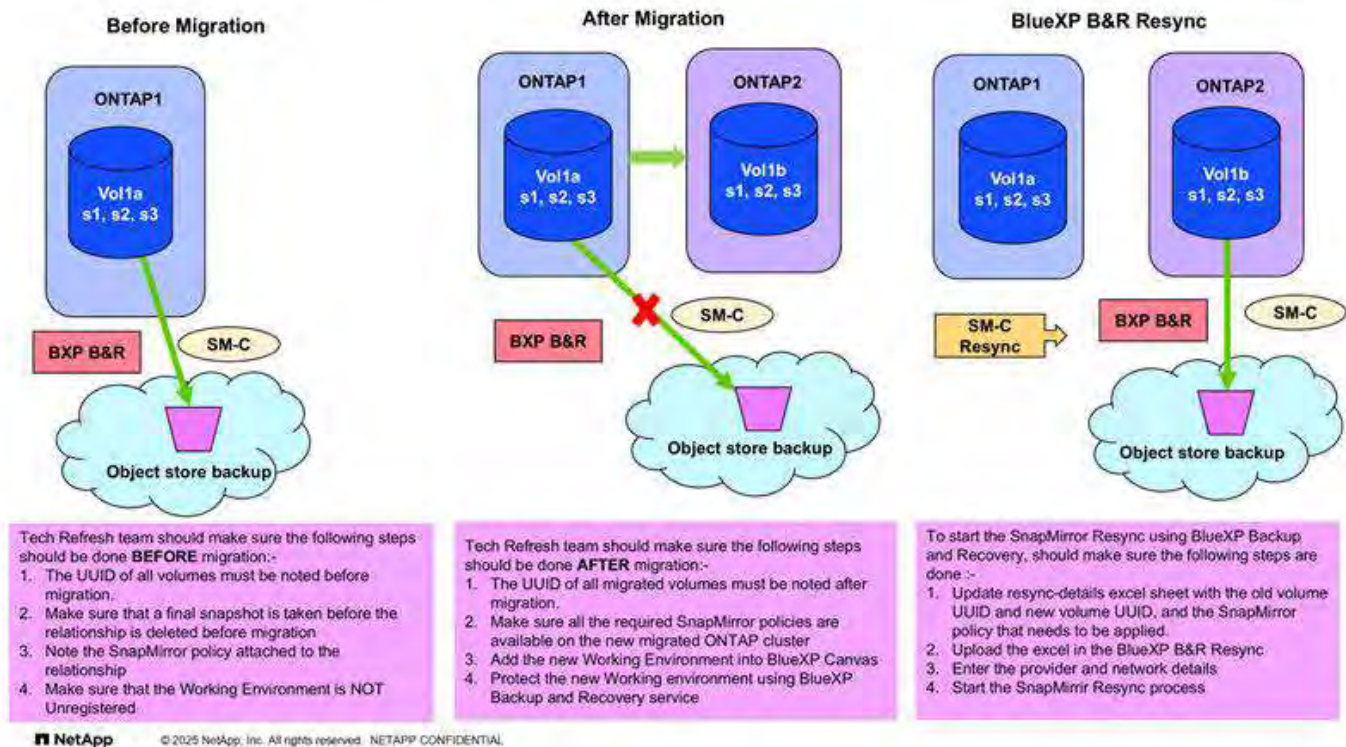
- Migrations and tech refreshes are not performed using BlueXP backup and recovery. They should be carried out by a professional services team or a qualified storage administrator.
- A NetApp migration team is responsible for creating the SnapMirror relationship between the source and destination ONTAP clusters to facilitate volume migration.
- Ensure that the migration during a tech refresh is based on SnapMirror-based migration.

## How to migrate volumes using SnapMirror to Cloud Resync

Migrating volumes using SnapMirror to Cloud Resync involves the following major steps, each described in more detail below:

- **Follow a pre-migration checklist:** Before starting the migration, a NetApp Tech Refresh team ensures the following prerequisites are met to avoid data loss and ensure a smooth migration process.
- **Follow a post-migration checklist:** After the migration, a NetApp Tech Refresh team ensures the following steps are completed to establish protection and prepare for the resync.
- **Perform a SnapMirror to Cloud Resync:** After the migration, a NetApp Tech Refresh team performs a SnapMirror to Cloud Resync operation to resume cloud backups from the newly migrated volumes.





### Follow a pre-migration checklist

Before starting the migration, a NetApp Tech Refresh team ensures the following prerequisites are met to avoid data loss and ensure a smooth migration process.

1. Ensure all volumes that are to be migrated are protected using BlueXP backup and recovery.
2. Record volume instance UUIDs. Write down the Instance UUIDs of all volumes before starting the migration. These identifiers are crucial for mapping and resync operations later.
3. Take a final snapshot of each volume to preserve the latest state, before deleting any SnapMirror relationships.
4. Document SnapMirror policies. Record the SnapMirror policy currently attached to each volume's relationship. This will be needed later during the SnapMirror to Cloud Resync process.
5. Delete the SnapMirror Cloud relationships with the object store.
6. Create a standard SnapMirror relationship with the new ONTAP cluster to migrate the volume to the new target ONTAP cluster.

### Follow a post-migration checklist

After the migration, a NetApp Tech Refresh team ensures the following steps are completed to establish protection and prepare for the resync.

1. Record new volume instance UUIDs of all migrated volumes in the destination ONTAP cluster.
2. Confirm that all required SnapMirror policies that were available in the old ONTAP cluster are correctly configured in the new ONTAP cluster.
3. Add the new ONTAP cluster as a working environment in the BlueXP canvas.



The volume instance UUID should be used, not the volume ID. The volume instance UUID is a unique identifier that remains consistent across migrations, while the volume ID may change after migration.

### Perform a SnapMirror to Cloud Resync

After the migration, a NetApp Tech Refresh team performs a SnapMirror to Cloud Resync operation to resume cloud backups from the newly migrated volumes.

1. Add the new ONTAP cluster as a working environment in the BlueXP canvas.
2. Look at the BlueXP backup and recovery Volumes page to ensure that the old source working environment details are available.
3. From the BlueXP backup and recovery Volumes page, select **Backup Settings**.
  - Within the Backup Settings page, select **View all**.
  - From the Actions ... menu to the right of the *new* source, select **Resync backup**.
4. In the Resync Working Environment page, do the following:
  - a. **New source working environment:** Enter the new ONTAP cluster where the volumes have been migrated.
  - b. **Existing Target Object Store:** Select the target object store that contains the backups from the old source working environment.
5. Select **Download CSV Template** to download the Resync Details Excel sheet. Use this sheet to enter the details of the volumes to be migrated. In the CSV file, enter the following details:
  - The old volume instance UUID from the source cluster
  - The new volume instance UUID from the destination cluster
  - The SnapMirror policy to be applied to the new relationship.
6. Select **Upload** under the **Upload Volume Mapping Details** to upload the completed CSV sheet into the BlueXP backup and recovery UI.



The volume instance UUID should be used, not the volume ID. The volume instance UUID is a unique identifier that remains consistent across migrations, while the volume ID may change after migration.

7. Enter provider and network configuration information required for the resync operation.
8. Select **Submit** to start the validation process.

BlueXP backup and recovery validates that each volume selected for resync is the latest snapshot and has at least one common snapshot. This ensures that the volumes are ready for the SnapMirror to Cloud Resync operation.

9. Review validation results including the new source volume names and the resync status for each volume.
10. Check volume eligibility. The system checks if the volumes are eligible for resync. If a volume is not eligible, it means that it isn't the latest snapshot or no common snapshot was found.



To ensure that volumes remain eligible for the SnapMirror to Cloud Resync operation, take a final snapshot of each volume before deleting any SnapMirror relationships during the pre-migration phase. This preserves the latest state of the data.

11. Select **Resync** to start the resync operation. The system uses the latest and common snapshot to transfer only the incremental changes, ensuring backup continuity.
12. Monitor the resync process in the Job Monitor page.

## Restore BlueXP backup and recovery configuration 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, you can deploy a new Connector and restore the critical BlueXP backup and recovery data.



This procedure applies only to ONTAP volume data.

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 two 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 BlueXP Connector

If your on-premises BlueXP 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.

You'll need to perform the following tasks 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 one of the following commands, depending on whether you are using a Docker or Podman container:

```
docker cp CBS_DB_Backup_23_05_2023.sql ds_mysql_1:/.
```

```
podman cp CBS_DB_Backup_23_05_2023.sql ds_mysql_1:/.
```

3. Enter the MySQL container shell using one of the following commands, depending on whether you are using a Docker or Podman container:

```
docker exec -it ds_mysql_1 sh
```

```
podman 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](#).

The following information applies to private mode installations starting from BlueXP 3.9.xx. For older versions, use the following procedure: [DarkSite Cloud Backup: MySQL and Indexed Catalog Backup and Restore](#).

You'll need to perform these steps for each 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 '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"}' > '
```

While the IP address, username, and passwords are custom values, the account name is not. The account name is always "account-DARKSITE1". Also, the username must use an email-formatted name.

This API will return a response like the following. You can retrieve the authorization token as shown below.

```
{"expires_in":21600,"access_token":"eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjRiIn0eyJzdWIiOiJvY2NtYXV0aHwxIiwiaXVkiIjpbImh0dHBzOi8vYXBpLmNsb3VkLm5ldGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm5ldGFwcC5jb20vZnVsbnVsbF9uYW11IjoiYWRTaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS9lbWFpbCI6ImFkbWluQG5ldGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwcm9maWx1IiwiaWF0IjoxNjcyNzY2MDIzLCJleHAiOiJlE2NzI3NTc2MjMsImIzcyI6Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9CJtRpRDY23PokyLg1if67bmgnMcYxdCvBOY-ZUYWzhrWbbY_hqUH4T-114v_pNDsPyNDyWqHaKizThdjYHYHxm56vTz_Vdn4NqjaBDPwN9KAnC6Z88WA1cJ4WRQqj5yKODNDmrv5At_f9HHp0-xVMYHqywZ4nNFAlMvAh4xESc5jfoKOZc-IOQdWm4F4LHpMzs4qFzCYthTuSKLYtqSTURZB81-o-ipvrOqSolIwIeHXZJJV-Uswun9daNgiYd_wX-4WWJVIGEnDzzwOKfUoUoelFg3ch--7JFkFl-rrXDOjk1sUMumN3WHV9usp1PgBE5HAcJPrEBm0ValSZcUbiA"}
```

## 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 eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjRiIn0eyJzdWIiOiJvY2NtYXV0aHwxIiwiaXVkiIjpbImh0dHBzOi8vYXBpLmNsb3VkLm5ldGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm5ldGFwcC5jb20vZnVsbnVsbF9uYW11IjoiYWRTaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS9lbWFpbCI6ImFkbWluQG5ldGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwcm9maWx1IiwiaWF0IjoxNjcyNzY2MDIzLCJleHAiOiJlE2NzI3NDQzMTMsImIzcyI6Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdu_kN-flWpdJJX98HODwPpVUitLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-yE0fH9gr4XgkdsWjWcNvw2rRkFzjHpWrETgfqAMkZcAukV4DHuxogHWh6-DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zC-sp81GaqMahPf0KcFVYjbBL4krOewgKHGfO_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-KbZqmmBX9vDnYp7SSxClhHJRdstcFgJLdJHtowweNH2829KsjEGBTtcBd08SvIDtctNH_GAxwSgMT3zUfwaOimPw'
```

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 eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjRiIn0eyJzdWIiOiJvY2NtYXV0aHwxIiwiaXVkaWVzImh0dHBzOi8vYXBpLmNsb3VkLm5ldGFwC5jb20iXSwiaHR0cDovL2Nsb3VkLm5ldGFwC5jb20vZnVsbF9uYW11IjoiYWRTaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS9lbWVpY29wZSI6Im9wZW5pZCBwcm9maWxlIiwiaWF0IjoxNjcyNzIyNzI3NDQzMjM5ImlzcyI6Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdu_kN-fLWpdJJX98HODwPpVUitLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-yE0fH9gr4XgkdsWjWcNvw2rRkFzjHpWrETgfqAMkZcAukV4DHuxogHWh6-DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zC-sp81GaqMahPf0KcFVYjBBL4krOewgKHGfo_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-KbZqmmBX9vDnYp7SSxClhHJRDStcFgJLdJHTowweNH2829KsjEGBTtcBd08SvIDtctNH_GAxwSgMT3zUfwaOimPw' \  
> --header 'x-agent-id: vB_1xShPpBtUosjD7wfB1LIhqDgIPA0wclients' \  
> -d '{ "storage-server" : "sr630ip15.rtp.eng.netapp.com:10443", "access-key": "2ZMYOAVAS5E70MCNH9", "secret-password": "uk/6ikd4Lj1XQOfnzSzP/T0zR4ZQ1G0w1xgWsB" }'
```

### 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.

### Manage backups for your ONTAP systems with BlueXP backup and recovery

With BlueXP backup and recovery, 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, force deleting backups, and more.

This includes all types of backups, including snapshot copies, replicated volumes, and backup files in object storage. You can also unregister BlueXP backup and recovery.



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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

## 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. Select 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 the plus sign.

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.

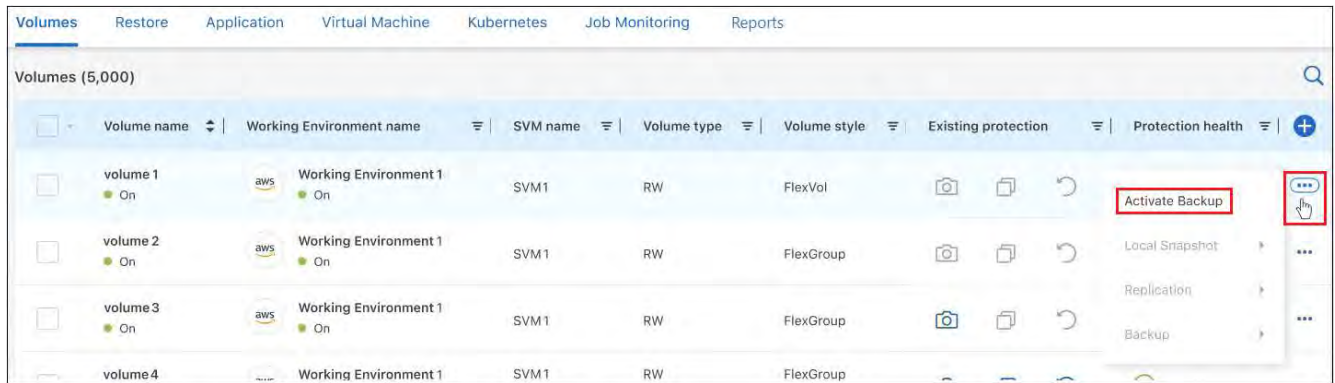
## Activate backup on additional volumes in a working environment

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**.





- 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 select **Next**.
- Review the backup settings for this volume, and then select **Activate Backup**.

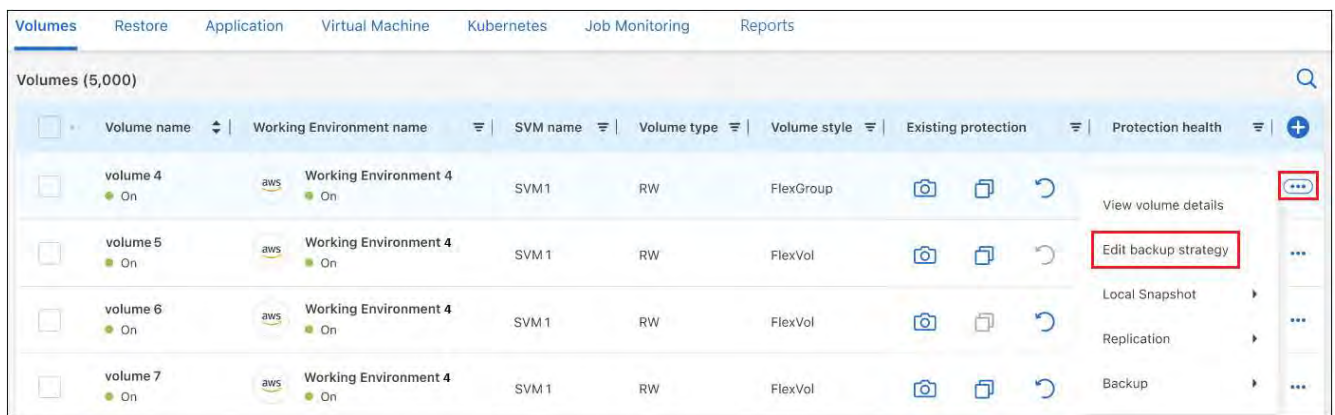
### Change the backup settings assigned to existing volumes

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

- 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**.



- In the *Edit backup strategy* page, make changes to the existing backup policies for Local Snapshot copies, Replicated volumes, and Backup files and select **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.

- Review the backup settings for this volume, and then select **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)

When more than five volumes are enabled for backup, BlueXP backup and recovery initializes only five volumes at a time. When those are finished, it creates the next batch of five subjobs to start the next set and continues until all volumes are initialized.

### 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...	Select this button...
Manage snapshot backup settings	<b>Manage Local Snapshots</b>
Manage replication backup settings	<b>Manage Replication</b>
Manage backup to cloud backup settings	<b>Manage Backup</b>
Manage multiple types of backup settings. This option enables you to change the backup architecture as well.	<b>Manage Backup and Recovery</b>

4. In the backup page that appears, make changes to the existing backup policies for Local Snapshot copies, Replicated volumes, or Backup files and select **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.](#)

When you create an ad-hoc backup, a snapshot is created on the source volume. Because 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, select **...** for the volume and select **Backup > Create Ad-hoc Backup**.

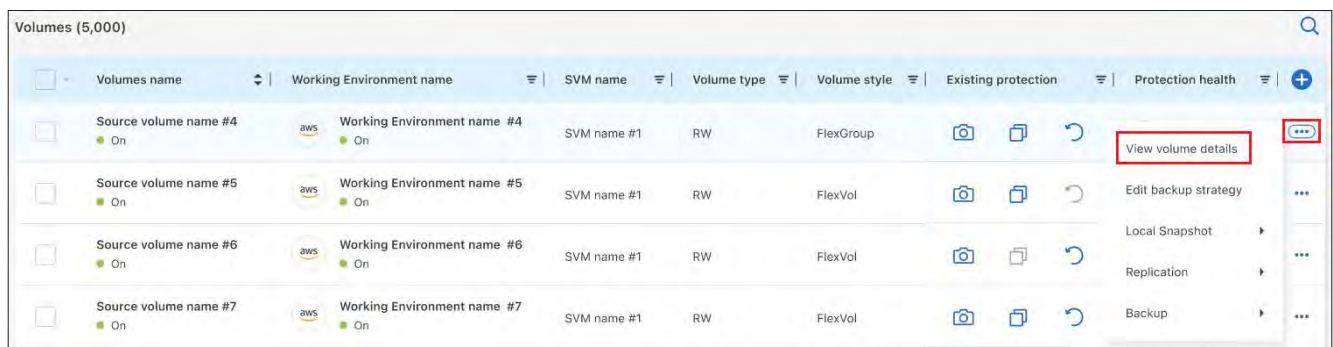
The Backup Status column for that volume displays "In Progress" until the backup is created.

## 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, select **...** for the source volume and select **View volume details**.



The details for the volume and the list of snapshot copies are displayed.

2. Select **Snapshot**, **Replication**, or **Backup** to see the list of all backup files for each type of backup.

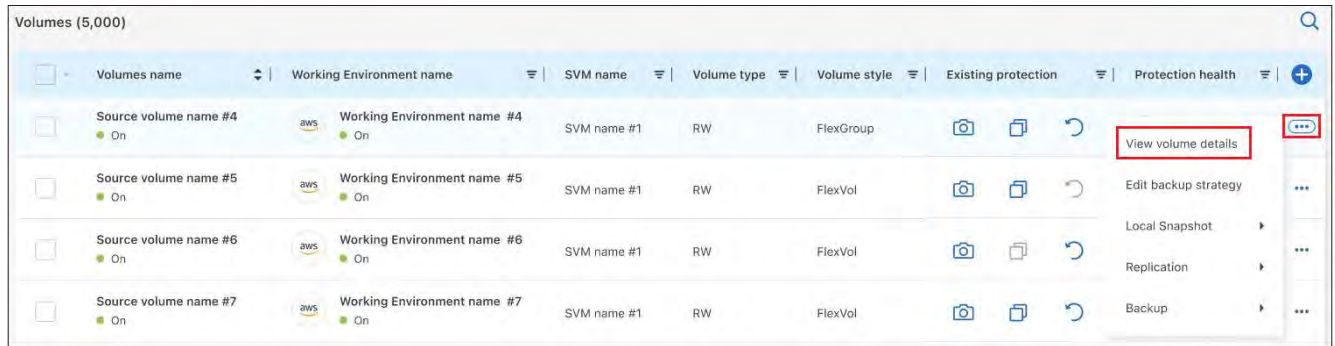
## Run a ransomware scan on a volume backup in object storage

BlueXP backup and recovery 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 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.

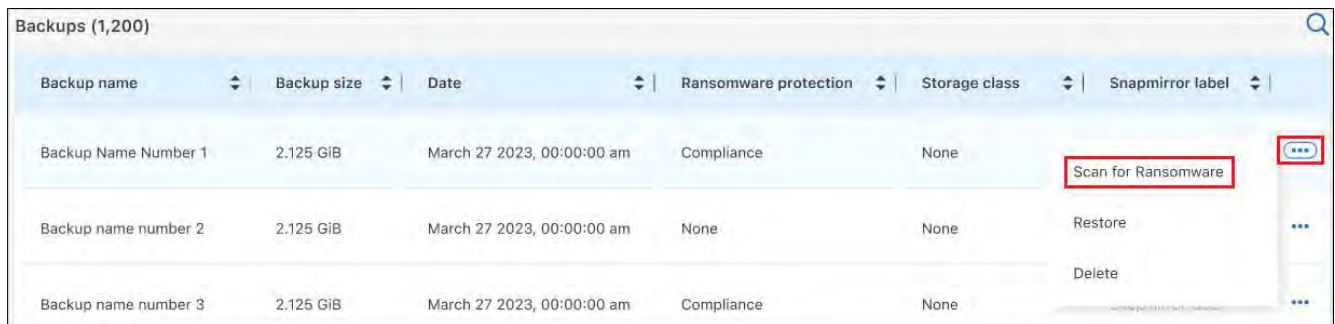
## Steps

1. From the **Volumes** tab, select **...** 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. Select **...** for the volume backup file you want to scan for ransomware and click **Scan for Ransomware**.



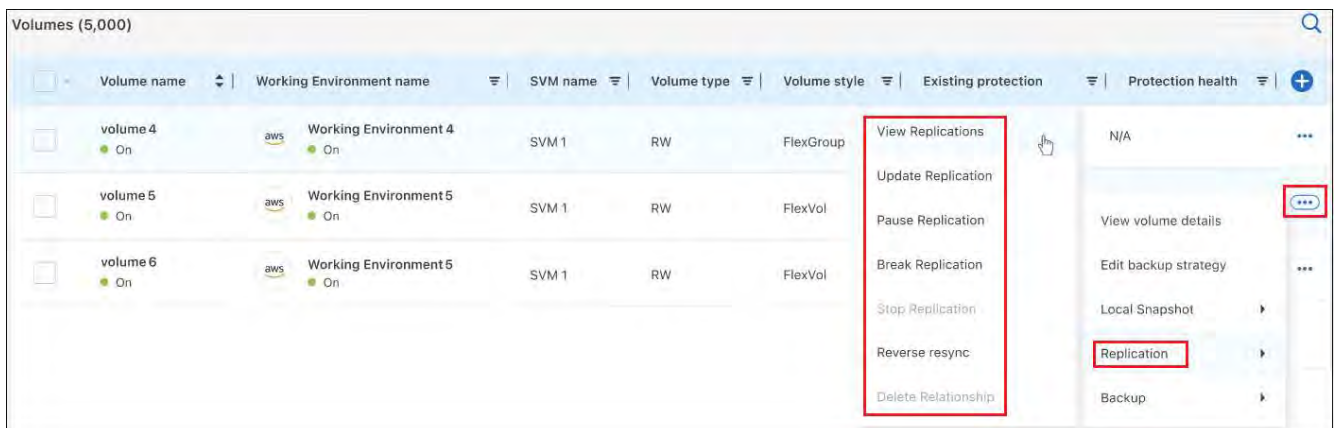
The Ransomware Protection column shows that the scan is In Progress.

## 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, select **...** 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	Breaks the relationship between the source and destination volumes, and activates the destination volume for data access - makes it read-write.  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.  <a href="#">Learn how to configure a destination volume for data access and reactivate a source volume in the ONTAP documentation</a>
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	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.  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.
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.

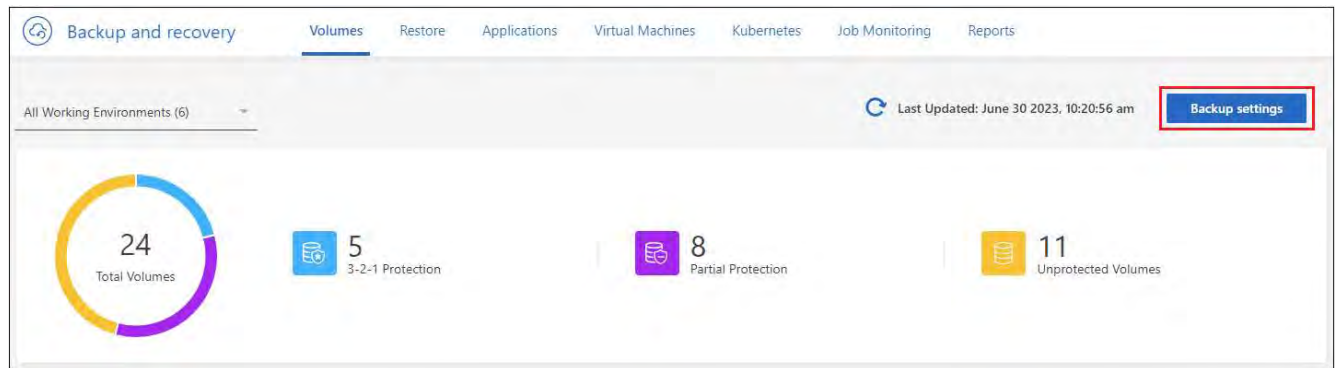




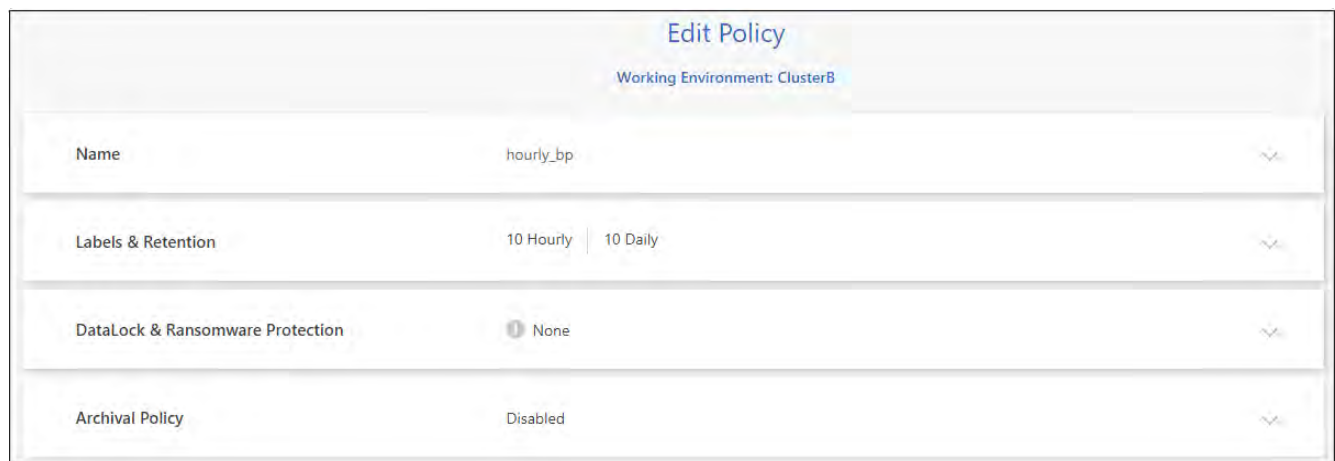
- If you enabled *DataLock and Ransomware Protection* in the initial policy when activating BlueXP backup and recovery for this cluster, any policies that you edit must be configured with the same DataLock setting (Governance or Compliance). And if you did not enable *DataLock and Ransomware Protection* 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, select **...** for the working environment where you want to change the policy settings, and select **Manage Policies**.
3. From the *Manage Policies* page, select **Edit** for the backup policy you want to change in that working environment.
4. From the *Edit Policy* page, select the down arrow to expand the *Labels & Retention* section to change the schedule and/or backup retention, and select **Save**.



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.)

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.

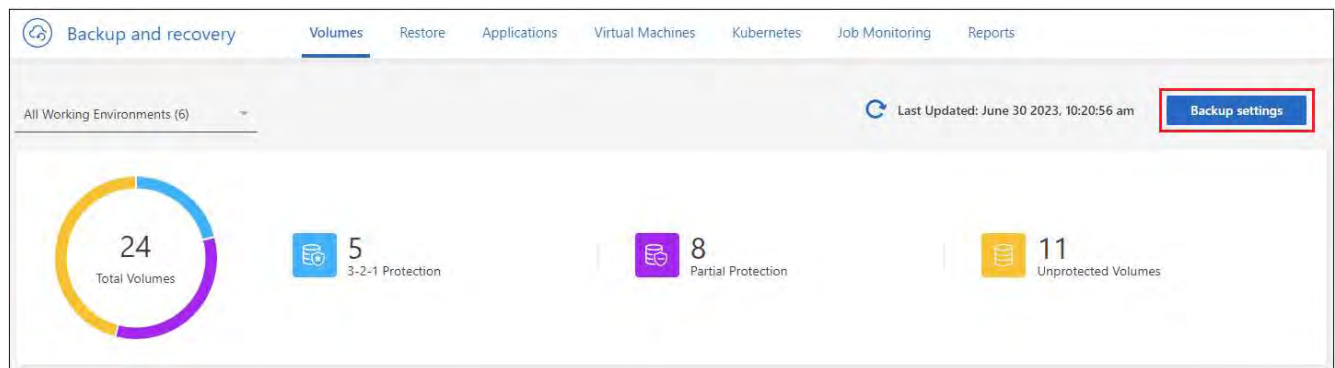
If you want to apply a new backup policy to certain volumes in a working environment, you first need to add the backup policy to the working environment. Then you can [apply the policy to volumes in that working environment](#).



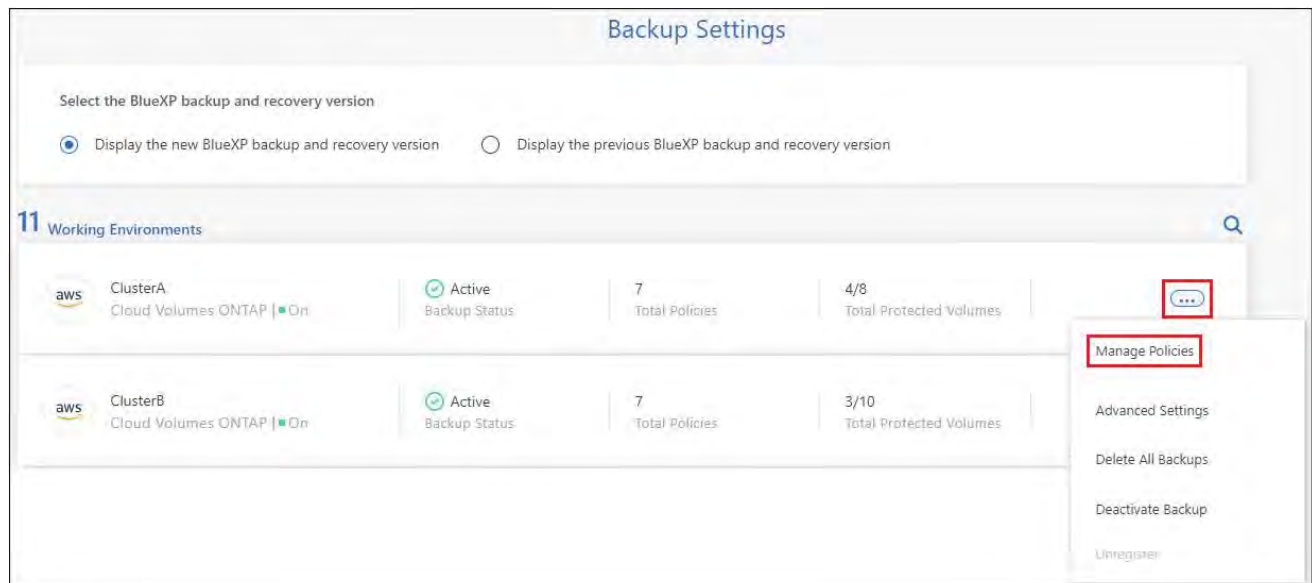
- If you enabled *DataLock and Ransomware Protection* in the initial policy when activating BlueXP backup and recovery for this cluster, any additional policies you create must be configured with the same DataLock setting (Governance or Compliance). And if you did not enable *DataLock and Ransomware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

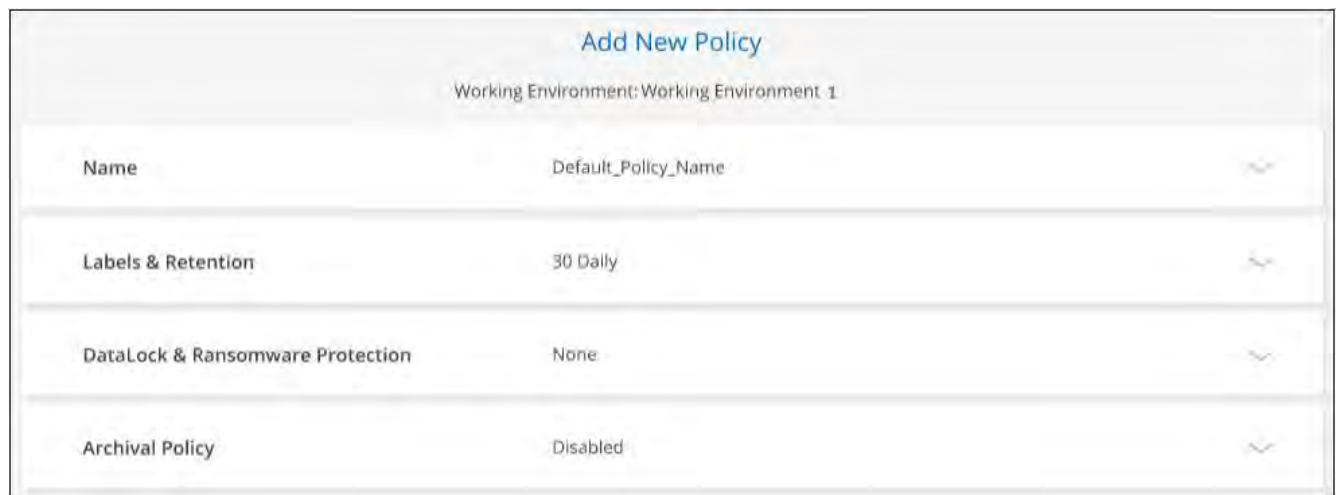
1. From the **Volumes** tab, select **Backup Settings**.



2. From the *Backup Settings* page, select **...** for the working environment where you want to add the new policy, and select **Manage Policies**.



3. From the *Manage Policies* page, select **Add New Policy**.
4. From the *Add New Policy* page, select down arrow to expand the *Labels & Retention* section to define the schedule and backup retention, and select **Save**.



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.)

## 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.

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 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.

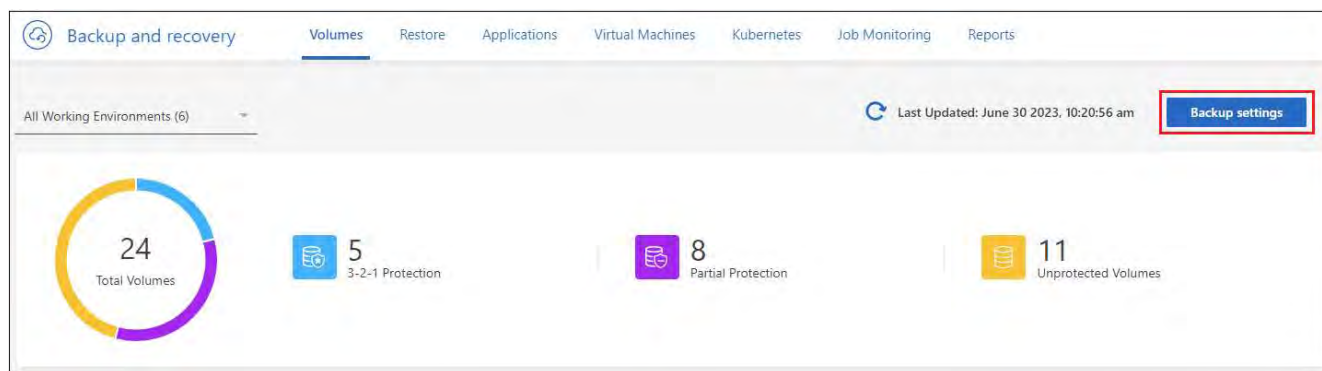
### Delete all backup files for a working environment

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. Select **...** for the working environment where you want to delete all backups and select **Delete All Backups**.

### Delete all backups

Deleting all backups for a working environment performs the following actions:

- Deletes all backup files from object storage.
- Disables future backups of those volumes.
- Disables the automatic backup feature for newly created volumes (if it was previously enabled).

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

Type the name of the of Working Environment in order to delete all backups:

Enter Working Environment Name

**Advanced settings** ▼

**Force Delete Backups** ⓘ

Disabled

Enabled

**ⓘ Note:** Please enable this option only if you are facing error deleting backups in regular way. **This option is Irreversible.** Backup and Recovery will forget/delete backup records for this working environment, even if it is unable to delete the backup from backup target. You will need to manually delete backups from the backup target.

**ⓘ Note:** You can activate backups for these volumes later from the Volumes page.

3. In the confirmation dialog box, enter the name of the working environment.
4. Select **Advanced settings**.
5. **Force delete backups:** Indicate whether or not you want to force the deletion of all backups.

In some extreme cases, you might want BlueXP backup and recovery not to have access to backups any longer. This might happen for example, if the service no longer has access to the backup bucket or backups are DataLock protected but you don't want them anymore. Previously, you could not delete these yourself and needed to call NetApp Support. With this release, you can use the option to force delete backups (at volume and work environment levels).



Use this option carefully and only in extreme cleanup needs. BlueXP backup and recovery will not have access to these backups any longer even if they are not deleted in the object storage. You will need to go to your cloud provider and manually delete the backups.

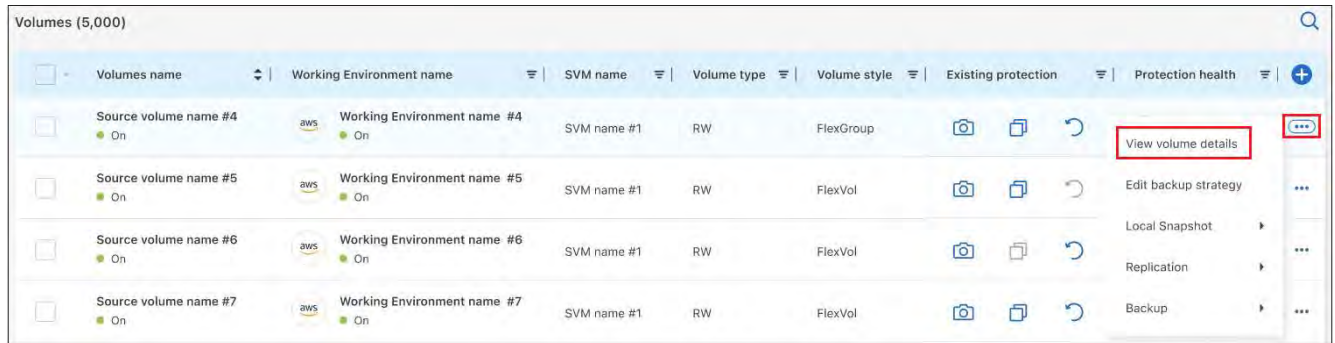
6. Select **Delete**.

**Delete all backup files for a volume**

Deleting all backups for a volume also disables future backups for that volume.

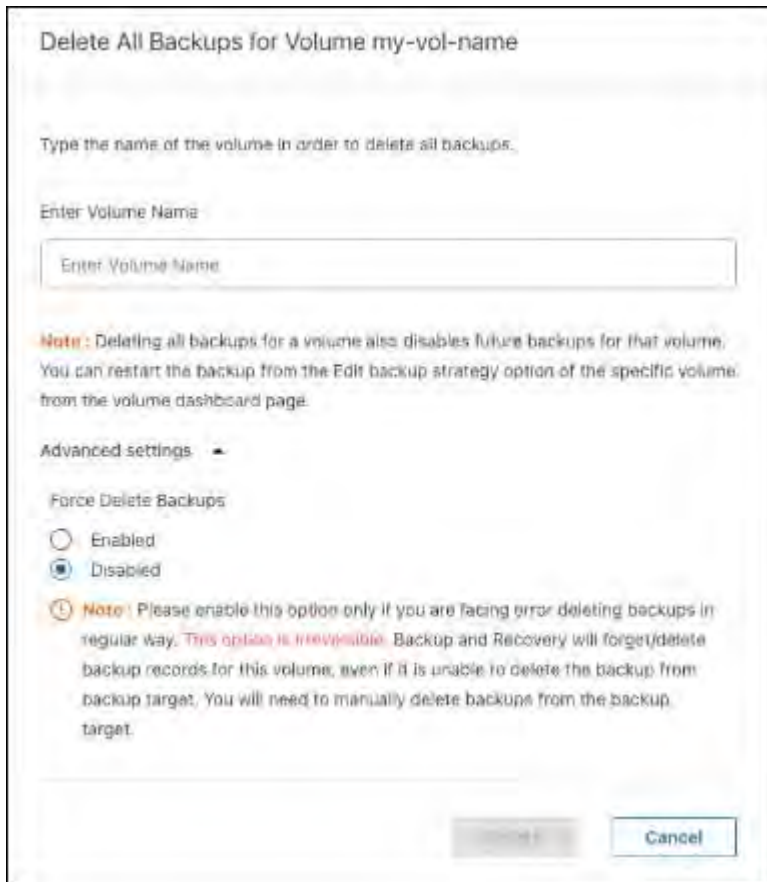
**Steps**

1. From the **Volumes** tab, click **...** for the source volume and select **Details & Backup List**.



The list of all backup files is displayed.

2. Select **Actions > Delete all Backups**.



3. Enter the volume name.
4. Select **Advanced settings**.

5. **Force delete backups:** Indicate whether or not you want to force the deletion of all backups.

In some extreme cases, you might want BlueXP backup and recovery not to have access to backups any longer. This might happen for example, if the service no longer has access to the backup bucket or backups are DataLock protected but you don't want them anymore. Previously, you could not delete these yourself and needed to call NetApp Support. With this release, you can use the option to force delete backups (at volume and work environment levels).



Use this option carefully and only in extreme cleanup needs. BlueXP backup and recovery will not have access to these backups any longer even if they are not deleted in the object storage. You will need to go to your cloud provider and manually delete the backups.

6. Select **Delete**.

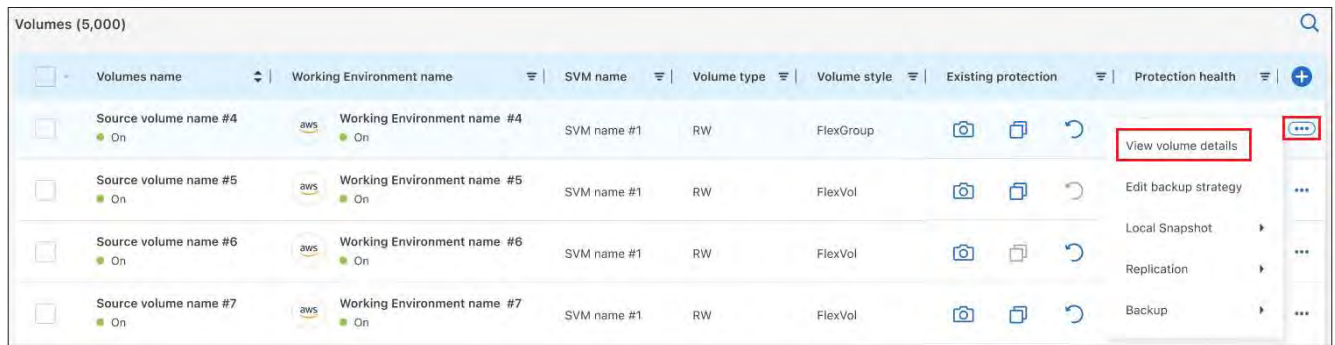
### Delete a single backup file for a volume

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, select **...** for the source volume and select **View volume details**.



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.

2. Select **Snapshot** or **Backup** to see the type of backup files that you want to delete.
3. Select **...** for the volume backup file you want to delete and select **Delete**.
4. In the confirmation dialog box, select **Delete**.

### Delete volume backup relationships

Deleting the backup relationship for a volume provides you with an archiving mechanism if you want to stop the creation of new backup files and delete the source volume, but retain all the existing backup files. This gives you the ability to restore the volume from the backup file in the future, if needed, while clearing space from your source storage system.

You don't necessarily need to delete the source volume. You can delete the backup relationship for a volume and retain the source volume. In this case you can "Activate" backup on the volume at a later time. The original baseline backup copy continues to be used in this case - a new baseline backup copy is not created and exported to the cloud. Note that if you do reactivate a backup relationship, the volume is assigned the default

backup policy.

This feature is available only if your system is running ONTAP 9.12.1 or greater.

You can't delete the source volume from the BlueXP backup and recovery user interface. However, you can open the Volume Details page on the Canvas, and [delete the volume from there](#).



You can't delete individual volume backup files once the relationship has been deleted. You can, however, you can delete all backups for the volume.

### Steps

1. From the **Volumes** tab, select **...** for the source volume and select **Backup > Delete relationship**.

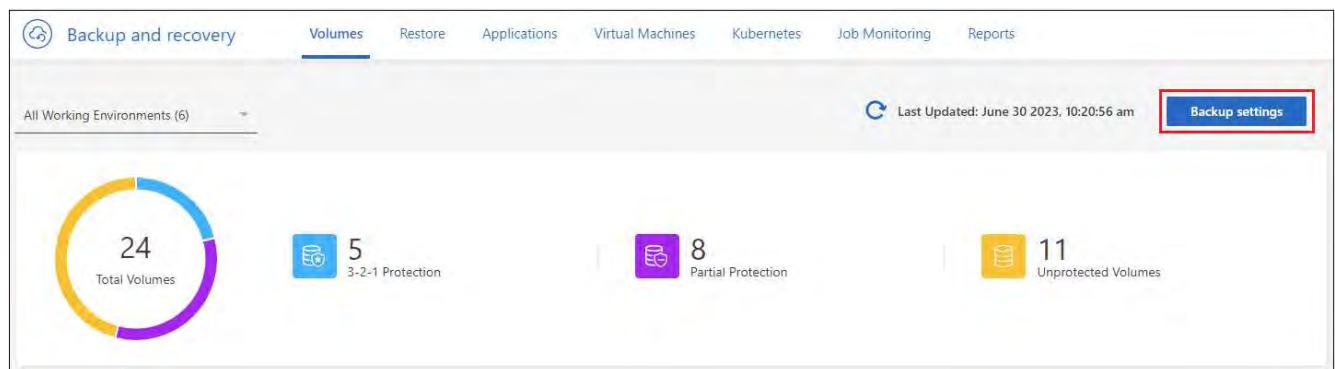
### Deactivate BlueXP backup and recovery for a working environment

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, select **...** for the working environment where you want to disable backups and select **Deactivate Backup**.
3. In the confirmation dialog box, select **Deactivate**.



An **Activate Backup** button appears for that working environment while backup is disabled. You can select 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.

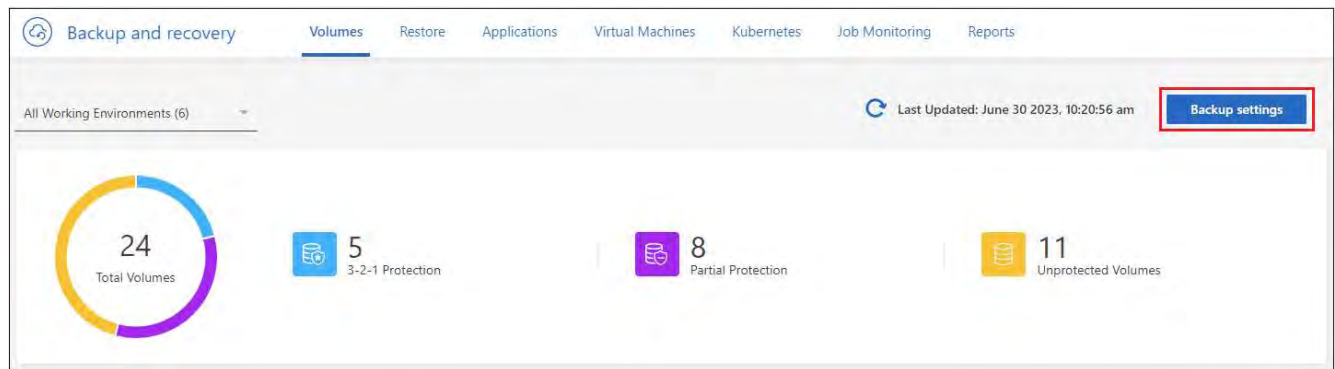
Before you can unregister BlueXP backup and recovery, you must perform the following steps, in this order:

- 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*, select **...** for the working environment where you want to unregister the backup service and select **Unregister**.
3. In the confirmation dialog box, select **Unregister**.

## Restore ONTAP data from backup files with BlueXP backup and recovery

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. With BlueXP backup and recovery, restore an entire ONTAP volume from a backup file, or if you only need to restore a few files, restore a folder or individual files.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

- 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




If the restore operation does not complete, do not try the restore process again until the Job Monitor shows that the restore operation has failed. If you try the restore process again before the Job Monitor shows that the restore operation has failed, the restore operation will fail again. When you see the Job Monitor status as "Failed," you can try the restore process again.



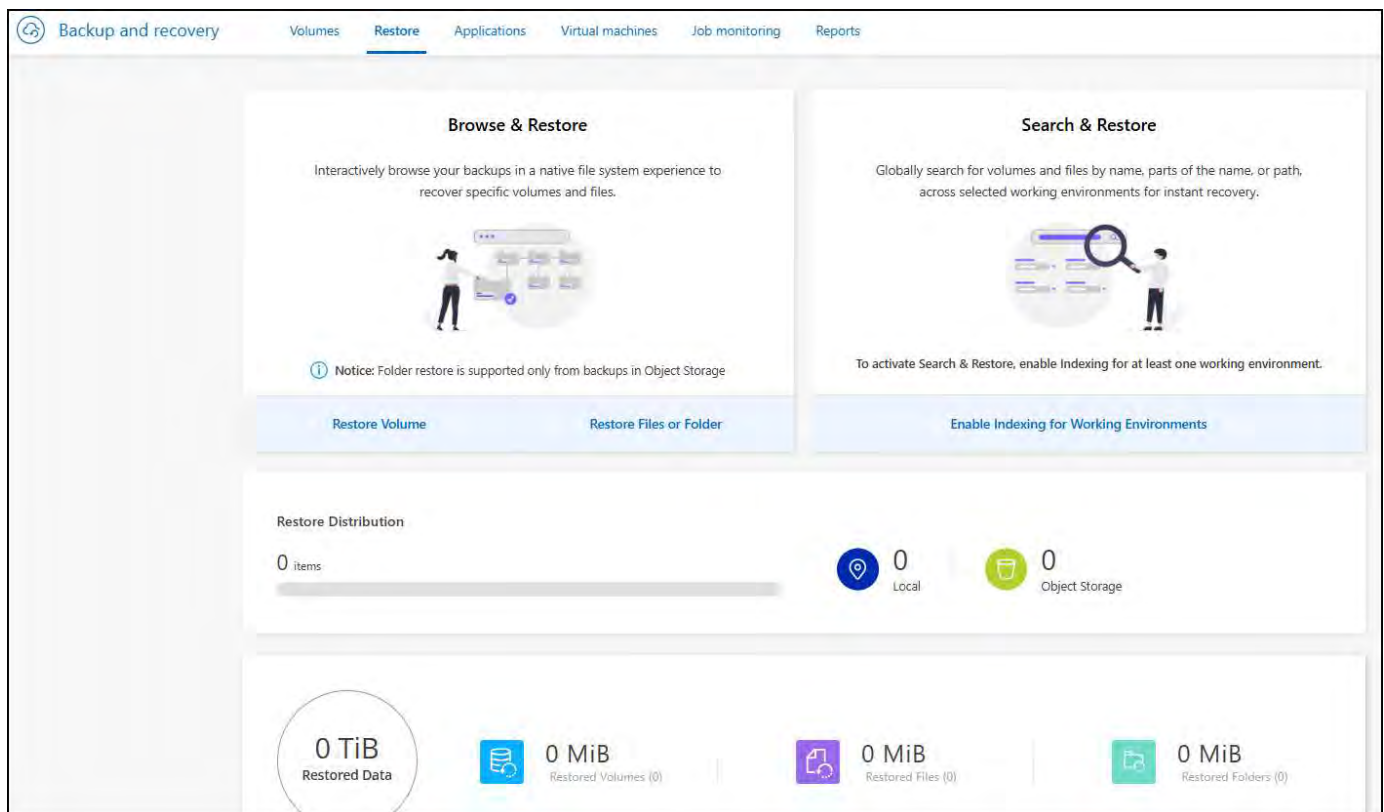
For limitations related to restoring ONTAP data, see [Backup and restore limitations for ONTAP volumes](#).

## The Restore Dashboard

You use the Restore Dashboard to perform volume, folder, and file restore operations. You access the Restore Dashboard by clicking **Backup and recovery** from the BlueXP menu, and then clicking the **Restore** tab. You can also click  > **View Restore Dashboard** from the Backup and recovery service from the Services panel.



BlueXP backup and recovery must already be activated for at least one working environment and initial backup files must exist.



The screenshot shows the 'Restore' dashboard with the following components:

- Navigation:** Backup and recovery, Volumes, **Restore**, Applications, Virtual machines, Job monitoring, Reports.
- Browse & Restore:** Interactively browse your backups in a native file system experience to recover specific volumes and files. Includes a 'Notice: Folder restore is supported only from backups in Object Storage'. Buttons: Restore Volume, Restore Files or Folder.
- Search & Restore:** Globally search for volumes and files by name, parts of the name, or path, across selected working environments for instant recovery. Includes a button: Enable Indexing for Working Environments.
- Restore Distribution:** 0 items. Metrics: 0 Local, 0 Object Storage.
- Summary Row:** 0 TiB Restored Data, 0 MiB Restored Volumes (0), 0 MiB Restored Files (0), 0 MiB Restored Folders (0).

As you can see, the Restore Dashboard provides two 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 two 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 <b>all backup files</b> by partial or full volume name, partial or full folder/file name, size range, and additional search filters.
Does not handle file recovery if the file has been deleted or renamed, and the user doesn't know the original file name	Handles newly created/deleted/renamed directories and newly created/deleted/renamed files
No additional cloud provider resources required	When you restore from the cloud, additional bucket and public cloud provider resources required per account.
No additional cloud provider costs required	When you restore from the cloud, additional costs are required when scanning your backups and volumes for search results.
Quick restore is supported.	Quick restore is not supported.

This table provides a list of valid restore operations based on the location where your backup files reside.

Backup Type	Browse & Restore			Search & Restore		
	Restore volume	Restore files	Restore folder	Restore volume	Restore files	Restore folder
<b>Snapshot copy</b>	Yes	No	No	Yes	Yes	Yes
<b>Replicated volume</b>	Yes	No	No	Yes	Yes	Yes
<b>Backup file</b>	Yes	Yes	Yes	Yes	Yes	Yes

Before you can use either restore method, make sure you have configured your environment for the unique resource requirements. Those requirements are described in the sections below.

See the requirements and restore steps for the type of restore operation you want to use:

- [Restore volumes using Browse & Restore](#)
- [Restore folders and files using Browse & Restore](#)
- [Restore volumes, folders, and files using Search & Restore](#)



## Restore ONTAP data using Browse & Restore

Before you start restoring a volume, folder, or file, you should know the name of the volume from which you want to restore, the name of the working environment and SVM where the volume resides, and the approximate date of the backup file that you want to restore from. You can restore ONTAP data from a Snapshot copy, a replicated volume, or from backups stored in object storage.

**Note:** If the backup file containing the data that you want to restore resides in archival cloud storage (starting with ONTAP 9.10.1), the restore operation will take a longer amount of time and will incur a cost. Additionally, the destination cluster must also be running ONTAP 9.10.1 or greater for volume restore, 9.11.1 for file restore, 9.12.1 for Google Archive and StorageGRID, and 9.13.1 for folder restore.

[Learn more about restoring from AWS archival storage.](#)

[Learn more about restoring from Azure archival storage.](#)

[Learn more about restoring from Google archival storage.](#)



The High priority isn't supported when restoring data from Azure archival storage to StorageGRID systems.

### Browse & Restore supported working environments and object storage providers

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

**Note:** You can restore a volume from any type of backup file, but you can restore a folder or individual files only from a backup file in object storage at this time.

From Object Store (Backup)	From Primary (Snapshot)	From Secondary System (Replication)	To Destination Working Environment
Amazon S3	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Azure Blob
Cloud Volumes ONTAP in Azure On-premises ONTAP system	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Google Cloud Storage	Cloud Volumes ONTAP in Google On-premises ONTAP system
Cloud Volumes ONTAP in Google On-premises ONTAP system	NetApp StorageGRID	On-premises ONTAP system	On-premises ONTAP system Cloud Volumes ONTAP
To on-premises ONTAP system	ONTAP S3	On-premises ONTAP system	On-premises ONTAP system Cloud Volumes ONTAP

For Browse & Restore, the Connector can be installed in the following locations:

- For Amazon S3, the Connector can be deployed in AWS or in your premises
- For Azure Blob, the Connector can be deployed in Azure or in your premises

- For Google Cloud Storage, the Connector must be deployed in your Google Cloud Platform VPC
- For StorageGRID, the Connector must be deployed in your premises; with or without internet access
- For ONTAP S3, the Connector can be deployed in your premises (with or without internet access) or in a cloud provider environment

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



If the ONTAP version on your system is less than 9.13.1, then you can't restore folders or files if the backup file has been configured with DataLock & Ransomware. In this case, you can restore the entire volume from the backup file and then access the files you need.

### Restore volumes using Browse & Restore

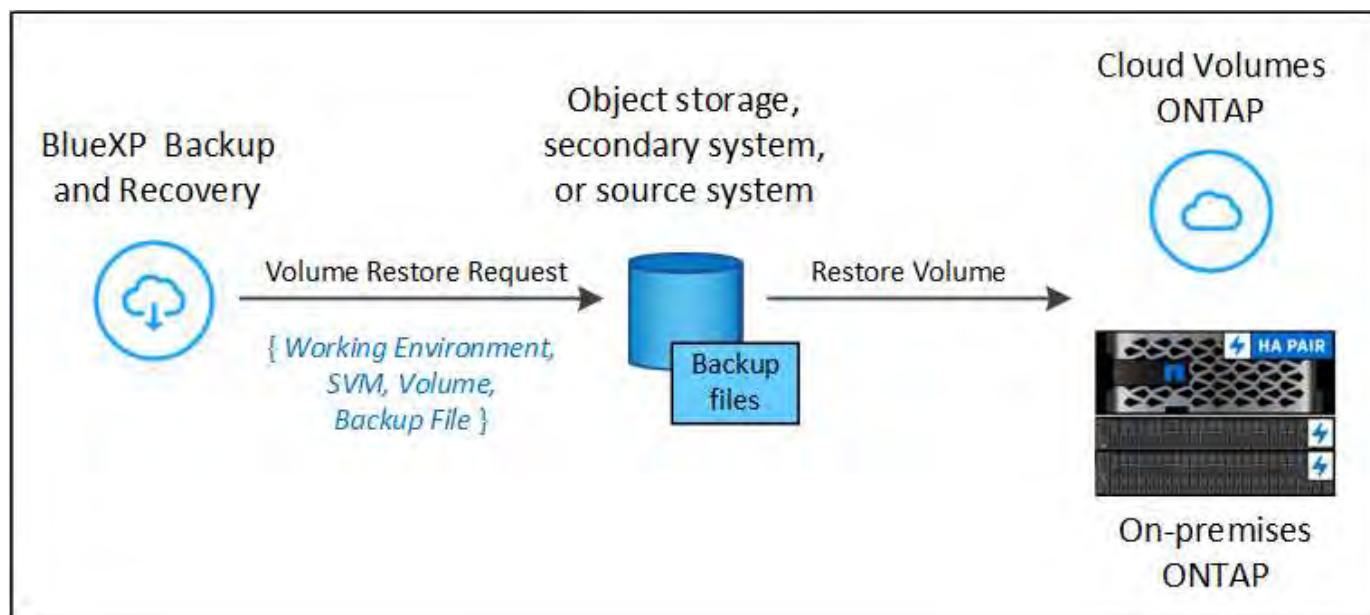
When you restore a volume from a backup file, BlueXP backup and recovery creates a *new* volume using the data from the backup. When using a backup from object storage, you can restore the data to a volume in the original working environment, to a different working environment that's located in the same cloud account as the source working environment, or to an on-premises ONTAP system.

When restoring a cloud backup to a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater or to an on-premises ONTAP system running ONTAP 9.14.1, you'll have the option to perform a *quick restore* operation. The quick restore is ideal for disaster recovery situations where you need to provide access to a volume as soon as possible. A quick restore restores the metadata from the backup file to a volume instead of restoring the entire backup file. Quick restore is not recommended for performance or latency-sensitive applications, and it is not supported with backups in archived storage.



Quick restore is supported for FlexGroup volumes only if the source system from which the cloud backup was created was running ONTAP 9.12.1 or greater. And it is supported for SnapLock volumes only if the source system was running ONTAP 9.11.0 or greater.

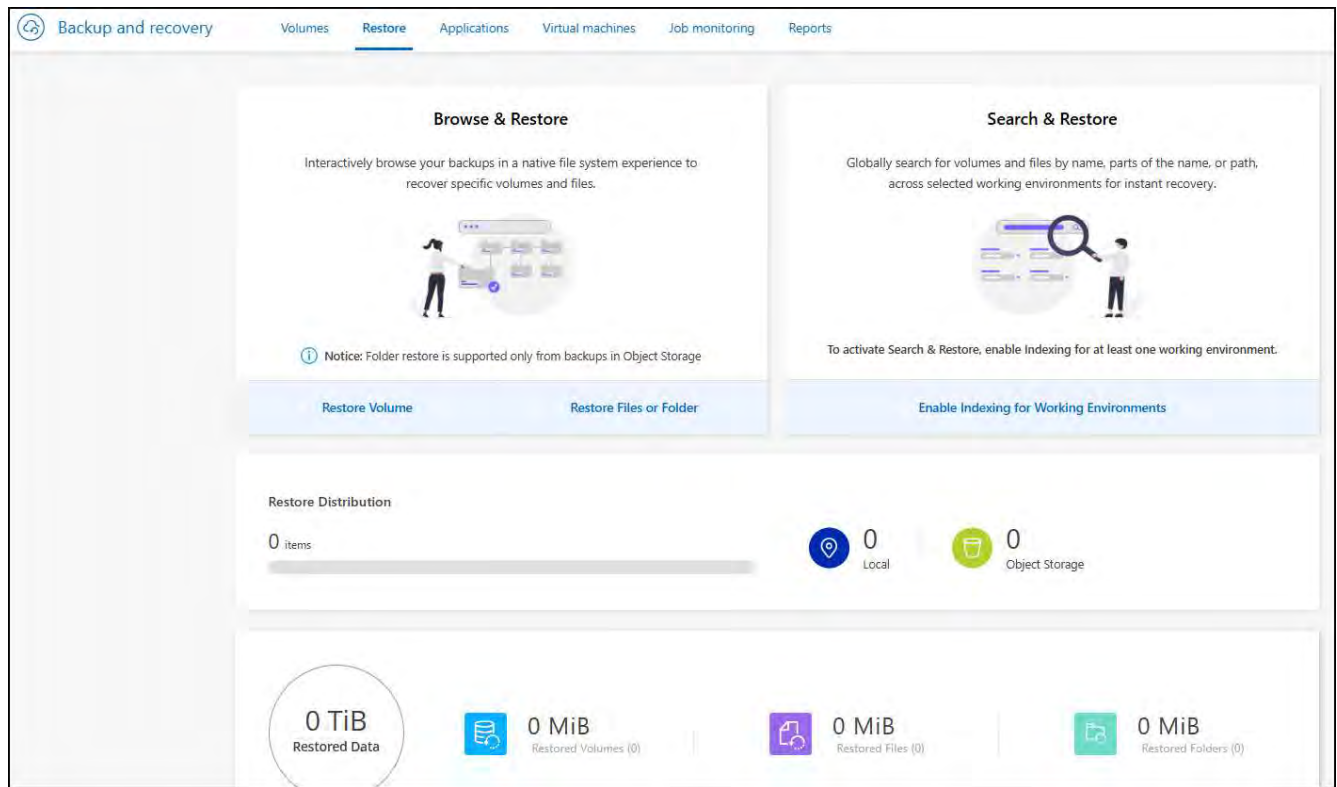
When restoring from a replicated volume, you can restore the volume to the original working environment or to a Cloud Volumes ONTAP or on-premises ONTAP system.



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

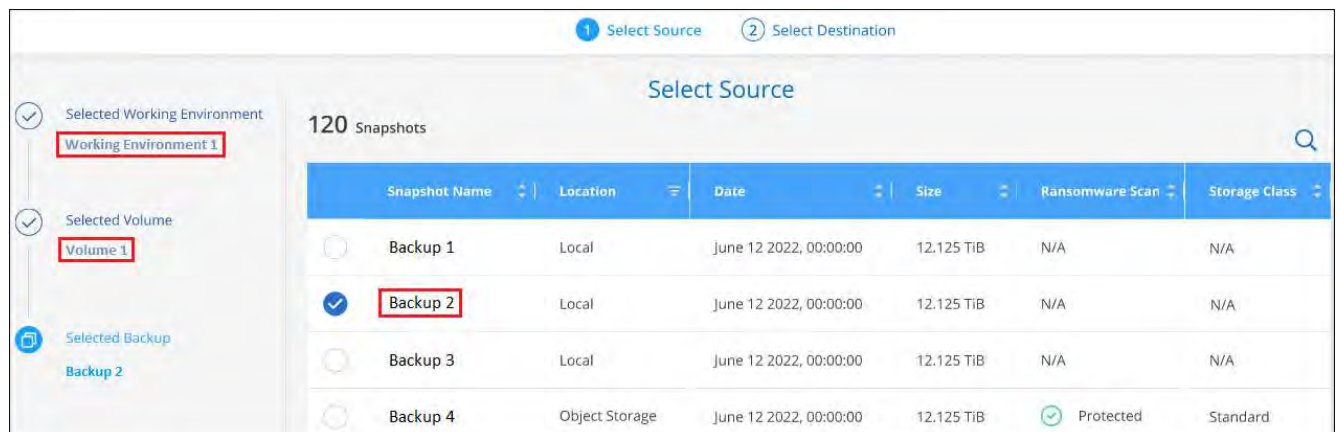
## Steps

1. From the BlueXP menu, select **Protection > Backup and recovery**.
2. Select the **Restore** tab and the Restore Dashboard is displayed.
3. From the *Browse & Restore* section, select **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.

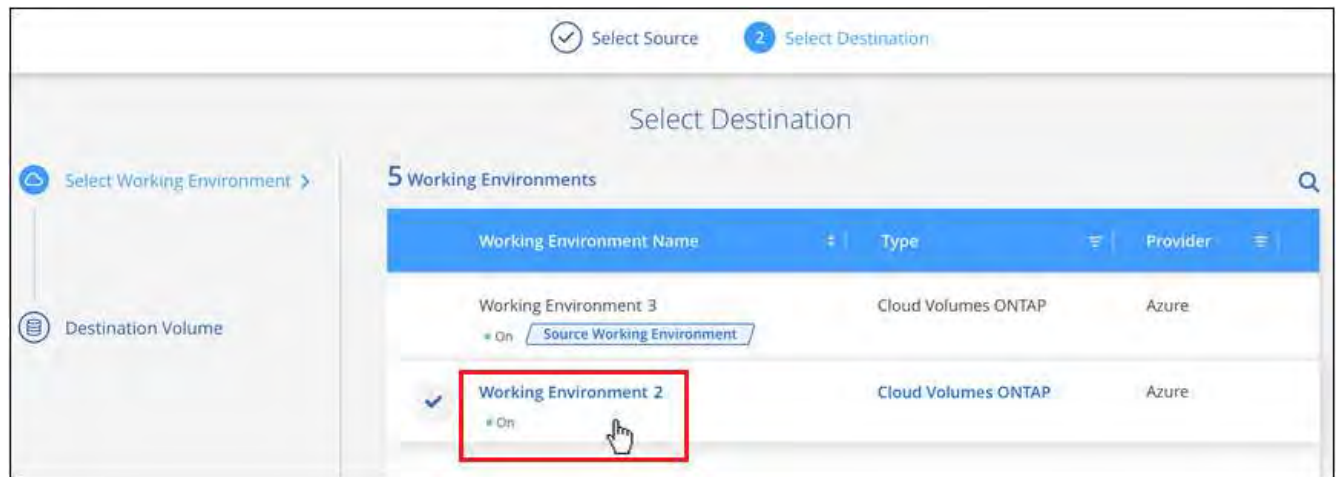


5. Select **Next**.

Note that if you select a backup file in object storage, and ransomware protection is active for that backup

(if you enabled DataLock and Ransomware Protection in the backup policy), then you are prompted to run an additional ransomware scan on the backup file before restoring the data. We recommend that you scan the backup file for ransomware. (You'll incur extra egress costs from your cloud provider to access the contents of the backup file.)

6. In the *Select Destination* page, select the **Working Environment** where you want to restore the volume.



7. When restoring a backup file from object storage, if you select an on-premises ONTAP system and you haven't already configured the cluster connection to the object storage, you are prompted for additional information:
  - When restoring from Amazon S3, select the IPspace in the ONTAP cluster where the destination volume will reside, enter the access key and secret key for the user you created to give the ONTAP cluster access to the S3 bucket, and optionally choose a private VPC endpoint for secure data transfer.
  - When restoring from Azure Blob, select the IPspace in the ONTAP cluster where the destination volume will reside, select the Azure Subscription to access the object storage, and optionally choose a private endpoint for secure data transfer by selecting the VNet and Subnet.
  - When restoring from Google Cloud Storage, select the Google Cloud Project and the Access Key and Secret Key to access the object storage, the region where the backups are stored, and the IPspace in the ONTAP cluster where the destination volume will reside.
  - When restoring from StorageGRID, enter the FQDN of the StorageGRID server and the port that ONTAP should use for HTTPS communication with StorageGRID, select the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume will reside.
  - When restoring from ONTAP S3, enter the FQDN of the ONTAP S3 server and the port that ONTAP should use for HTTPS communication with ONTAP S3, select the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume will reside.
8. Enter the name you want to use for the restored volume, and select the Storage VM and Aggregate where the volume will reside. When restoring a FlexGroup volume you'll need to select multiple aggregates. By default, **<source\_volume\_name>\_restore** is used as the volume name.

Select Destination

Selected Working Environment  
Working Environment Name 2

Destination Volume >  
General\_restore

*i* A new volume will be created in the working environment based on the backup you selected

Volume Name  
General\_restore

Storage VM  
svm1

Aggregate  
aggr2

Restore Priority  
Low

Volume Information
Volume Size: 50.00 GB
Backup Policy: CloudBackupService
Protocol: NFS
<b>Disk Type: RW</b>

When restoring a backup from object storage to a Cloud Volumes ONTAP system using ONTAP 9.13.0 or greater or to an on-premises ONTAP system running ONTAP 9.14.1, you'll have the option to perform a *quick restore* operation.

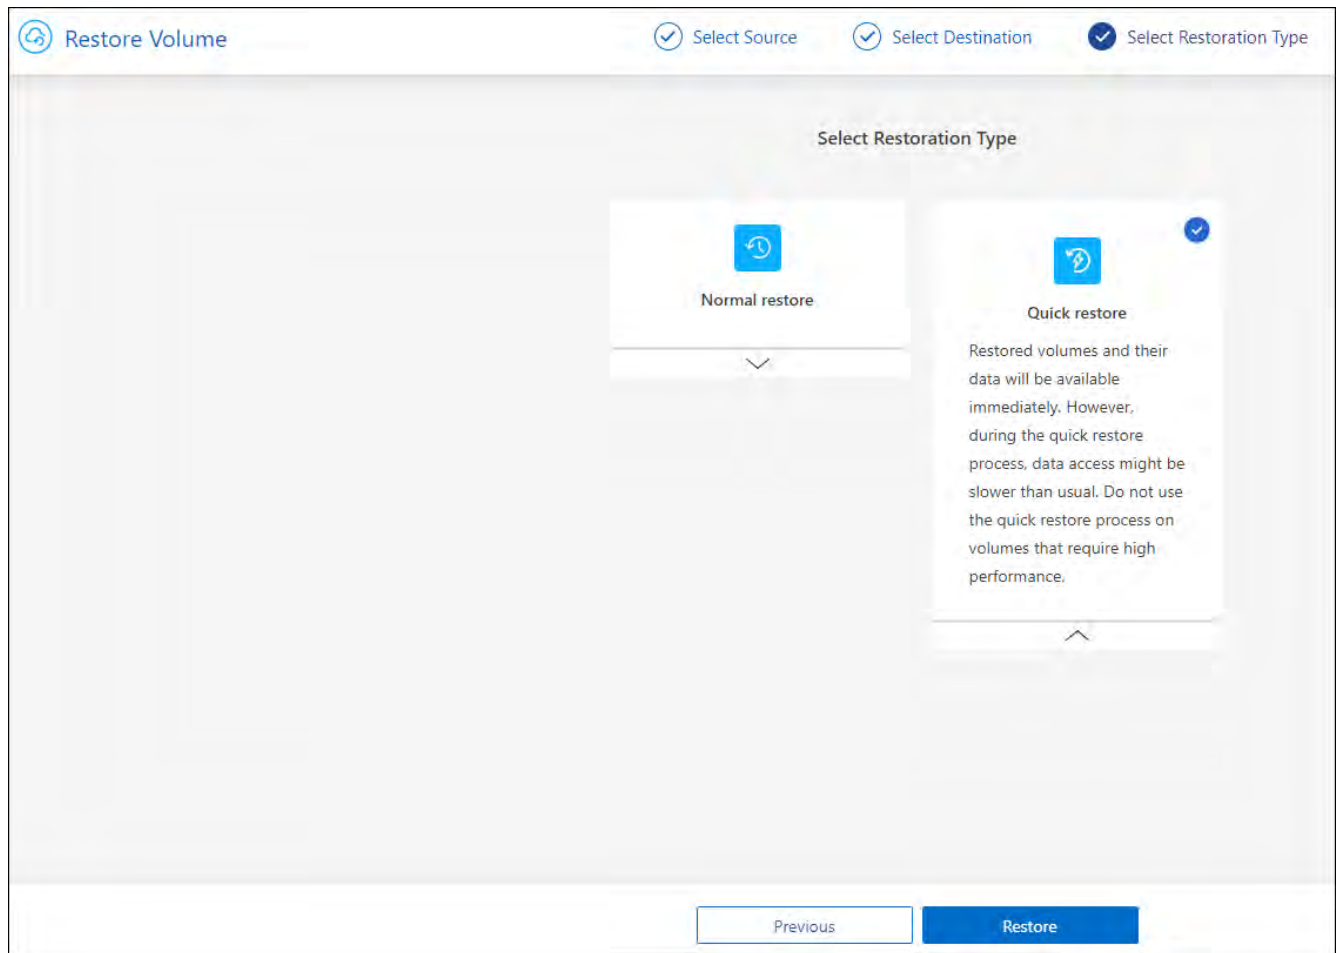
And if you are restoring the volume from a backup file that resides in an archival storage tier (available starting with ONTAP 9.10.1), then you can select the Restore Priority.

[Learn more about restoring from AWS archival storage.](#)

[Learn more about restoring from Azure archival storage.](#)

[Learn more about restoring from Google archival storage.](#) Backup files in the Google Archive storage tier are restored almost immediately, and require no Restore Priority.

9. Select **Next** to choose whether you want to do a Normal restore or a Quick Restore process:



- **Normal restore:** Use normal restore on volumes that require high performance. Volumes will not be available until the restore process is complete.
- **Quick restore:** Restored volumes and data will be available immediately. Do not use this on volumes that require high performance because during the quick restore process, access to the data might be slower than usual.

10. Select **Restore** and you return 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 select the **Job Monitoring** tab to see the restore progress.

### Restore folders and files using Browse & Restore

If you need to restore only a few files from an ONTAP volume backup, you can choose to restore a folder or individual files instead of restoring the entire volume. You can restore folders and files to an existing volume in the original working environment, or to a different working environment that's using the same cloud account. You can also restore folders and files to a volume on an on-premises ONTAP system.





You can restore a folder or individual files only from a backup file in object storage at this time. Restoring files and folders is not currently supported from a local snapshot copy or from a backup file that resides in a secondary working environment (a replicated volume).

If you select multiple files, all the files are restored to the same destination volume that you choose. So if you want to restore files to different volumes, you'll need to run the restore process multiple times.

When using ONTAP 9.13.0 or greater, you can restore a folder along with all files and sub-folders within it. When using a version of ONTAP before 9.13.0, only files from that folder are restored - no sub-folders, or files in sub-folders, are restored.



- If the backup file has been configured with DataLock & Ransomware protection, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the entire volume from the backup file and then access the folder and files you need.
- If the backup file resides in archival storage, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the folder from a newer backup file that has not been archived, or you can restore the entire volume from the archived backup and then access the folder and files you need.
- With ONTAP 9.15.1, you can restore FlexGroup folders using the "Browse and restore" option. This feature is in a Technology Preview mode.

You can test it using a special flag described in the [BlueXP backup and recovery July 2024 Release blog](#).

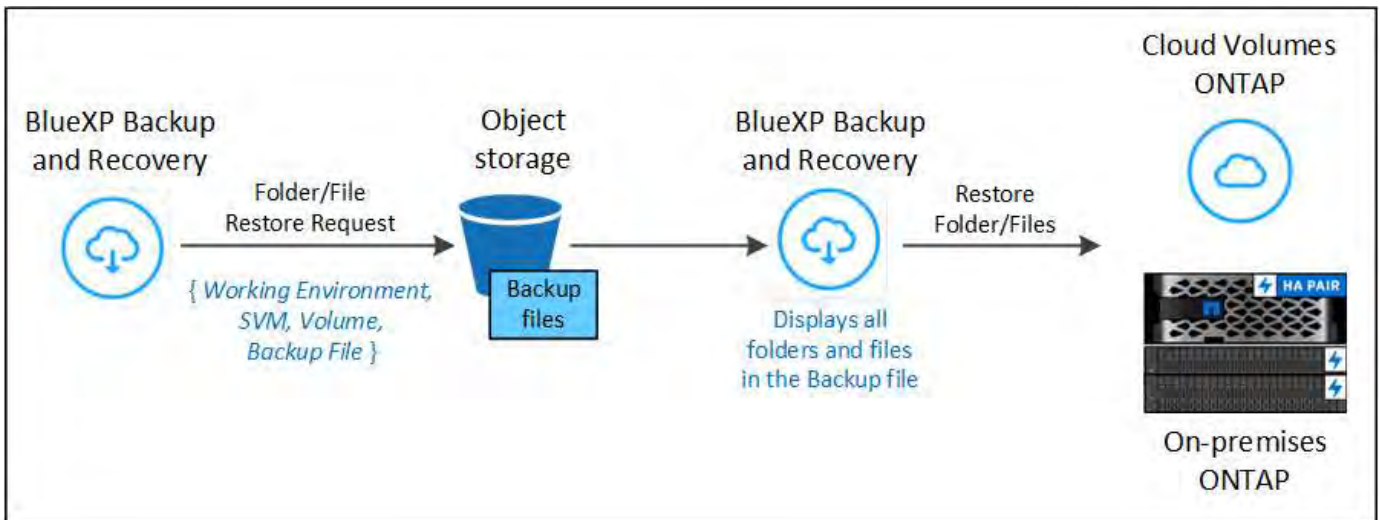
## Prerequisites

- The ONTAP version must be 9.6 or greater to perform *file* restore operations.
- The ONTAP version must be 9.11.1 or greater to perform *folder* restore operations. ONTAP version 9.13.1 is required if the data is in archival storage, or if the backup file is using DataLock and Ransomware protection.
- The ONTAP version must be 9.15.1 p2 or greater to restore FlexGroup directories using the Browse and restore option.

## Folder and file restore process

The process goes like this:

1. When you want to restore a folder, or one or more files, from a volume backup, click the **Restore** tab, and click **Restore Files or Folder** under *Browse & Restore*.
2. Select the source working environment, volume, and backup file in which the folder or file(s) reside.
3. BlueXP backup and recovery displays the folders and files that exist within the selected backup file.
4. Select the folder or file(s) that you want to restore from that backup.
5. Select the destination location where you want the folder or file(s) to be restored (the working environment, volume, and folder), and click **Restore**.
6. The file(s) are restored.



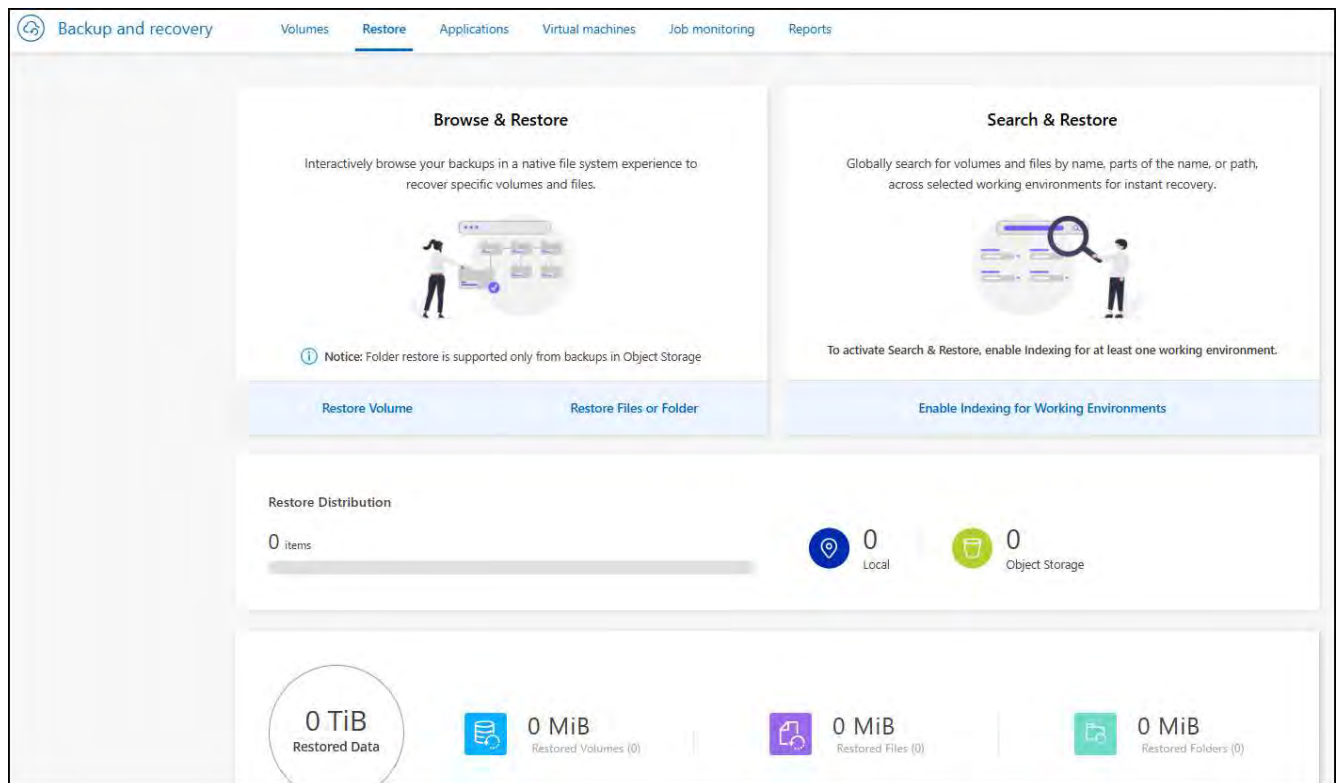
As you can see, you need to know the working environment name, volume name, backup file date, and folder/file name to perform a folder or file restore.

### Restore folders and files

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

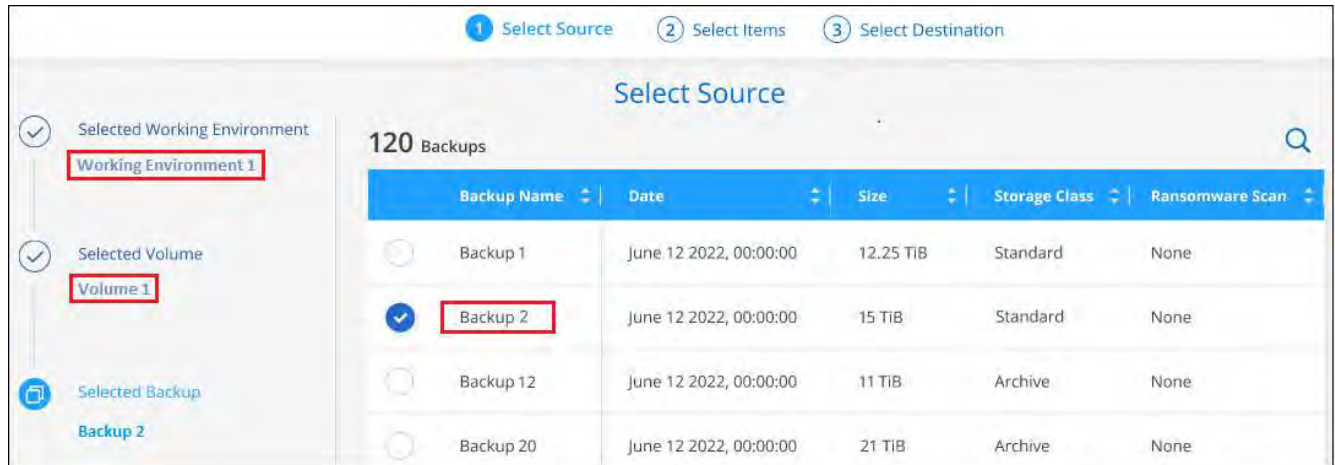
#### Steps

1. From the BlueXP menu, select **Protection > Backup and recovery**.
2. Select the **Restore** tab and the Restore Dashboard is displayed.
3. From the *Browse & Restore* section, select **Restore Files or Folder**.





- 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.



- Select **Next** and the list of folders and files from the volume backup are displayed.

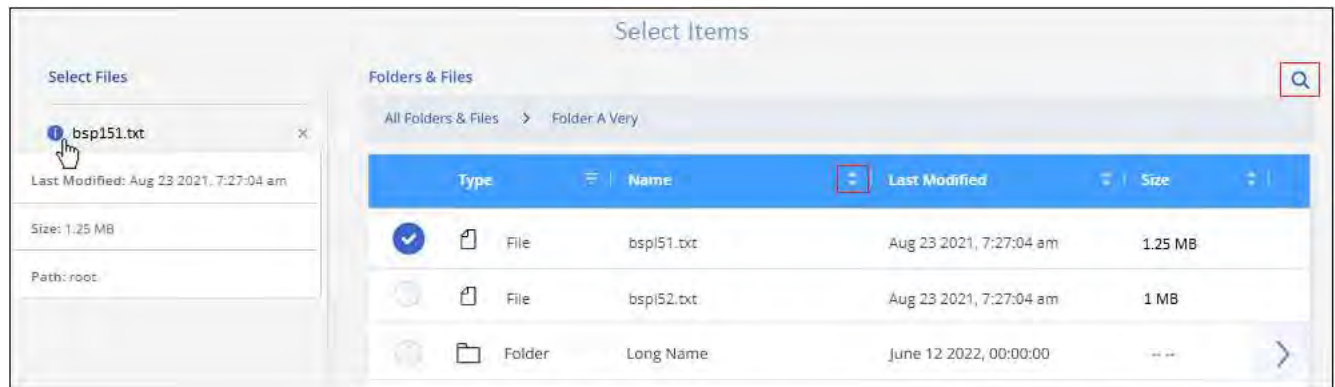
If you are restoring folders or files from a backup file that resides in an archival storage tier, then you can select the Restore Priority.

[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.)

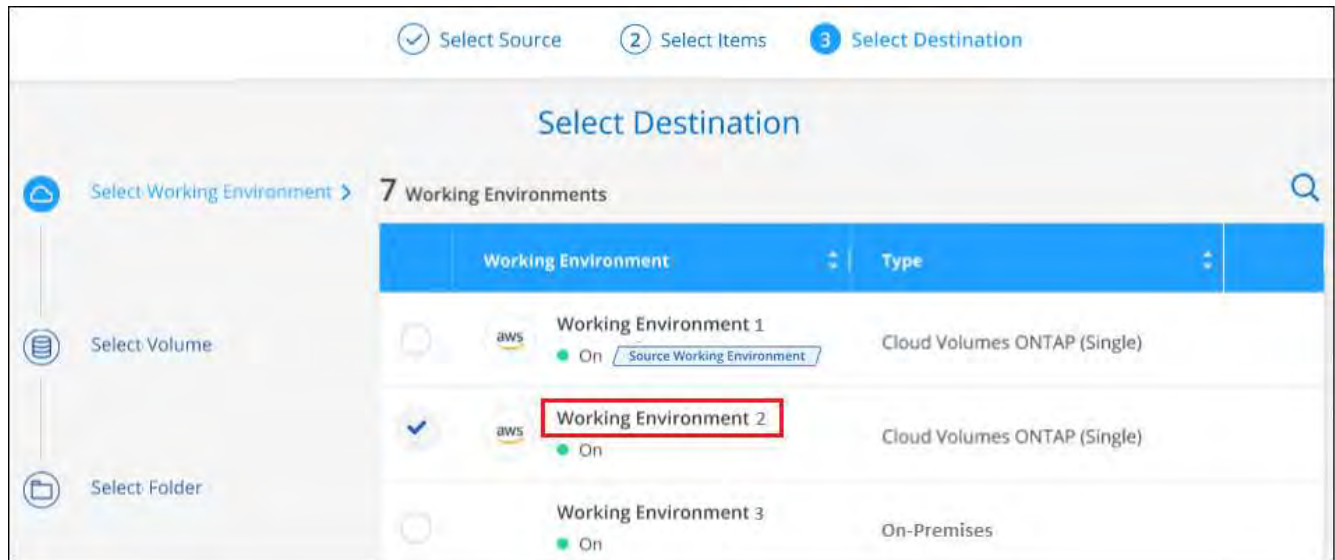


- In the *Select Items* page, select the folder or file(s) that you want to restore and select **Continue**. To assist you in finding the item:

- You can select the folder or file name if you see it.
- You can select 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 Down arrow 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 selecting 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. Select **Restore** and you are returned to the Restore Dashboard so you can review the progress of the restore operation. You can also click the **Job Monitoring** tab to see the restore progress.

## Restore ONTAP data using Search & Restore

You can restore a volume, folder, or files from an ONTAP backup file using Search & Restore. Search & Restore enables you to search for a specific volume, folder, or file from all backups, and then perform a restore. You don't need to know the exact working environment name, volume name, or file name - the search looks through all volume backup files.

The search operation looks across all local snapshot copies that exist for your ONTAP volumes, all replicated volumes on secondary storage systems, and all backup files that exist in object storage. Since restoring data from a local Snapshot copy or replicated volume can be faster and less costly than restoring from a backup file in object storage, you may want to restore data from these other locations.

When you restore a *full volume* from a backup file, BlueXP backup and recovery creates a *new* volume using the data from the backup. You can restore the data as a volume in the original working environment, to a different working environment that's located in the same cloud account as the source working environment, or to an on-premises ONTAP system.

You can restore *folders or files* to the original volume location, to a different volume in the same working environment, to a different working environment that's using the same cloud account, or to a volume on an on-premises ONTAP system.

When using ONTAP 9.13.0 or greater, you can restore a folder along with all files and sub-folders within it. When using a version of ONTAP before 9.13.0, only files from that folder are restored - no sub-folders, or files in sub-folders, are restored.

If the backup file for the volume that you want to restore resides in archival storage (available starting with ONTAP 9.10.1), the restore operation will take a longer amount of time and will incur additional cost. Note that the destination cluster must also be running ONTAP 9.10.1 or greater for volume restore, 9.11.1 for file restore, 9.12.1 for Google Archive and StorageGRID, and 9.13.1 for folder restore.

[Learn more about restoring from AWS archival storage.](#)

[Learn more about restoring from Azure archival storage.](#)

[Learn more about restoring from Google archival storage.](#)



- If the backup file in object storage has been configured with DataLock & Ransomware protection, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the entire volume from the backup file and then access the folder and files you need.
- If the backup file in object storage resides in archival storage, then folder-level restore is supported only if the ONTAP version is 9.13.1 or greater. If you are using an earlier version of ONTAP, you can restore the folder from a newer backup file that has not been archived, or you can restore the entire volume from the archived backup and then access the folder and files you need.
- The "High" restore priority is not supported when restoring data from Azure archival storage to StorageGRID systems.
- Restoring folders is not currently supported from volumes in ONTAP S3 object storage.

Before you start, you should have some idea of the name or location of the volume or file you want to restore.

### 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 & Restore 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](#).

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.

For information about the Indexed Catalog v2, see the section below about how to enable the Indexed Catalog.

- 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

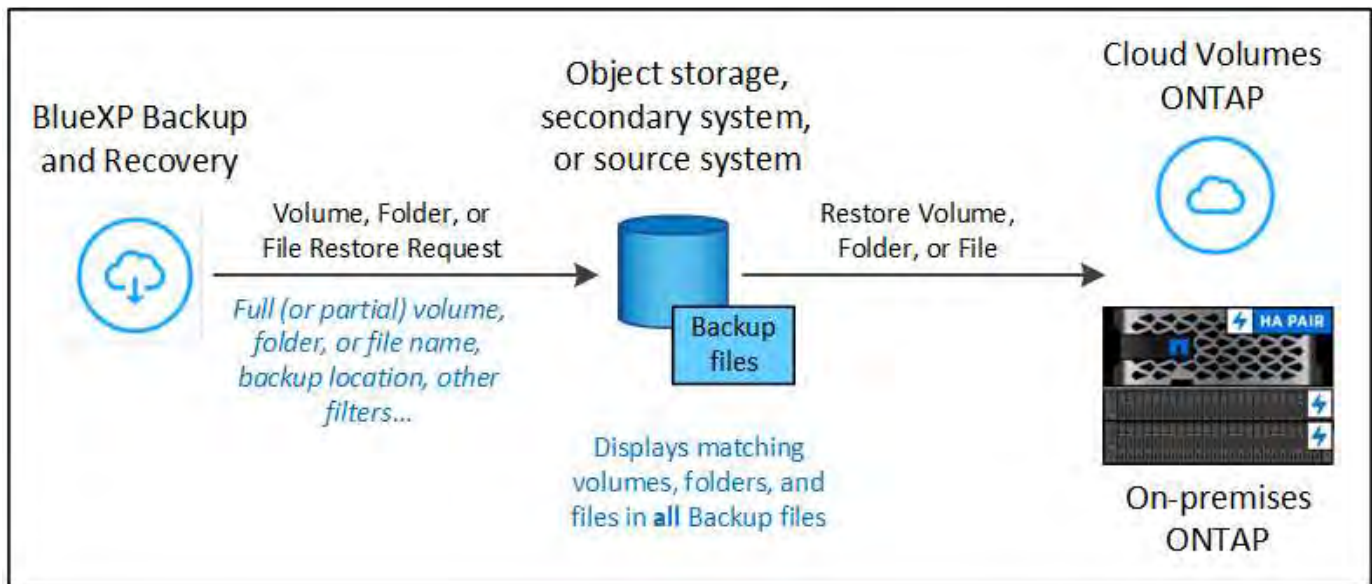


every volume.

2. When you want to restore a volume or files from a volume backup, under *Search & Restore*, select **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 select **Search**.

The Search Results page displays all the locations that have a file or volume that matches your search criteria.

4. Select **View All Backups** for the location you want to use to restore the volume or file, and then select **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 select **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.

The Indexed Catalog is a database that stores metadata about all the volumes and backup files in your working environment. It is used by the Search & Restore functionality to quickly find the backup files that contain the data you want to restore.

### Indexed Catalog v2 features

The Indexed Catalog v2, released in February 2025 and updated in June 2025, includes features that make it more efficient and easier to use. This version has a significant performance enhancement and is enabled by default for all new customers.

Review the following considerations regarding v2:

- The Indexed Catalog v2 is available in preview mode.

- If you are an existing customer and want to use the Catalog v2, you need to completely re-index your environment.
- The Catalog v2 indexes only those snapshots that have a snapshot label.
- BlueXP backup and recovery does not index snapshots with "hourly" SnapMirror labels. If you want to index snapshots with the "hourly" SnapMirror label, you need to enable it manually while the v2 is in preview mode.
- BlueXP backup and recovery will index volumes and snapshots associated with working environments protected by BlueXP backup and recovery only with the Catalog v2. Other working environments discovered on the BlueXP platform will not be indexed.
- Data indexing with Catalog v2 occurs in on-premises environments and in Amazon Web Services, Microsoft Azure, and Google Cloud Platform (GCP) environments.

The Indexed Catalog v2 supports the following:

- Global search efficiency in less than 3 minutes
- Up to 5 billion files
- Up to 5000 volumes per cluster
- Up to 100K snapshots per volume
- Maximum time for baseline indexing is less than 7 days. The actual time will vary depending on your environment.

### Enabling the Indexed Catalog for a working environment

The service does not provision a separate bucket when you use the Indexed Catalog v2. Instead, for backups stored in AWS, Azure, Google Cloud Platform, StorageGRID, or ONTAP S3, the service provisions space on the Connector or on the cloud provider environment.

If you enabled the Indexed Catalog prior to the v2 release, the following occurs with working environments:

- 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.

### Steps to enable Indexing for a working environment:

1. Do one of the following:
  - If no working environments have been indexed, on the Restore Dashboard under *Search & Restore*, select **Enable Indexing for Working Environments**.
  - If at least one working environment has already been indexed, on the Restore Dashboard under *Search & Restore*, select **Indexing Settings**.
2. Select **Enable Indexing** for the working environment.

### Result

After all the services are provisioned and the Indexed Catalog has been activated, the working environment is shown as "Active".

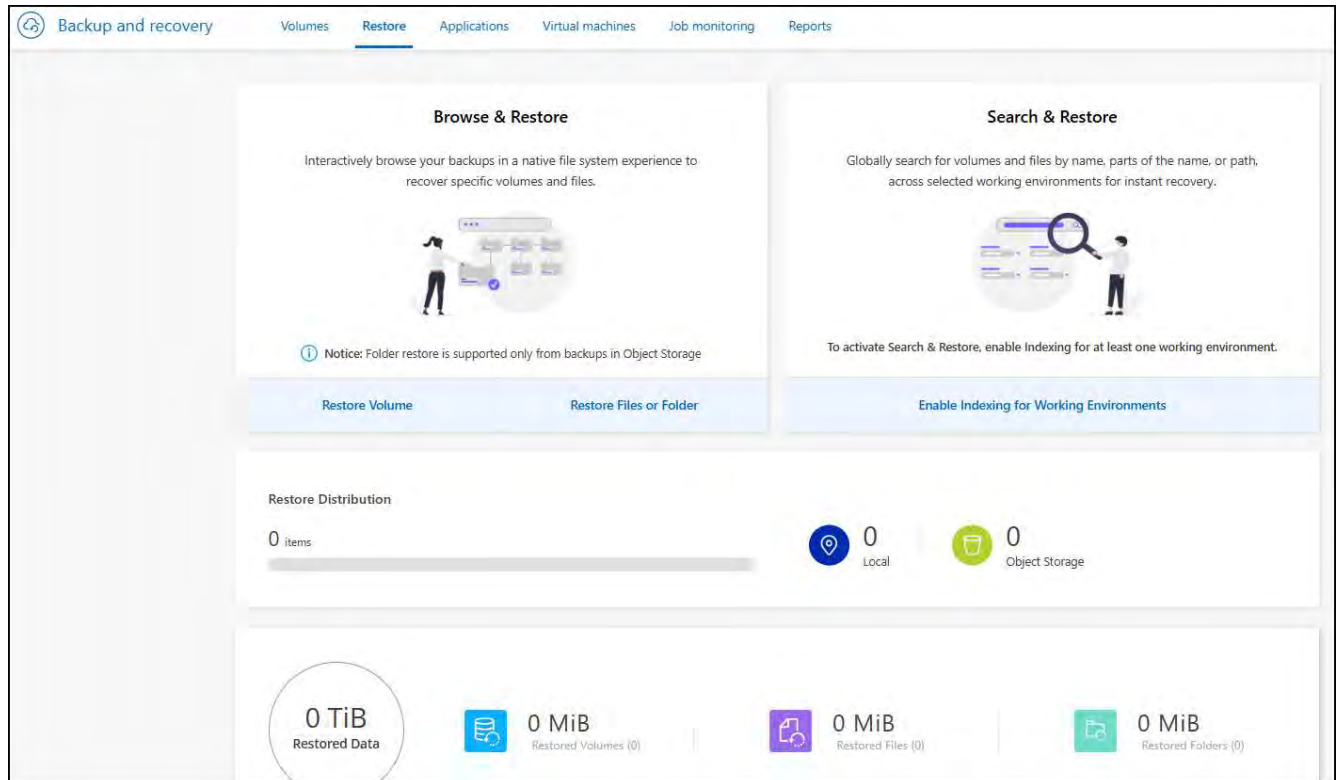
Depending on the size of the volumes in the working environment, and the number of backup files in all 3 backup locations, the initial indexing process could take up to an hour. After that it is transparently updated hourly with incremental changes to stay current.

### 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. Select the **Restore** tab and the Restore Dashboard is displayed.
3. From the *Search & Restore* section, select **Search & Restore**.
4. From the *Search & Restore* section, select **Search & Restore**.



5. From the Search & 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.
6. Select **Search** and the Search Results area displays all the resources that have a file, folder, or volume



that matches your search.

7. Locate the resource that has the data you want to restore and select **View All Backups** to display all the backup files that contain the matching volume, folder, or file.
8. Locate the backup file that you want to use to restore the data and select **Restore**.

Note that the results identify local volume Snapshot copies and remote Replicated volumes that contain the file in your search. You can choose to restore from the cloud backup file, from the Snapshot copy, or from the Replicated volume.

9. Select the destination location where you want the volume, folder, or file(s) to be restored and select **Restore**.
  - For volumes, you can select the original destination working environment or you can select an alternate working environment. When restoring a FlexGroup volume you'll need to choose multiple aggregates.
  - For folders, you can restore to the original location or you can select an alternate location; including the working environment, volume, and folder.
  - For files, you can restore to the original location or you can select an alternate location; including the working environment, volume, and folder. When selecting the original location, you can choose to overwrite the source file(s) or to create new file(s).

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

- When restoring from Amazon S3, select the IPspace in the ONTAP cluster where the destination volume will reside, enter the access key and secret key for the user you created to give the ONTAP cluster access to the S3 bucket, and optionally choose a private VPC endpoint for secure data transfer. [See details about these requirements.](#)
- When restoring from Azure Blob, select the IPspace in the ONTAP cluster where the destination volume will reside, and optionally choose a private endpoint for secure data transfer by selecting the VNet and Subnet. [See details about these requirements.](#)
- When restoring from Google Cloud Storage, select the IPspace in the ONTAP cluster where the destination volume will reside, and the Access Key and Secret Key to access the object storage. [See details about these requirements.](#)
- When restoring from StorageGRID, enter the FQDN of the StorageGRID server and the port that ONTAP should use for HTTPS communication with StorageGRID, enter the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume resides. [See details about these requirements.](#)
- When restoring from ONTAP S3, enter the FQDN of the ONTAP S3 server and the port that ONTAP should use for HTTPS communication with ONTAP S3, select the Access Key and Secret Key needed to access the object storage, and the IPspace in the ONTAP cluster where the destination volume will reside. [See details about these requirements.](#)

## Results

The volume, folder, or file(s) are restored and you are returned to the Restore Dashboard so you can review the progress of the restore operation. You can also select the **Job Monitoring** tab to see the restore progress. See [Job monitor page](#).

## Protect Microsoft SQL Server workloads

## Protect Microsoft SQL workloads overview with BlueXP backup and recovery

Protect your Microsoft SQL Server applications data from on-premises ONTAP systems to Amazon Web Services, Microsoft Azure, or StorageGRID using BlueXP backup and recovery. Backups are automatically generated and stored in an object store in your public or private cloud account based on the policies you create. 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.

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 available just in case the onsite copy is compromised.

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

You can accomplish the following protection goals:

- [Configure additional items if importing from SnapCenter](#)
- [Discover Microsoft SQL Server workloads and optionally import SnapCenter resources](#)
- [Back up workloads with local snapshots on local ONTAP primary storage](#)
- [Replicate workloads to ONTAP secondary storage](#)
- [Back up workloads to an object store location](#)
- [Back up workloads now](#)
- [Restore workloads](#)
- [Clone workloads](#)
- [Manage inventory of workloads](#)
- [Manage snapshots](#)

To back up workloads, typically you create policies that govern the backup and restore operations. See [Create policies](#) for more information.

### Supported backup destinations

BlueXP backup and recovery enables you to back up Microsoft SQL Server instances and databases 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 ONTAP S3

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

### Supported restore destinations

You can restore Microsoft SQL Server instances and databases from a backup that resides in primary storage or 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.

From Backup File Location		To Destination Working Environment
Object Store (Backup)	Secondary System (Replication)	
Amazon S3	Cloud Volumes ONTAP in AWS On-premises ONTAP system	Cloud Volumes in AWS On-premises ONTAP system ONTAP S3
Azure Blob	Cloud Volumes ONTAP in Azure On-premises ONTAP system	Cloud Volumes ONTAP in Azure On-premises ONTAP system ONTAP S3
StorageGRID	Cloud Volumes ONTAP On-premises ONTAP system	On-premises ONTAP system ONTAP S3
Amazon FSx for NetApp ONTAP	Amazon FSx for NetApp ONTAP	NA



References to "on-premises ONTAP systems" include FAS and AFF systems.

## Prerequisites for importing from the Plug-in service into BlueXP backup and recovery

If you are going to import resources from the SnapCenter Plug-in service for Microsoft SQL Server into BlueXP backup and recovery, you'll need to configure a few more items.

### Create working environments in BlueXP Canvas first

If you are going to import resources from SnapCenter, you should create working environments in BlueXP Canvas for all on-premises SnapCenter cluster storage first before importing from SnapCenter. This ensures that host resources can be discovered and imported correctly.

## Ensure host requirements to install the SnapCenter Plug-in

To import resources from the SnapCenter Plug-in for Microsoft SQL Server, ensure host requirements to install the SnapCenter Plug-in for Microsoft SQL Server are met.

Check specifically for the SnapCenter requirements in [BlueXP backup and recovery prerequisites](#).

## Disable User Account Control remote restrictions

Before you import resources from SnapCenter, disable User Account Control (UAC) remote restrictions on the SnapCenter Windows host. Disable UAC if you use a local administrative account to connect remotely to the SnapCenter Server host or the SQL host.

## Security considerations

Consider the following issues before disabling UAC remote restrictions:

- Security risks: Disabling token filtering can expose your system to security vulnerabilities, especially if local administrative accounts are compromised by malicious actors.
- Use with caution:
  - Modify this setting only if it is essential for your administrative tasks.
  - Ensure that strong passwords and other security measures are in place to protect administrative accounts.

## Alternative solutions

- If remote administrative access is required, consider using domain accounts with appropriate privileges.
- Use secure remote management tools that adhere to best security practices to minimize risks.

## Steps to disable User Account Control remote restrictions

1. Modify the `LocalAccountTokenFilterPolicy` registry key on the SnapCenter Windows host.

Do this by using one of the following, with instructions next:

- Method 1: Registry Editor
- Method 2: PowerShell script

### Method 1: Disable User Account Control by using the Registry Editor

This is one of the methods that you can use to disable User Account Control.

#### Steps

1. Open the Registry Editor on the SnapCenter Windows host by doing the following:
  - a. Press `Windows+R` to open the Run dialog box.
  - b. Type `regedit` and press `Enter`.
2. Navigate to the Policy Key:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System
```

3. Create or modify the `DWORD` value:

- a. **Locate:** LocalAccountTokenFilterPolicy
  - b. If it doesn't exist, create a new DWORD (32-bit) Value named LocalAccountTokenFilterPolicy.
4. The following values are supported. For this scenario, set the value to 1:
- 0 (Default): UAC remote restrictions are enabled. Local accounts have filtered tokens when accessing remotely.
  - 1: UAC remote restrictions are disabled. Local accounts bypass token filtering and have full administrative privileges when accessing remotely.
5. Click **OK**.
  6. Close the Registry Editor.
  7. Restart the SnapCenter Windows host.

### Example registry modification

This example sets LocalAccountTokenFilterPolicy to "1", disabling UAC remote restrictions.

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]

"LocalAccountTokenFilterPolicy"=dword:00000001
```

### Method 2: Disable User Account Control by using a PowerShell script

This is another method that you can use to disable User Account Control.



Running PowerShell commands with elevated privileges can affect system settings. Ensure you understand the commands and their implications before running them.

### Steps

1. Open a PowerShell window with administrative privileges on the SnapCenter Windows host:
  - a. Click on the **Start** menu.
  - b. Search for **PowerShell 7** or **Windows Powershell**.
  - c. Right-click on that option and select **Run as administrator**.
2. Ensure that PowerShell is installed on your system. After installation, it should appear in the **Start** menu.



PowerShell is included by default in Windows 7 and later versions.

3. To disable UAC remote restrictions, set LocalAccountTokenFilterPolicy to "1" by running the following command:

```
Set-ItemProperty -Path
"HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System" -Name
"LocalAccountTokenFilterPolicy" -Value 1 -Type DWord
```

4. Verify that the current value is set to "1" in LocalAccountTokenFilterPolicy` by running:

```
Get-ItemProperty -Path
"HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System" -Name
"LocalAccountTokenFilterPolicy"
```

- If the value is 1, UAC remote restrictions are disabled.
- If the value is 0, UAC remote restrictions are enabled.

5. To apply the changes, restart your computer.

#### Example PowerShell 7 commands to disable UAC remote restrictions:

This example with the value set to "1" indicates that UAC remote restrictions are disabled.

```
# Disable UAC remote restrictions

Set-ItemProperty -Path
"HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System" -Name
"LocalAccountTokenFilterPolicy" -Value 1 -Type DWord

# Verify the change

Get-ItemProperty -Path
"HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System" -Name
"LocalAccountTokenFilterPolicy"

# Output

LocalAccountTokenFilterPolicy : 1
```

## Discover Microsoft SQL Server workloads and optionally import from SnapCenter in BlueXP backup and recovery

The BlueXP backup and recovery service needs to first discover Microsoft SQL Server workloads in order for you to use the service. You can optionally import backup data and policies from SnapCenter if you already have SnapCenter installed.

### Required BlueXP role

This task requires the data services Backup and recovery super admin role. Learn about [Backup and recovery data services roles and privileges](#). [Learn about BlueXP access roles for all services](#).

### Discover Microsoft SQL Server workloads and optionally import SnapCenter resources

During discovery, BlueXP backup and recovery analyzes Microsoft SQL Server instances and databases in working environments within your organization.

BlueXP backup and recovery assesses Microsoft SQL Server applications. The service assesses the existing protection level including the current backup protection policies, snapshot copies, and backup and recovery options.

Discovery occurs in the following ways:

- If you already have SnapCenter, import SnapCenter resources into BlueXP backup and recovery by using the BlueXP backup and recovery UI.



If you already have SnapCenter, first check to be sure you've met the prerequisites before importing from SnapCenter. For example, you should create working environments in BlueXP Canvas for all on-premises SnapCenter cluster storage first before importing from SnapCenter. See [Prerequisites for importing resources from SnapCenter](#).

- If you don't already have SnapCenter, you can still discover workloads within your working environments by adding a vCenter manually and performing discovery.

### **If SnapCenter is already installed, import SnapCenter resources into BlueXP backup and recovery**

If you already have SnapCenter installed, import SnapCenter resources into BlueXP backup and recovery using these steps. The BlueXP service discovers resources, hosts, credentials, and schedules from SnapCenter; you don't have to recreate all that information.

You can do this in the following ways:

- During discovery, select an option to import resources from SnapCenter.
- After discovery, from the Inventory page, select an option to import SnapCenter resources.
- After discovery, from the Settings menu, select an option to import SnapCenter resources. For details, see [Configure BlueXP backup and recovery](#).

This is a two-part process:

- Import SnapCenter Server application and host resources
- Manage selected SnapCenter host resources

### **Import SnapCenter Server application and host resources**

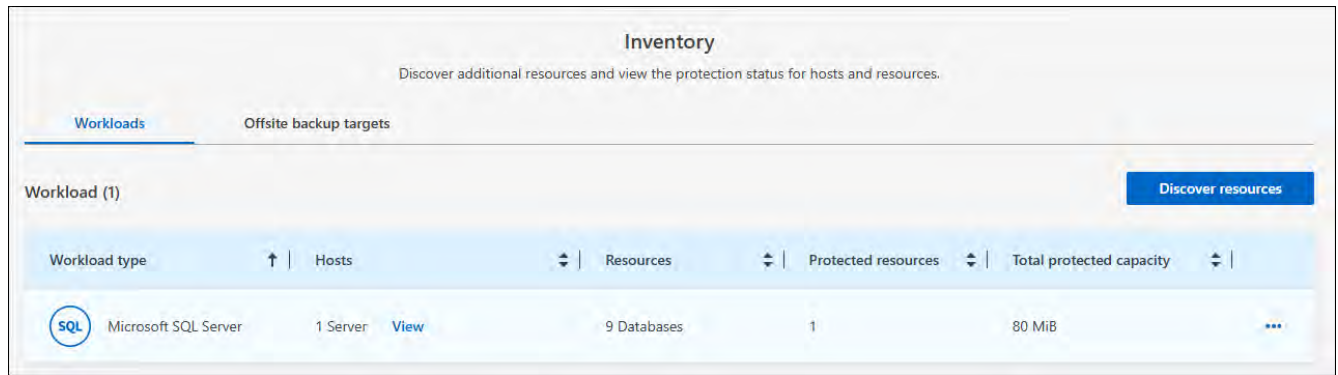
This first step imports host resources from SnapCenter and displays those resources in the BlueXP backup and recovery Inventory page. At that point, the resources are not yet managed by BlueXP backup and recovery.



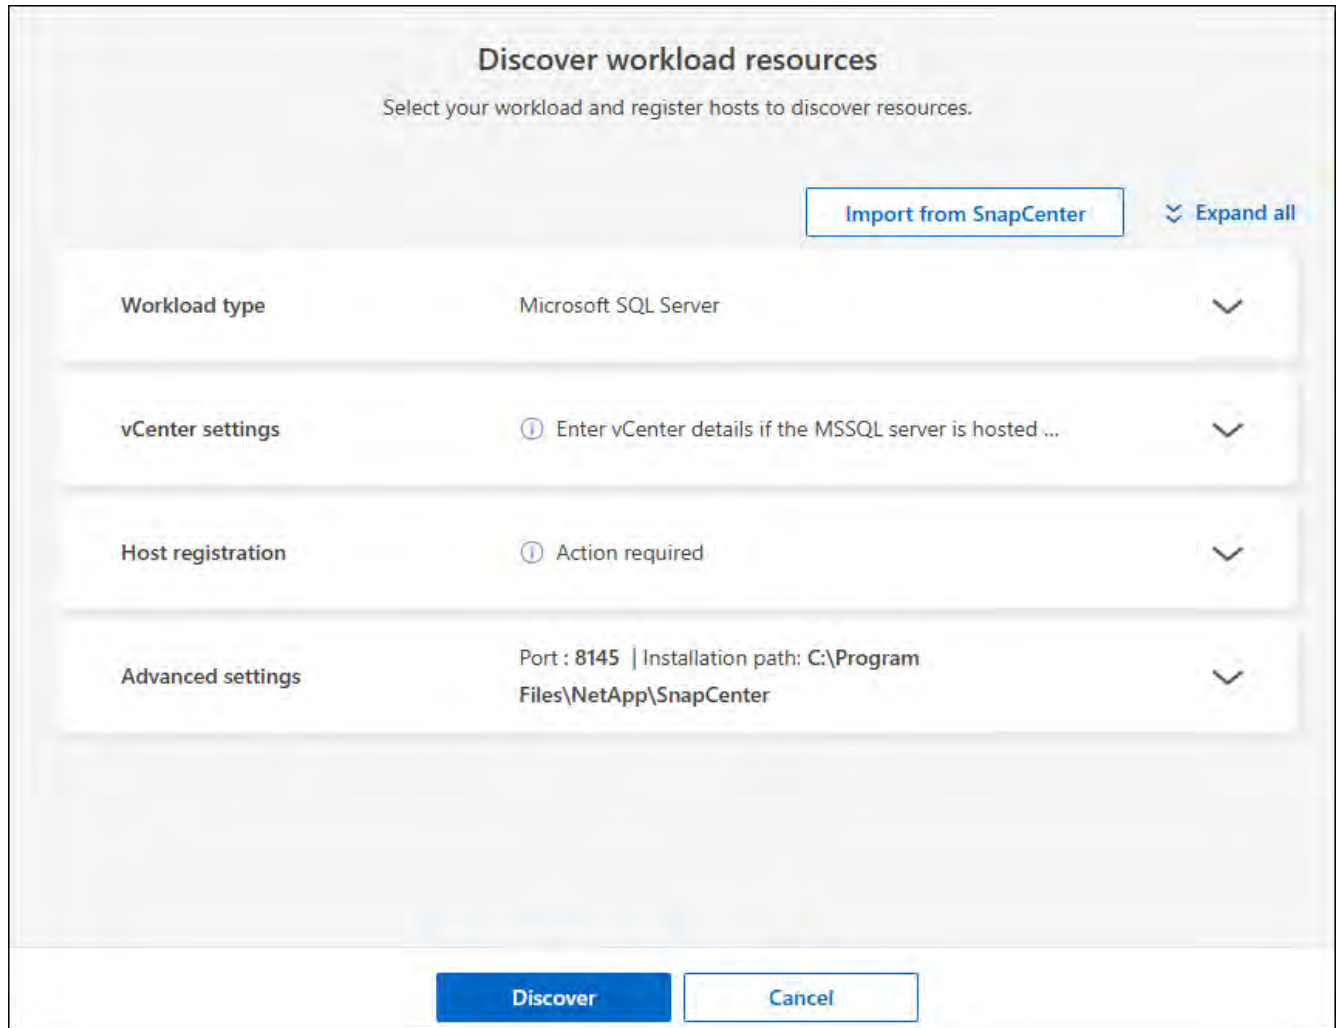
After you import SnapCenter host resources, BlueXP backup and recovery does not take over protection management automatically. To do so, you must explicitly select to manage the imported resources in BlueXP backup and recovery. This ensures that you are ready to have those resources backed up by BlueXP backup and recovery.

### **Steps**

1. From the BlueXP left navigation, select **Protection > Backup and recovery**.
2. From the top menu, select **Inventory**.



3. From the top menu, select **Discover resources**.



4. From the BlueXP backup and recovery Discover workload resources page, select **Import from SnapCenter**.



### Import from SnapCenter

Enter the SnapCenter application credentials to connect securely and import SnapCenter managed applications.

**Import from:**  SnapCenter

**SnapCenter application credentials**

Enter the SnapCenter connection details to establish a secure connection and import SnapCenter managed application hosts to BlueXP backup and recovery.

SnapCenter FQDN or IP Address	SnapCenter port number
<input type="text" value="Enter FQDN or IP address"/>	<input type="text" value="8146"/>
SnapCenter user name	SnapCenter password
<input type="text" value="Enter user name"/>	<input type="password" value="Enter password"/> <input type="checkbox"/>

Connectors

TestathonConnect X ▼

**SnapCenter server host credentials**

You can use SnapCenter host credentials you already added or supply additional credentials.

Existing credentials     Add new credentials

Credentials name Authentication mode

5. Enter **SnapCenter application credentials**:
  - a. **SnapCenter FQDN or IP address**: Enter the FQDN or IP address of the SnapCenter application itself.
  - b. **Port**: Enter the port number for the SnapCenter Server.
  - c. **Username** and **Password**: Enter the username and password for the SnapCenter Server.
  - d. **Connector**: Select the BlueXP Connector for SnapCenter.
6. Enter **SnapCenter server host credentials**:
  - a. **Existing credentials**: If you select this option, you can use the existing credentials that you have already added. Choose the credentials name.
  - b. **Add new credentials**: If you don't have existing SnapCenter host credentials, you can add new credentials. Enter the credentials name, authentication mode, user name, and password.
7. Select **Import** to validate your entries and register the SnapCenter Server.



If the SnapCenter Server is already registered, you can update the existing registration details.

## Result

The Inventory page shows the imported SnapCenter resources that include MS SQL hosts, instances, and databases.

Workload type	Hosts	Resources	Protected resources	Total protected capacity
Microsoft SQL Server	1 Server <a href="#">View</a>	9 Databases	1	80 MiB

To see the details of the imported SnapCenter resources, select the **View details** option from the Actions menu.

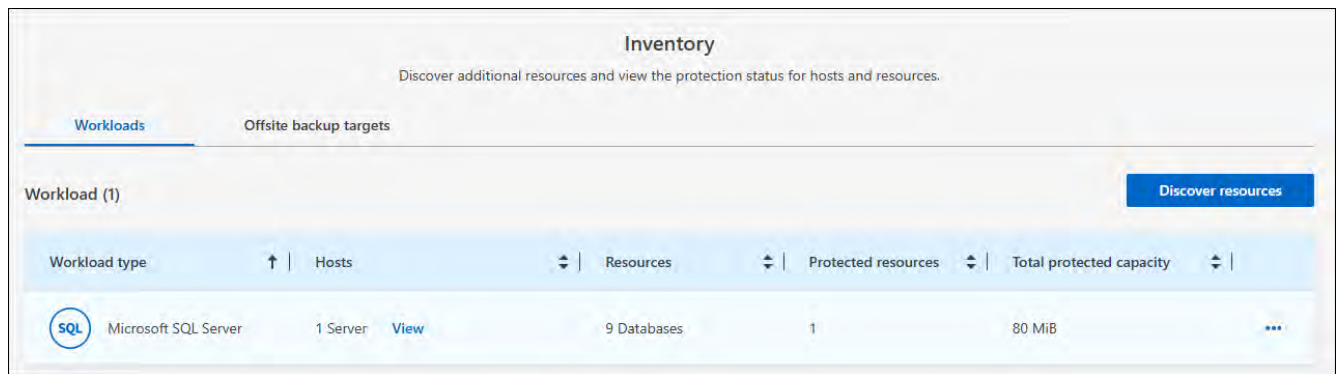
Host name	Instances	Deployment model	Nodes assigned	Connector
Host_name Unmanaged		Cluster	3 <a href="#">View</a>	Connector_1
Host_name Unmanaged		Standalone		Connector_2
Host_name Unmanaged		Standalone		Connector_3
Host_name Unmanaged		Standalone		Connector_1

## Manage SnapCenter host resources

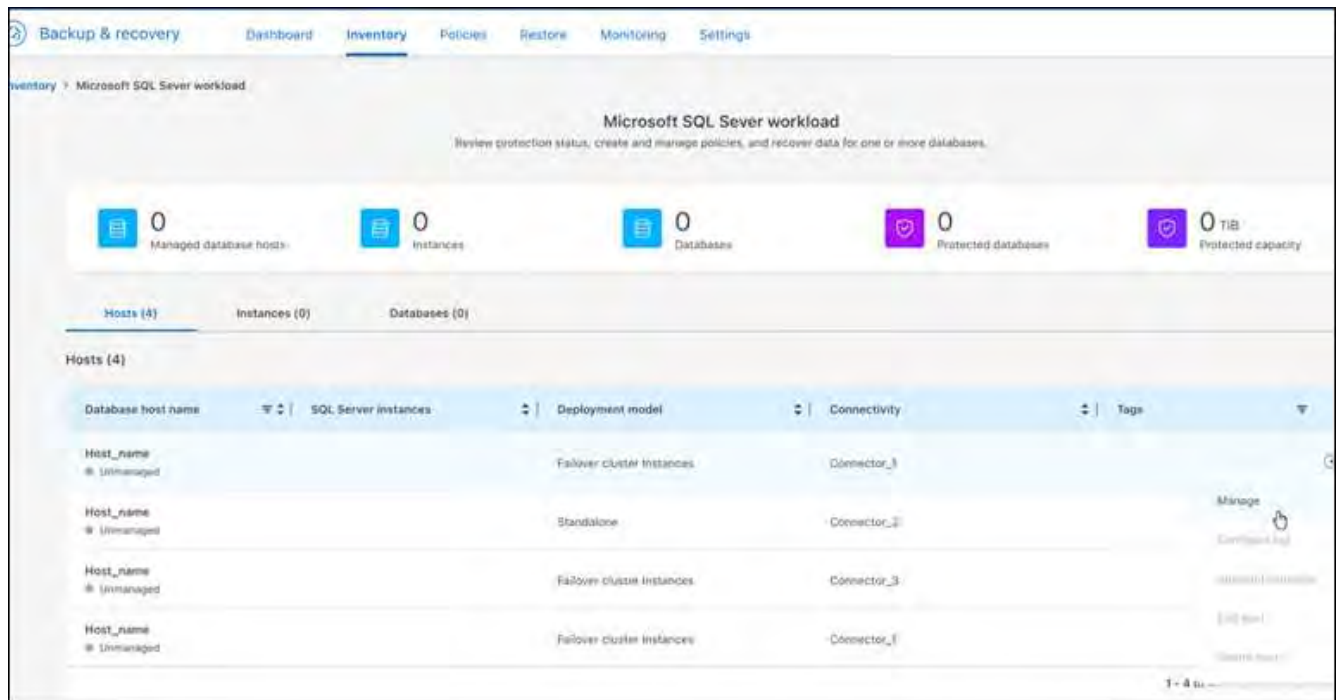
After you import the SnapCenter resources, manage those host resources in BlueXP backup and recovery. After you select to manage those resources, BlueXP backup and recovery is able to back up and recover the resources that you imported from SnapCenter. You no longer manage those resources in SnapCenter Server.

### Steps

1. After you import the SnapCenter resources, from the top menu, select **Inventory**.
2. From the Inventory page, select the imported SnapCenter host that you want to have BlueXP backup and recovery to manage from now on.



3. Select the Actions icon **...** > **View details** to display the workload details.



4. From the Inventory > workload page, select the Actions icon **...** > **Manage** to display the Manage host page.

5. Select **Manage**.

6. In the Manage host page, select either to use an existing vCenter or add a new vCenter.

7. Select **Manage**.

The Inventory page shows the newly managed SnapCenter resources.

You can optionally create a report of the managed resources by selecting the **Generate reports** option from the Actions menu.


### Import SnapCenter resources after discovery from the Inventory page

If you have already discovered resources, you can import SnapCenter resources from the Inventory page.

#### Steps

1. From the BlueXP left navigation, select **Protection** > **Backup and recovery**.

2. From the top menu, select **Inventory**.



3. From the Inventory page, select **Import SnapCenter resources**.

4. Follow the steps in the **Import SnapCenter resources** section above to import SnapCenter resources.

#### If you don't have SnapCenter installed, add a vCenter and discover resources

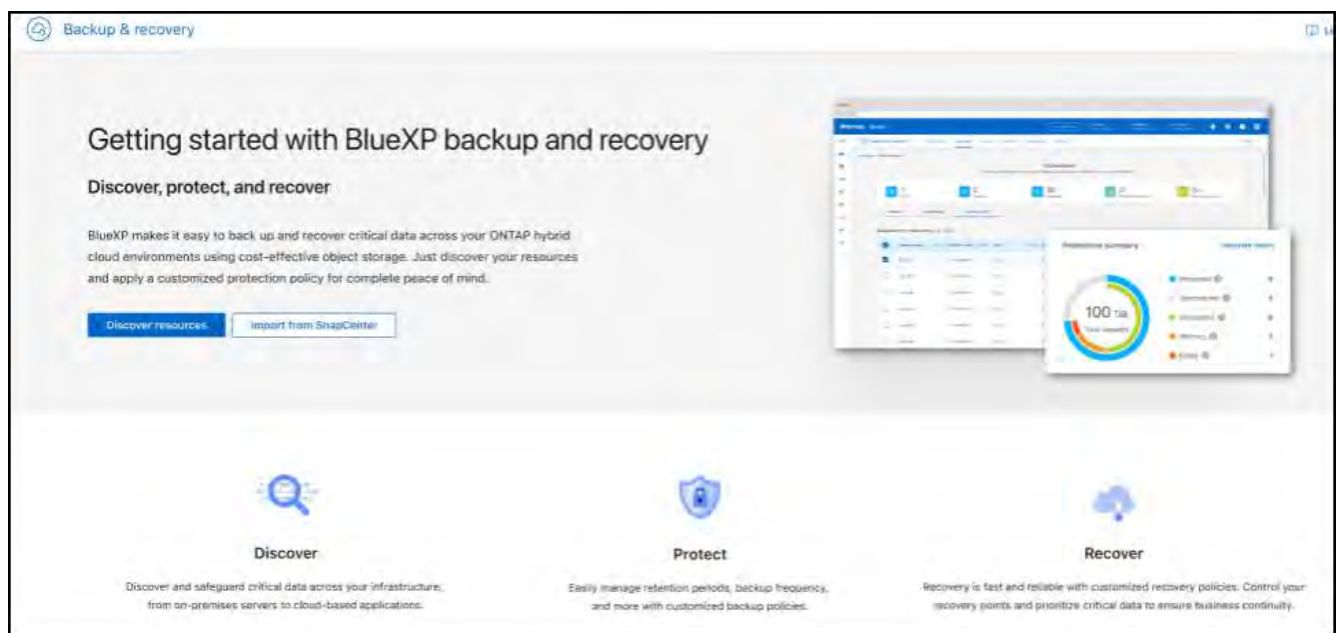
If you don't already have SnapCenter installed, you can add vCenter information and have BlueXP backup and recovery discover workloads. Within each BlueXP Connector, select the working environments where you want to discover workloads.

This is optional if you have a VMware environment.

#### Steps

1. From the BlueXP left navigation, select **Protection > Backup and recovery**.

If this is your first time logging in to this service, you already have a working environment in BlueXP, but haven't discovered any resources, the "Welcome to the new BlueXP backup and recovery" landing page appears and shows an option to **Discover resources**.



2. Select **Discover resources**.

### Discover workload resources

Select your workload and register hosts to discover resources.

Import from SnapCenter
⌵ Expand all

<b>Workload type</b>	Microsoft SQL Server	⌵
<b>vCenter settings</b>	ⓘ Enter vCenter details if the MSSQL server is hosted ...	⌵
<b>Host registration</b>	ⓘ Action required	⌵
<b>Advanced settings</b>	Port : 8145   Installation path: C:\Program Files\NetApp\SnapCenter	⌵

Discover
Cancel

3. Enter the following information:

- a. **Workload type:** For this version, only Microsoft SQL Server is available.
- b. **vCenter settings:** Select an existing vCenter or add a new one. To add a new vCenter, enter the vCenter FQDN or IP address, user name, password, port, and protocol.



If you are entering vCenter information, enter information for both vCenter settings and Host registration. If you added or entered vCenter information here, you also need to add plugin information in Advanced Settings next.

- c. **Host registration:** Select **Add credentials** and enter information about the hosts containing the workloads you want to discover.



If you are adding a standalone server and not a vCenter server, enter only the host information.

4. Select **Discover**.



This process might take a few minutes.

5. Continue with Advanced Settings.



## Set Advanced settings options during discovery and install the plugin

With Advanced Settings, you can manually install the plugin agent on all servers being registered. This enables you to import all SnapCenter workloads into BlueXP backup and recovery so you can manage backups and restores there. BlueXP backup and recovery shows the steps needed to install the plugin.

### Steps

1. From the Discover resources page, continue to Advanced Settings by clicking the down arrow on the right.

**Discover workload resources**  
Select your workload and register hosts to discover resources.

[Import from SnapCenter](#) [Expand all](#)

Workload type: Microsoft SQL Server

vCenter settings: Enter vCenter details if the MSSQL server is hosted in virtualiz...

Host registration: Action required

Advanced settings

Plug-in port: 8145

Installation path: C:\Program Files\NetApp\SnapCenter

If you want to install the agent manually:

Use manual installation. [Show me how?](#)

Add all hosts in the cluster

Skip optional preinstall checks

[Discover](#) [Cancel](#)

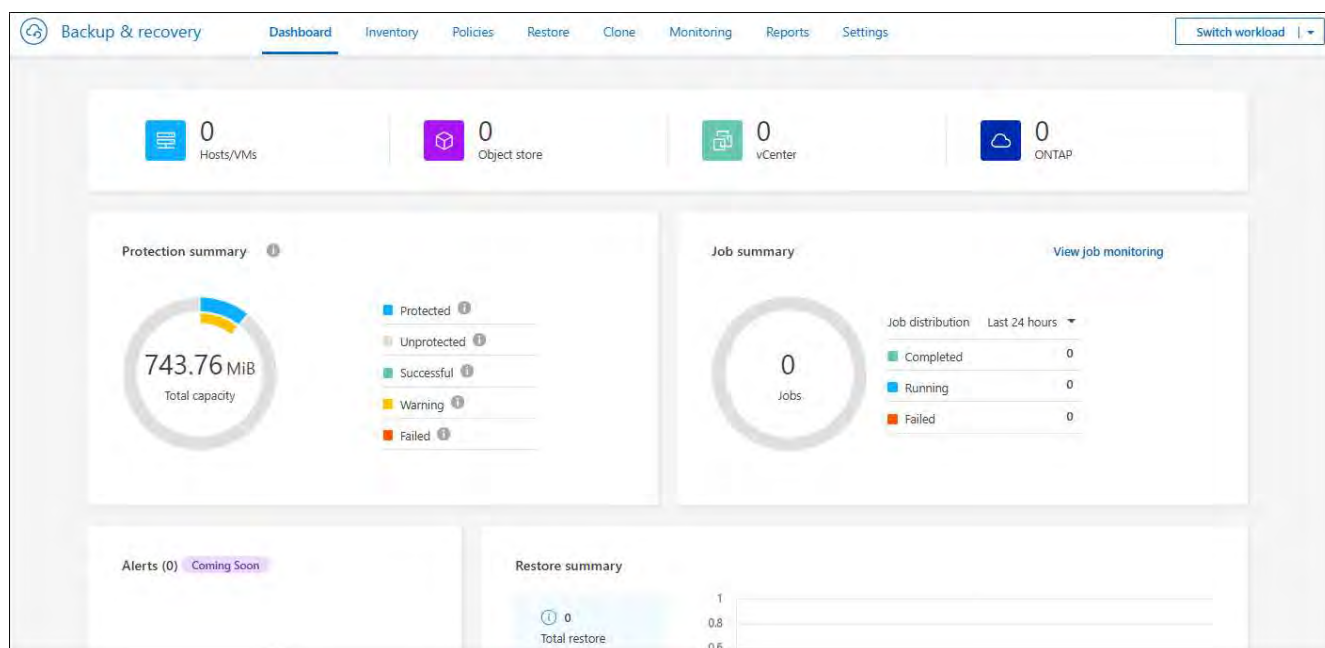
2. In the Discover workload resources page, enter the following information.
  - **Enter plug-in port number:** Enter the port number that the plugin uses.
  - **Installation path:** Enter the path where the plugin will be installed.
3. If you want to install the SnapCenter agent manually, check the boxes for the following options:
  - **Use manual installation:** Check this box to install the plugin manually.
  - **Add all hosts in the cluster:** Check this box to add all hosts in the cluster to BlueXP backup and recovery during discovery.

- **Skip optional preinstall checks:** Check this box to skip optional preinstall checks. You might want to do this for example, if you know that memory or space considerations will be changed in the near future and you want to install the plugin now.

#### 4. Select **Discover**.

### Continue to the BlueXP backup and recovery Dashboard

1. To display the BlueXP backup and recovery Dashboard, from the top menu, select **Dashboard**.
2. Review the health of data protection. The number of at risk or protected workloads increases based on the newly discovered, protected, and backed up workloads.



[Learn what the Dashboard shows you.](#)

## Back up Microsoft SQL Server workloads with BlueXP backup and recovery

Back up Microsoft SQL Server applications data from on-premises ONTAP systems to Amazon Web Services, Microsoft Azure, and StorageGRID to ensure that your data is protected. Backups are automatically generated and stored in an object store in your public or private cloud account.

- To back up workloads on a schedule, create policies that govern the backup and restore operations. See [Create policies](#) for instructions.
- Configure the log directory for discovered hosts before you initiate a backup.
- Back up workloads now (create an on-demand backup now).

### View workload protection status

Before you initiate a backup, view the protection status of your workloads.

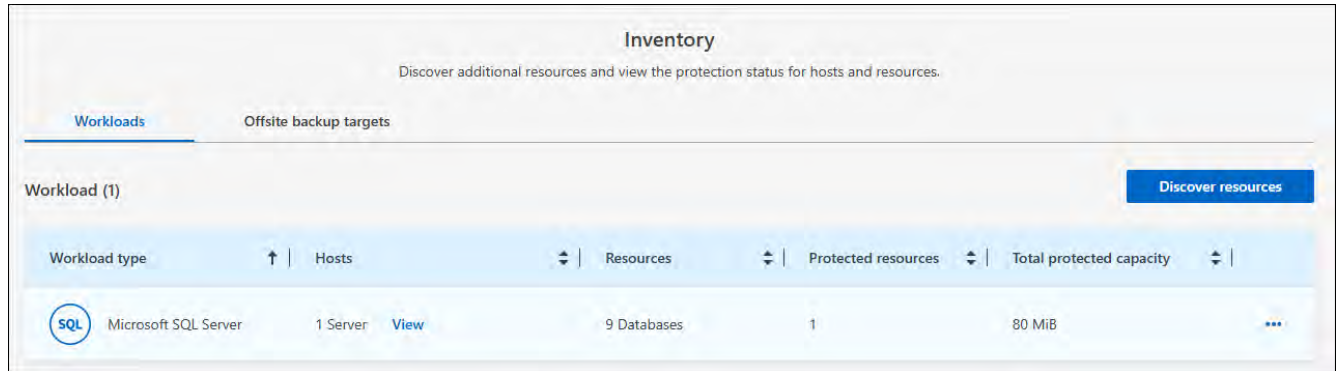
### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery

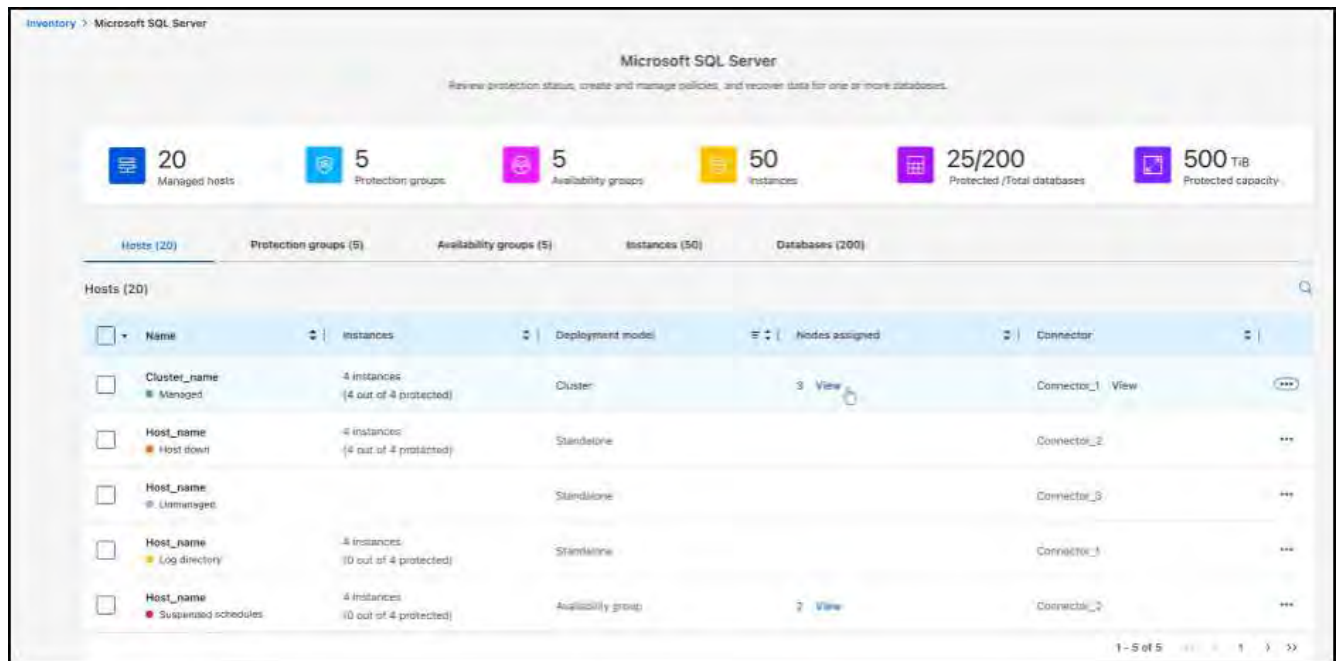
backup admin, Backup and Recovery restore admin, Backup and Recovery clone admin, or Backup and Recovery viewer role. Learn about [Backup and recovery roles and privileges](#). Learn about [BlueXP access roles for all services](#).

## Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.



2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.



4. Review details on the Hosts, Protection groups, Availability groups, Instances, and Databases tabs.

## Configure the log directory for discovered hosts

Before you back up your workloads, set the path for the activity logs for discovered hosts. This helps you to track the status of operations.

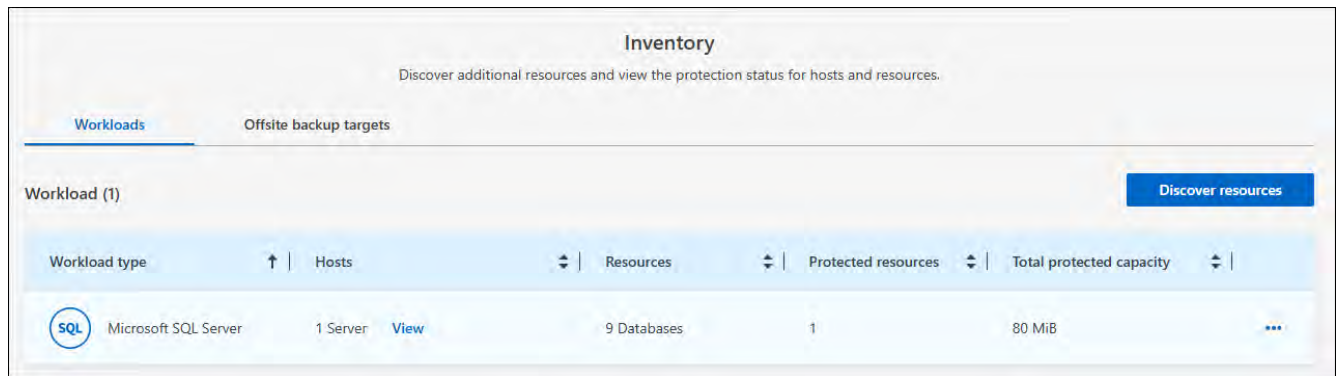
### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin, or Backup and Recovery restore admin role. [Learn about BlueXP access roles for all services](#).

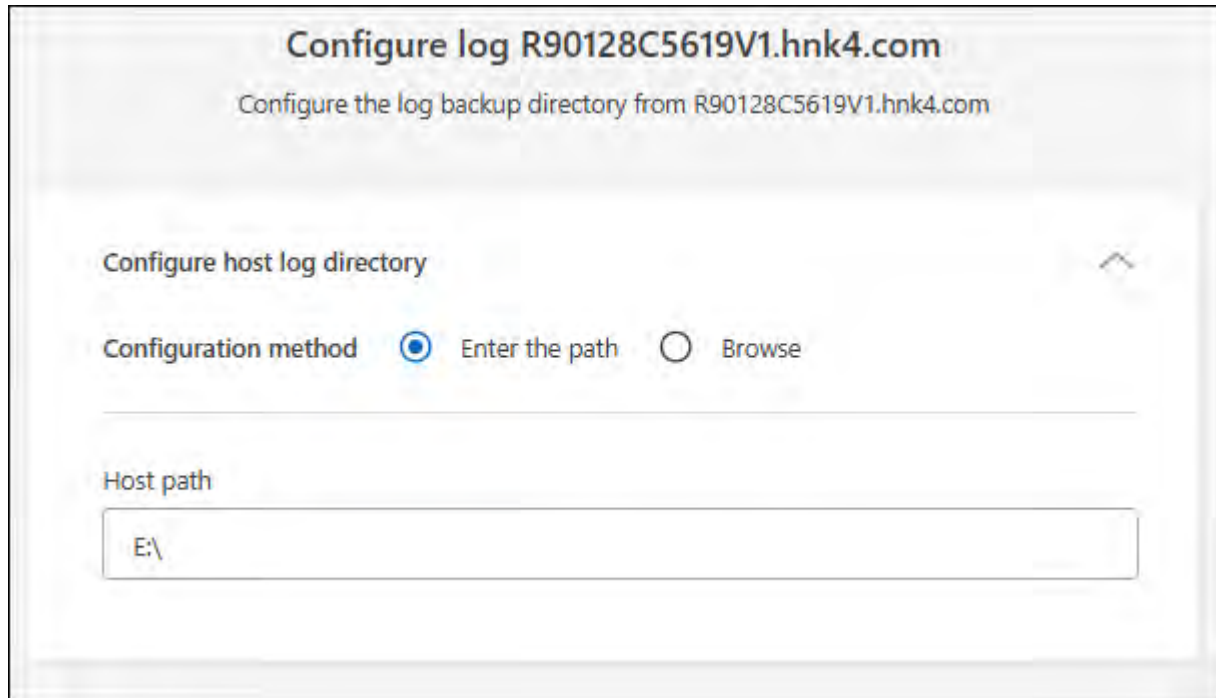
## Steps



1. From the BlueXP backup and recovery menu, select **Inventory**.



2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select a host.
5. Select the Actions icon **...** > **Configure log directory**.



6. Provide the host path or browse through a list of hosts or nodes hosts on the host to locate where you want the host log to be stored.
7. Select those on which you want to store the logs.



The fields that appear differ depending on the selected deployment model, for example, failover cluster instance or standalone.

8. Select **Save**.

## Create a protection group

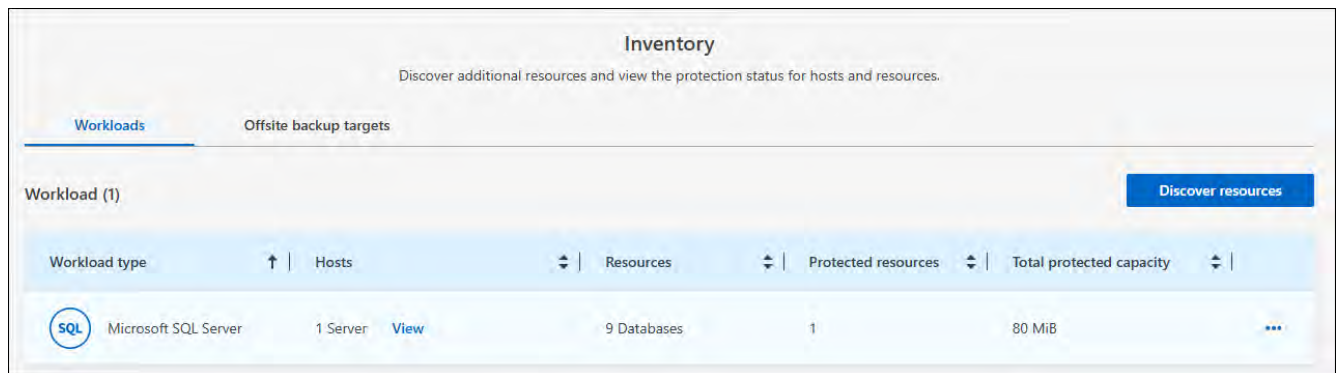
You can create a protection group to manage the backup and restore operations for a set of workloads. A protection group is a logical grouping of workloads that you want to protect together.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about BlueXP access roles for all services.](#)

### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.



2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select the **Protection groups** tab.
5. Select **Create protection group**.
6. Provide a name for the protection group.
7. Select the instances or databases that you want to include in the protection group.
8. Select **Next**.
9. Select the **Backup policy** that you want to apply to the protection group.

If you want to create a policy, select **Create new policy** and follow the prompts to create a policy. See [Create policies](#) for more information.

10. Select **Next**.
11. Review the configuration.
12. Select **Create** to create the protection group.

## Back up workloads now with an on-demand backup

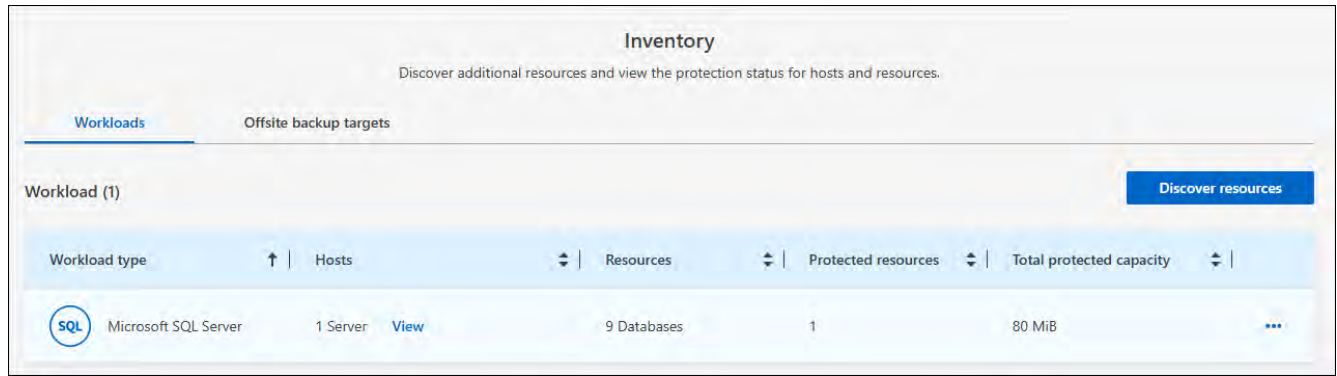
Create an on-demand backup immediately. You might want to run an on-demand backup if you're about to make changes to your system and want to ensure that you have a backup before you start.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about BlueXP access roles for all services.](#)

### Steps

1. From the menu, select **Inventory**.



2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select the **Protection Group**, **Instances** or **Databases** tab.
5. Select the instance or database you want to back up.
6. Select the Actions icon **...** > **Back up now**.
7. Select the policy that you want to apply to the backup.
8. Select the schedule tier.
9. Select **Back up now**.

### Suspend the backup schedule

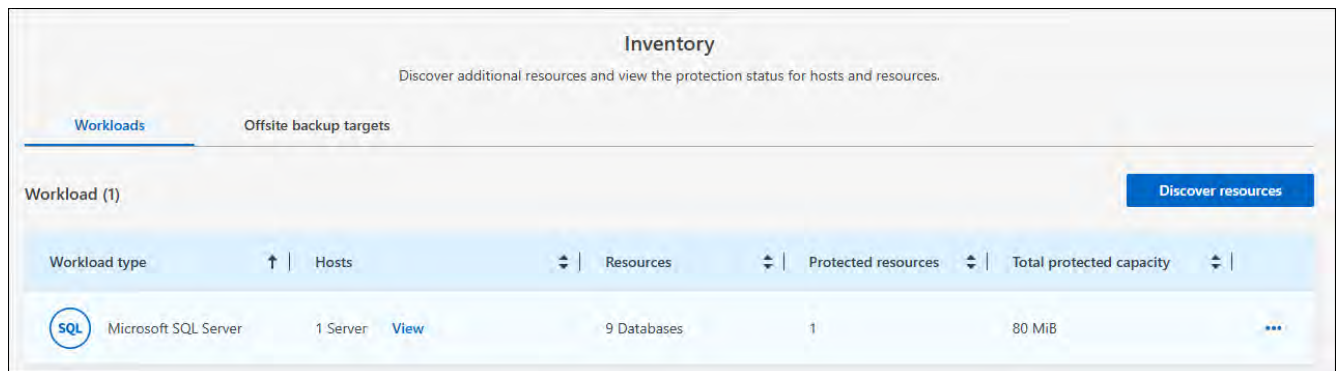
Suspending the schedule prevents the backup from running at the scheduled time temporarily. You might want to do this if you're performing maintenance on the system or if you're experiencing issues with the backup.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, or Backup and Recovery clone admin role. [Learn about BlueXP access roles for all services.](#)

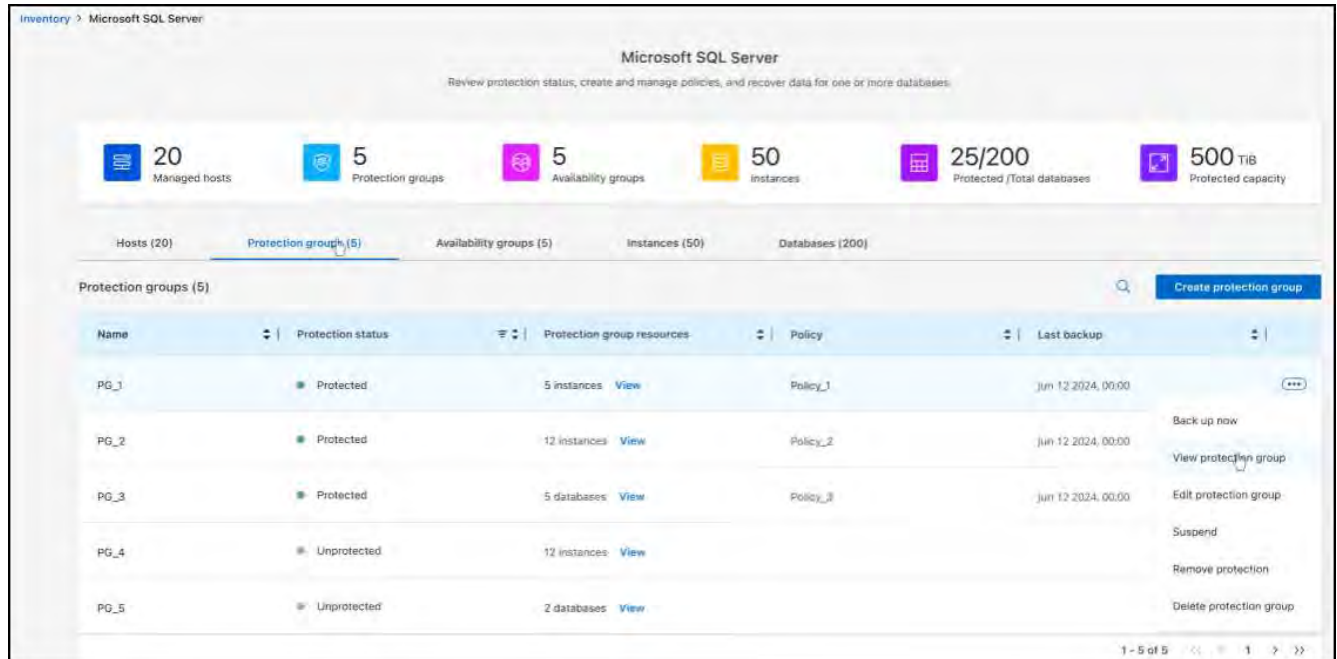
### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.



2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.

4. Select the **Protection Group, Instances** or **Databases** tab.
5. Select the protection group, instance, or database you want to suspend.



6. Select the Actions icon **...** > **Suspend**.

## Delete a protection group

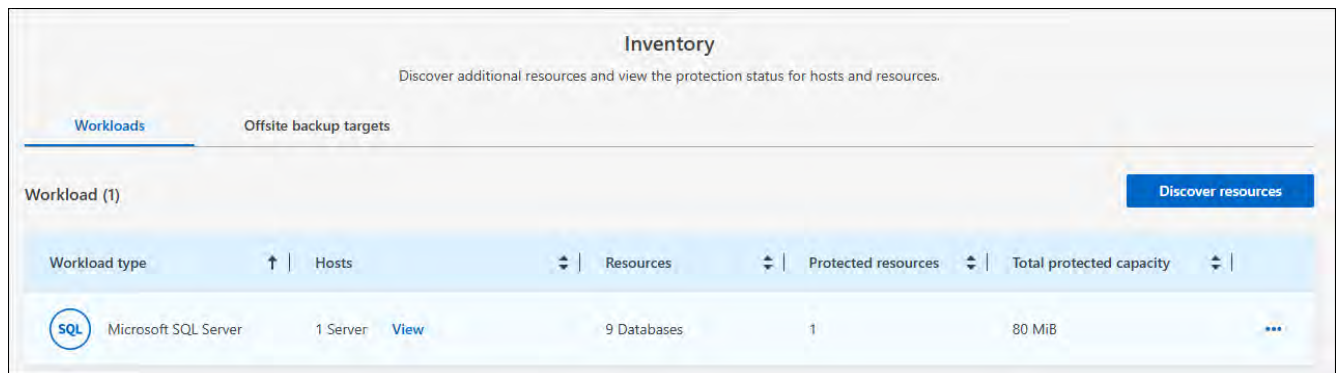
You can create a protection group to manage the backup and restore operations for a set of workloads. A protection group is a logical grouping of workloads that you want to protect together.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about BlueXP access roles for all services.](#)

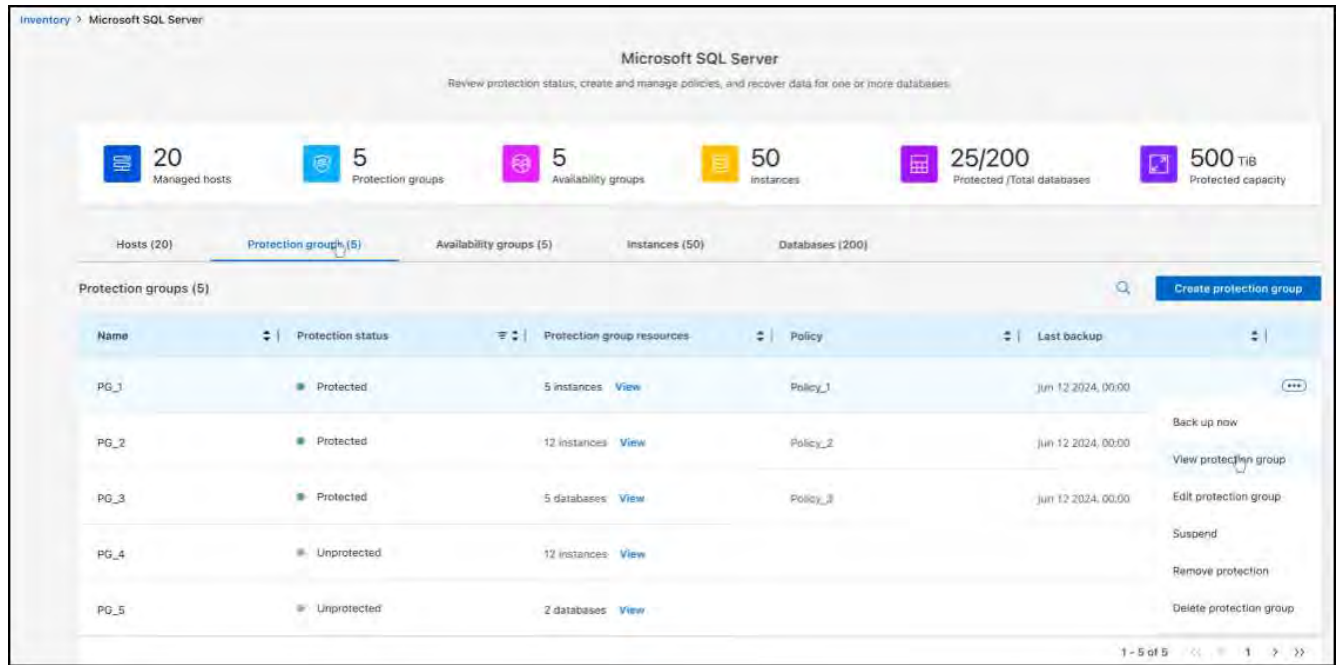
### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.



2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select the **Protection groups** tab.

5. Select the Actions icon **...** > **Delete protection group**.



## Remove protection from a workload

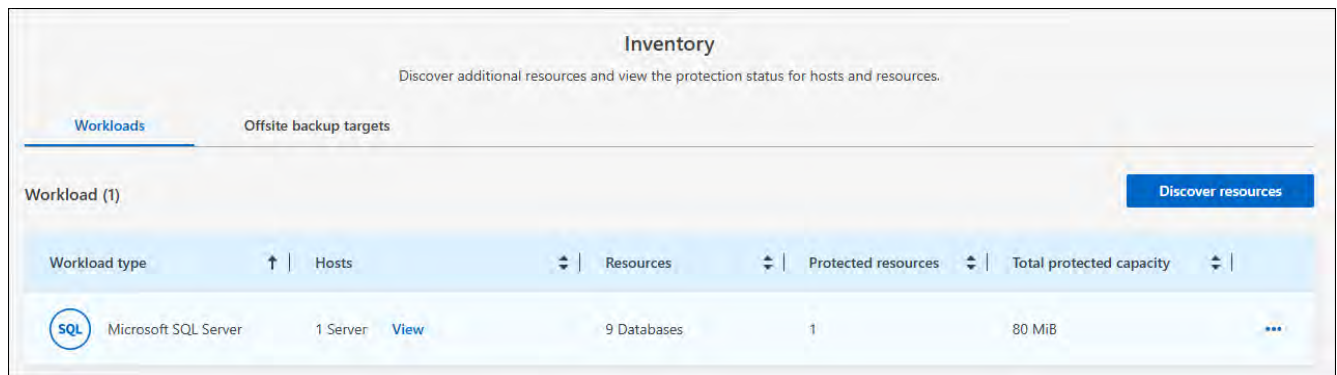
You can remove protection from a workload if you no longer want to back it up or if you want to stop managing it in BlueXP backup and recovery.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, or Backup and Recovery backup admin role. [Learn about BlueXP access roles for all services.](#)

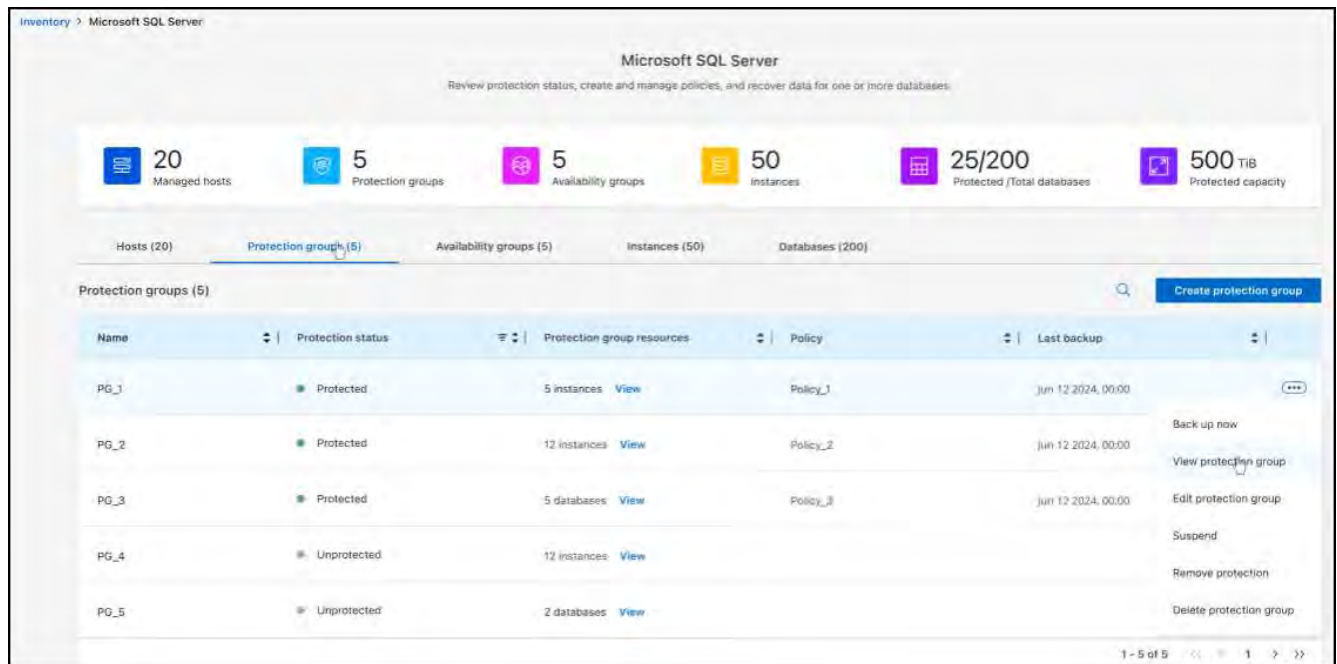
### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.



2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.
4. Select the **Protection Group**, **Instances** or **Databases** tab.
5. Select the protection group, instance, or database.





6. Select the Actions icon **...** > **Remove protection**.
7. In the Remove protection dialog box, select whether you want to keep backups and metadata or delete them.
8. Select **Remove** to confirm the action.

## Restore Microsoft SQL Server workloads with BlueXP backup and recovery

Restore Microsoft SQL Server workloads from snapshot copies, from a workload backup replicated to secondary storage, or from backups stored in object storage using BlueXP backup and recovery. You can restore a workload to the original working environment, to a different working environment that's using the same cloud account, or to an on-premises ONTAP system.

### Restore from these locations

You can restore workloads from different starting locations:

- Restore from a primary location
- Restore from a replicated resource
- Restore from an object store backup

### Restore to these points

You can restore data to the latest snapshot or to these points:

- Restore from snapshots
- Restore to a specific point in time. This is helpful if you know the name and location of the file, and the date when it was last in good shape.
- Restore to the latest backup

### Restore from object storage considerations

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 integrity check on the backup file before restoring the data. We recommend that you perform the scan.

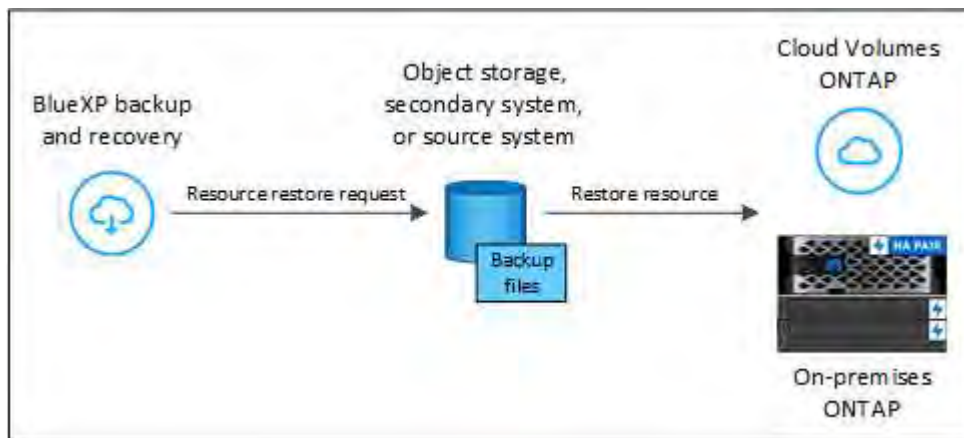


You'll incur extra egress costs from your cloud provider to access the contents of the backup file.

## How restoring workloads works

When you restore workloads, the following occurs:

- When you restore a workload from a backup file, BlueXP backup and recovery creates a *new* resource using the data from the backup.
- When you restore from a replicated workload, you can restore the workload to the original working environment or to an on-premises ONTAP system.



- When you restore a backup from object storage, you can restore the data to the original working environment or to an on-premises ONTAP system.

## Restore methods

You can restore workloads using one of the following methods. Typically, choose one of the following methods based on your restore needs:

- **From the Restore page:** Use this when you need to restore a resource, but you don't remember the exact name or the location in which it resides, or the date when it was last in good shape. You can search for the snapshot using filters.
- **From the Inventory page:** Use this when you need to restore a specific resource from the last week or month — and you know the name and location of the resource, and the date when it was last in good shape. You browse through a list of resources to find the one you want to restore.

## Required BlueXP role

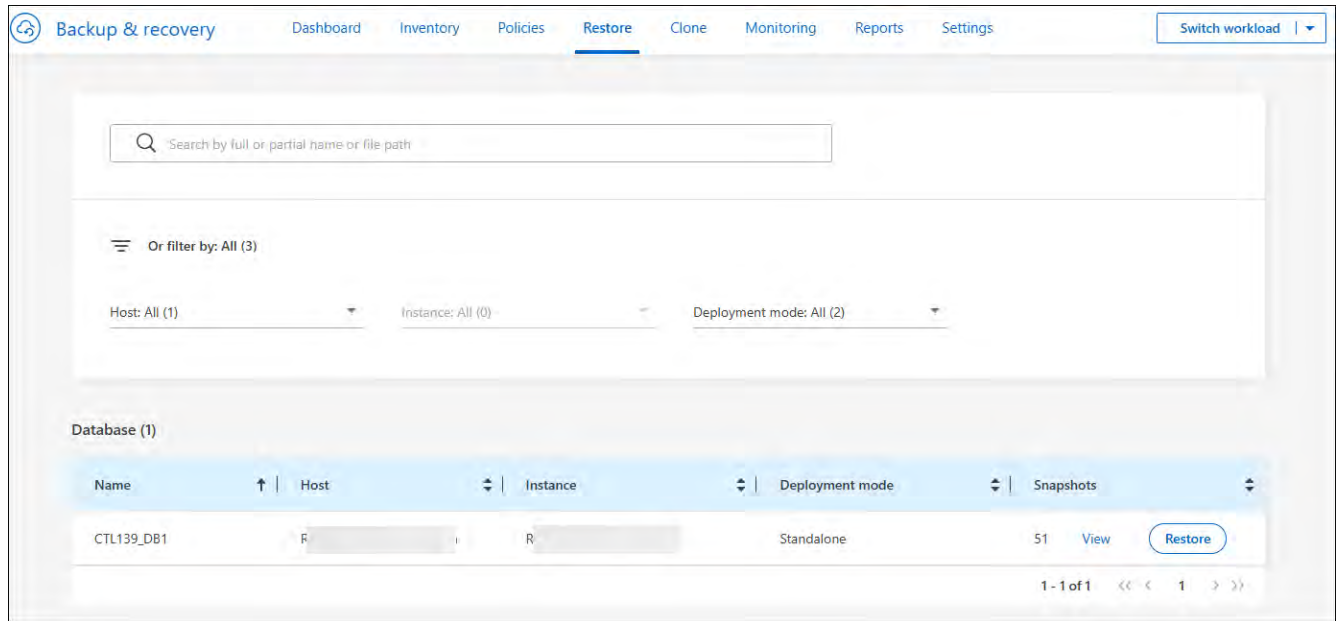
Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery restore admin role. [Learn about BlueXP access roles for all services.](#)

## Restore workload data from the Restore option

Restore database workloads using the Restore option.

## Steps

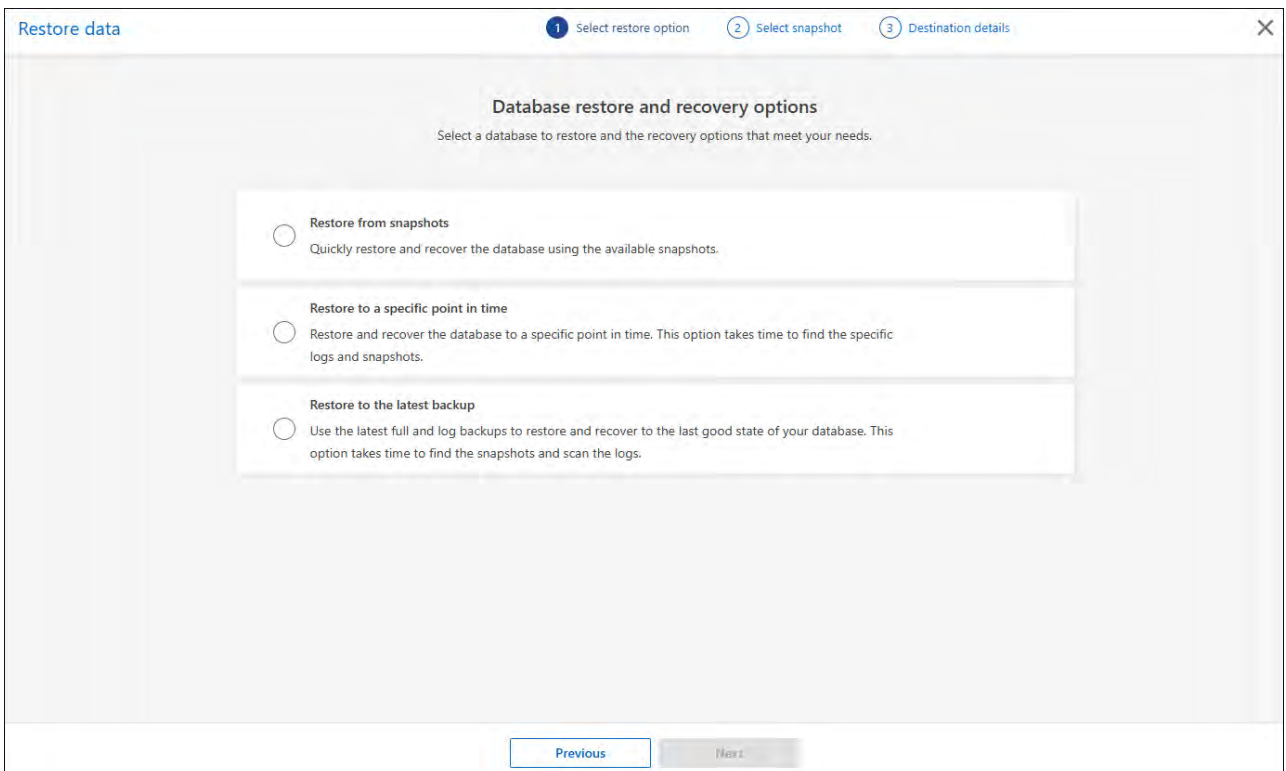
1. From the BlueXP backup and restore menu, select **Restore**.



2. Select the database that you want to restore. Use the filters to search.

3. Select the restore option:

- Restore from snapshots
- Restore to a specific point in time. This is helpful if you know the name and location of the file, and the date when it was last in good shape.
- Restore to the latest backup

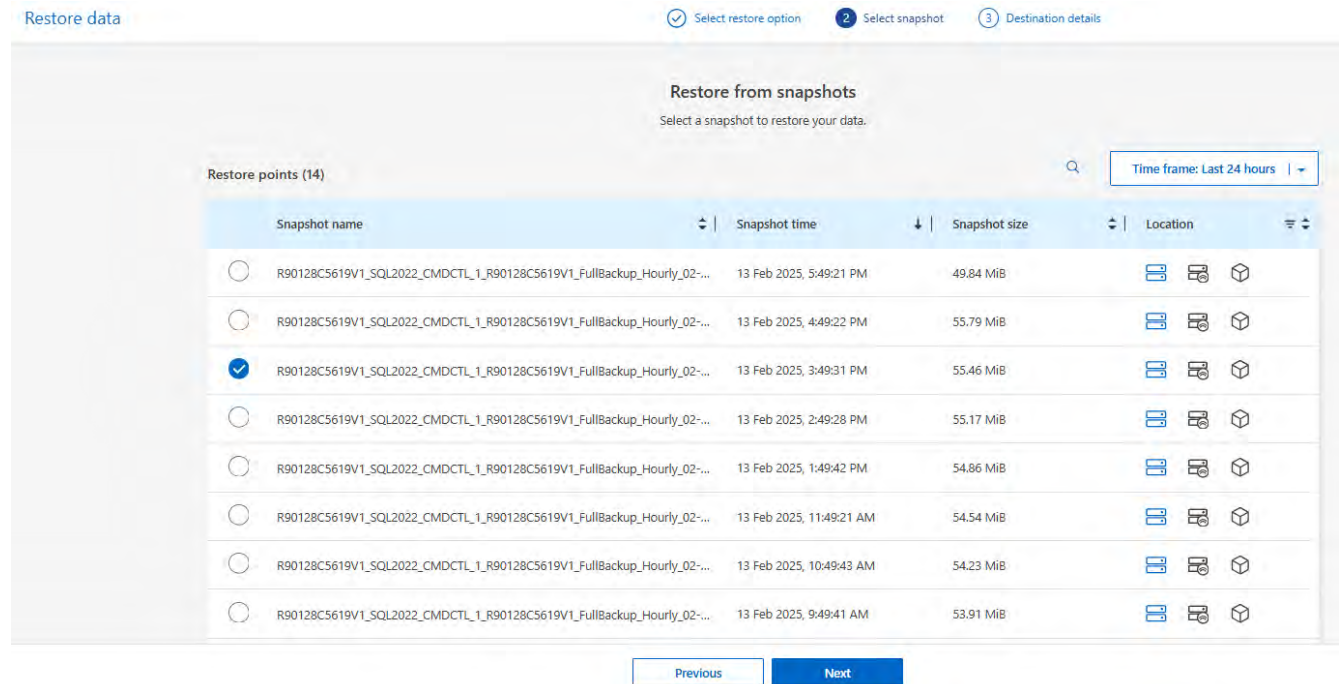




## Restore workloads from snapshots

1. Continuing from the Restore options page, select **Restore from snapshots**.

A list of snapshots appears.



The screenshot shows the 'Restore from snapshots' interface. At the top, there are three steps: '1 Select restore option', '2 Select snapshot', and '3 Destination details'. The current step is '2 Select snapshot'. Below the title, there is a search bar and a 'Time frame: Last 24 hours' dropdown. The main content is a table of 14 snapshots. The third snapshot is selected, indicated by a blue checkmark in a circle next to its name.

Snapshot name	Snapshot time	Snapshot size	Location
<input type="radio"/> R90128C5619V1_SQL2022_CMDCTL_1_R90128C5619V1_FullBackup_Hourly_02-...	13 Feb 2025, 5:49:21 PM	49.84 MiB	
<input type="radio"/> R90128C5619V1_SQL2022_CMDCTL_1_R90128C5619V1_FullBackup_Hourly_02-...	13 Feb 2025, 4:49:22 PM	55.79 MiB	
<input checked="" type="radio"/> R90128C5619V1_SQL2022_CMDCTL_1_R90128C5619V1_FullBackup_Hourly_02-...	13 Feb 2025, 3:49:31 PM	55.46 MiB	
<input type="radio"/> R90128C5619V1_SQL2022_CMDCTL_1_R90128C5619V1_FullBackup_Hourly_02-...	13 Feb 2025, 2:49:28 PM	55.17 MiB	
<input type="radio"/> R90128C5619V1_SQL2022_CMDCTL_1_R90128C5619V1_FullBackup_Hourly_02-...	13 Feb 2025, 1:49:42 PM	54.86 MiB	
<input type="radio"/> R90128C5619V1_SQL2022_CMDCTL_1_R90128C5619V1_FullBackup_Hourly_02-...	13 Feb 2025, 11:49:21 AM	54.54 MiB	
<input type="radio"/> R90128C5619V1_SQL2022_CMDCTL_1_R90128C5619V1_FullBackup_Hourly_02-...	13 Feb 2025, 10:49:43 AM	54.23 MiB	
<input type="radio"/> R90128C5619V1_SQL2022_CMDCTL_1_R90128C5619V1_FullBackup_Hourly_02-...	13 Feb 2025, 9:49:41 AM	53.91 MiB	

At the bottom of the table, there are two buttons: 'Previous' and 'Next'.

2. Select the snapshot you want to restore.
3. Select **Next**.

You'll see destination options next.


Restore data ✓ Select restore option   ✓ Select snapshot   3 Destination details

### Choose destination settings

Choose a recovery destination and operation speed for your data recovery.

Destination settings ^

Destination  Original location    Alternate location

 MDML\_DB1

---

R91115E55FFV1.hnk4.com                      R91115E55FFV1\SQL2022  
Host    Instance

Pre-restore options                      No action required v

Post-restore options                      No action required v

Previous
Next

4. In the Destination details page, enter the following information:

- **Destination settings:** Choose whether you want to restore the data to its original location or to an alternate location. For an alternate location, select the host name and instance, enter the database name, and enter the destination path where you want to restore the snapshot.
- **Pre-restore options:**
  - **Overwrite the database with the same name during restore:** During the restore, the original database name is preserved.
  - **Retain SQL database replication settings:** Keeps the replication settings for the SQL database after the restore operation.
  - **Create transaction log backup before restore:** Creates a transaction log backup before the restore operation.\* **Quit restore if transaction log backup before restore fails:** Stops the restore operation if the transaction log backup fails.
  - **Prescript:** Enter the full path for a script that should be run before the restore operation, any arguments that the script takes, and how long to wait for the script to complete.
- **Post-restore options:**
  - **Operational,** but unavailable for restoring additional transaction logs. This brings the database back online after transaction log backups are applied.
  - **Non-operational,** but available for restoring additional transaction logs. Maintains the database in a non-operational state after the restore operation while restoring transaction log backups. This option is useful for restoring additional transaction logs.

- **Read-only mode** and available for restoring additional transaction logs. Restores the database in a read-only mode and applies transaction log backups.
- **Postscript:** Enter the full path for a script that should be run after the restore operation and any arguments that the script takes.

5. Select **Restore**.

### Restore to specific point in time

BlueXP backup and recovery uses logs and the most recent snapshots to create a point-in-time restore of your data.

1. Continuing from the Restore options page, select **Restore to specific point in time**.
2. Select **Next**.

3. In the Restore to a specific point in time page, enter the following information:
  - **Date and time for data restoration:** Enter the exact date and time of the data that you want to restore. This date and time is from the Microsoft SQL Server Database host.
4. Select **Search**.
5. Select the snapshot that you want to restore.
6. Select **Next**.
7. In the Destination details page, enter the following information:
  - **Destination settings:** Choose whether you want to restore the data to its original location or to an alternate location. For an alternate location, select the host name and instance, enter the database name, and enter the destination path.
  - **Pre-restore options:**
    - **Preserve original database name:** During the restore, the original database name is preserved.
    - **Retain SQL database replication settings:** Keeps the replication settings for the SQL database after the restore operation.
    - **Prescript:** Enter the full path for a script that should be run before the restore operation, any arguments that the script takes, and how long to wait for the script to complete.
  - **Post-restore options:**
    - **Operational,** but unavailable for restoring additional transaction logs. This brings the database back online after transaction log backups are applied.

- **Non-operational**, but available for restoring additional transaction logs. Maintains the database in a non-operational state after the restore operation while restoring transaction log backups. This option is useful for restoring additional transaction logs.
- **Read-only mode** and available for restoring additional transaction logs. Restores the database in a read-only mode and applies transaction log backups.
- **Postscript**: Enter the full path for a script that should be run after the restore operation and any arguments that the script takes.

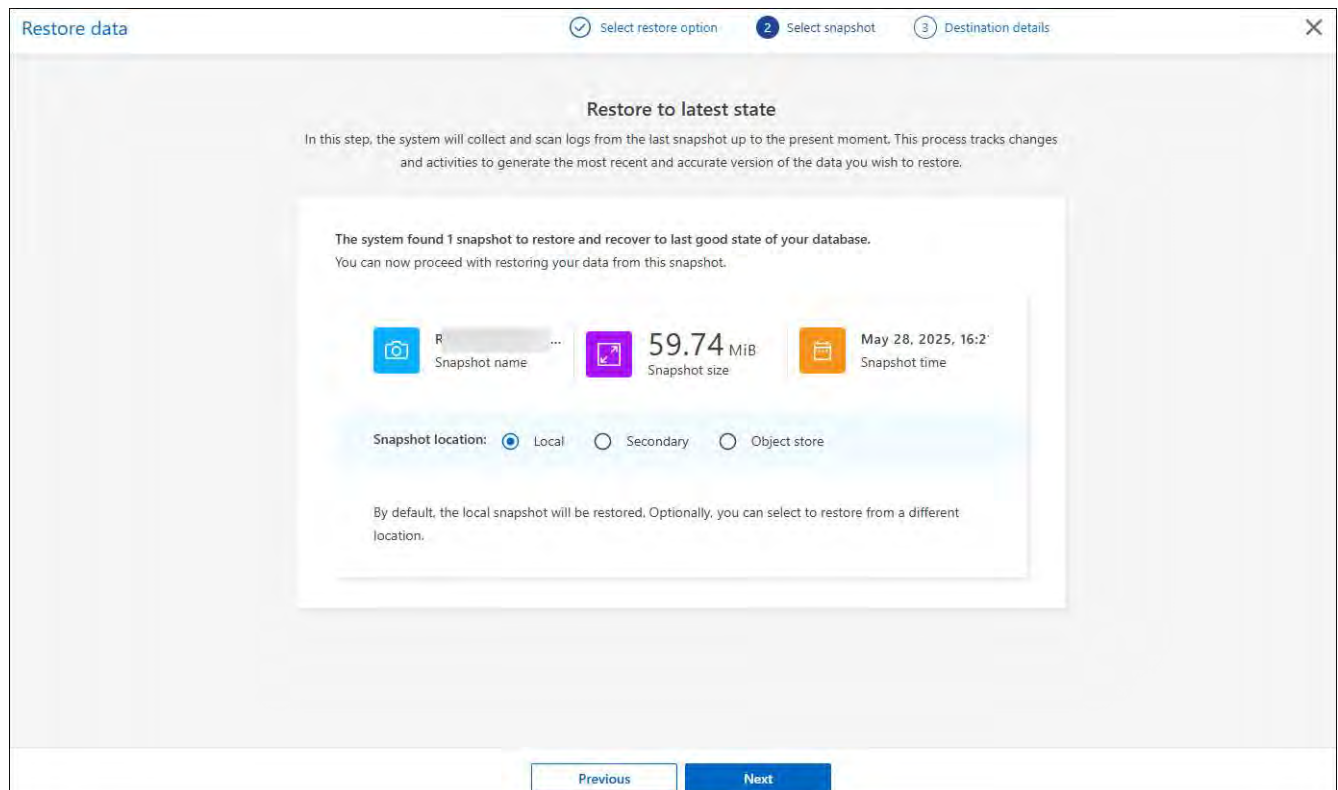
8. Select **Restore**.

### Restore to the latest backup

This option uses the latest full and log backups to restore your data to the last good state. The system scans logs from the last snapshot to the present. The process tracks changes and activities to restore the most recent and accurate version of your data.

1. Continuing from the Restore options page, select **Restore to the latest backup**.

BlueXP backup and recovery shows you the snapshots that are available for the restore operation.



2. In the Restore to the latest state page, select the snapshot location of local, secondary storage, or object storage.

3. Select **Next**.

4. In the Destination details page, enter the following information:

- **Destination settings**: Choose whether you want to restore the data to its original location or to an alternate location. For an alternate location, select the host name and instance, enter the database name, and enter the destination path.
- **Pre-restore options**:

- **Overwrite the database with the same name during restore:** During the restore, the original database name is preserved.
- **Retain SQL database replication settings:** Keeps the replication settings for the SQL database after the restore operation.
- **Create transaction log backup before restore:** Creates a transaction log backup before the restore operation.
- **Quit restore if transaction log backup before restore fails:** Stops the restore operation if the transaction log backup fails.
- **Prescript:** Enter the full path for a script that should be run before the restore operation, any arguments that the script takes, and how long to wait for the script to complete.
- **Post-restore options:**
  - **Operational,** but unavailable for restoring additional transaction logs. This brings the database back online after transaction log backups are applied.
  - **Non-operational,** but available for restoring additional transaction logs. Maintains the database in a non-operational state after the restore operation while restoring transaction log backups. This option is useful for restoring additional transaction logs.
  - **Read-only mode** and available for restoring additional transaction logs. Restores the database in a read-only mode and applies transaction log backups.
  - **Postscript:** Enter the full path for a script that should be run after the restore operation and any arguments that the script takes.


5. Select **Restore**.

### Restore workload data from the Inventory option

Restore database workloads from the Inventory page.

Using the Inventory option, you can restore only databases, not instances.

#### Steps

1. From the BlueXP backup and restore menu, select **Inventory**.
2. Choose the host where the resource that you want to restore is located.
3. Select the **Actions**  icon, and select **View details**.
4. On the Microsoft SQL Server page, select the **Databases** tab.
5. On the Databases tab, select the database that shows a "Protected" status indicating that there's a backup that you can restore.

Backup & recovery | Dashboard | **Inventory** | Policies | Restore | Monitoring | Settings

### Microsoft SQL Server

Review protection status, create and manage policies, and recover data for one or more databases.

1 Hosts

1 Instances

9 Databases

1 Protected resources

80 MiB Protected capacity

Hosts | Instances | **Databases**

Databases (9)

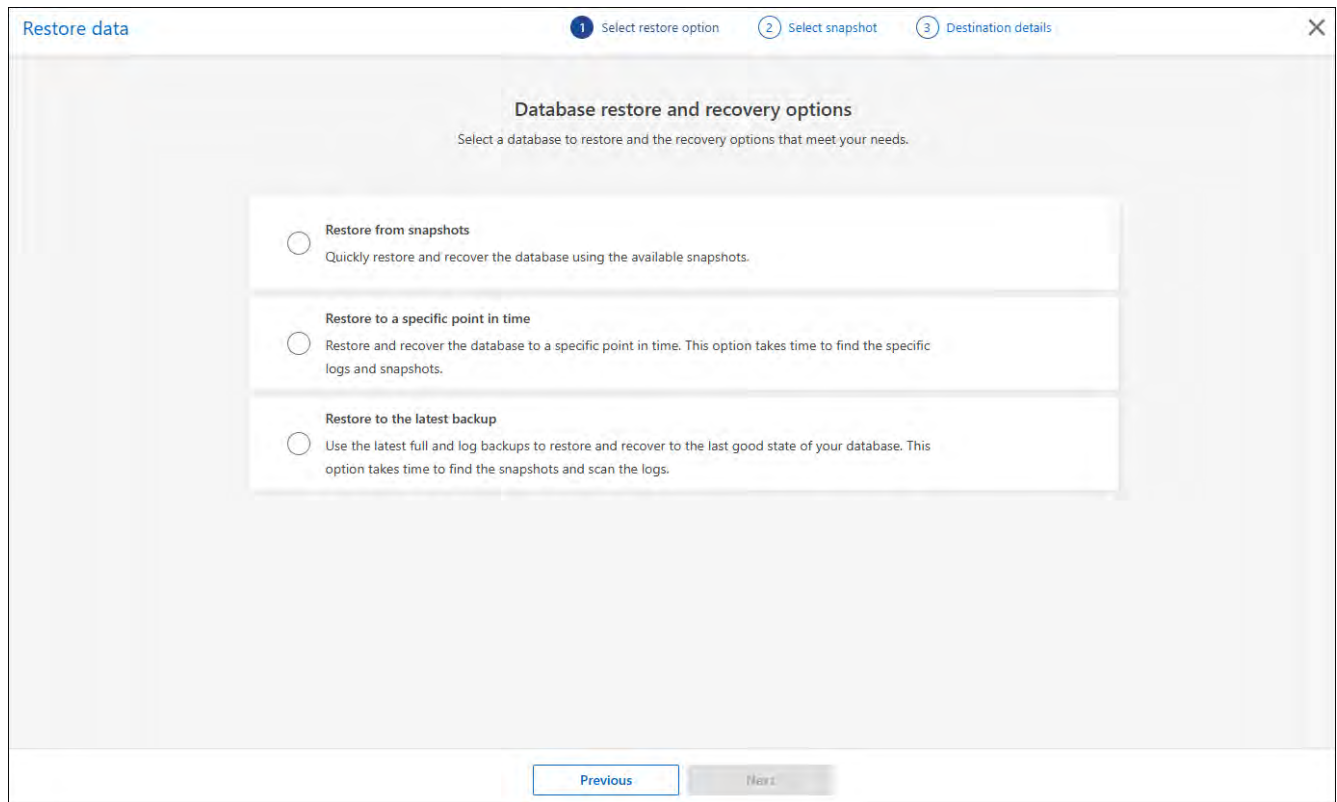
<input type="checkbox"/>	Database name ↑	Protection status ↓	Assigned host ↓	Assigned instance ↓	Storage type ↓	Capacity ↓	Policy ↓	
<input type="checkbox"/>	CMDCTL_1	Protected	R90128C5619V1.hnk...	R90128C5619V1\SQL2022	On-Premises ONTAP	80 MiB	HourlyDailyPolicy	
<input type="checkbox"/>	master	Not available for backup	R90128C5619V1.hnk...	R90128C5619V1\SQL2022	On-Premises ONTAP	5.25 MiB	Protect	
<input type="checkbox"/>	MDML_DB1	Unprotected	R90128C5619V1.hnk...	R90128C5619V1\SQL2022	On-Premises ONTAP	165.94 MiB	Restore	
<input type="checkbox"/>	MDML_DB2	Unprotected	R90128C5619V1.hnk...	R90128C5619V1\SQL2022	On-Premises ONTAP	165.94 MiB	View protection details	
<input type="checkbox"/>	MDSL_DB3	Unprotected	R90128C5619V1.hnk...	R90128C5619V1\SQL2022	On-Premises ONTAP	165.94 MiB	Edit protection	
<input type="checkbox"/>	MDSL_DB4	Unprotected	R90128C5619V1.hnk...	R90128C5619V1\SQL2022	On-Premises ONTAP	165.94 MiB	Backup now	

6. Select the **Actions** icon, and select **Restore**.

The same three options appear as when you restore from the Restore page:

- Restore from snapshots
- Restore to a specific point in time
- Restore to the latest backup

7. Continue with the same steps for the restore option from the Restore page



## Clone Microsoft SQL Server workloads with BlueXP backup and recovery

Clone Microsoft SQL Server applications data to the same or different VM for development, testing, or protection purposes using BlueXP backup and recovery. You can create clones from instant snapshots or existing snapshots of your Microsoft SQL Server workloads.

Choose between the following types of clones:

- **Instant snapshot and clone:** You can create a clone of your Microsoft SQL Server workloads from an instant snapshot. An instant snapshot is a point-in-time copy of the source data that is created from a backup. The clone is stored in an object store in your public or private cloud account. You can use the clone to restore your workloads in case of data loss or corruption.
- **Clone from an existing snapshot:** You can choose an existing snapshot from a list of snapshots that are available for the workload. This option is useful if you want to create a clone from a specific point in time. Clone to either primary or secondary storage.

You can accomplish the following protection goals:

- Create a clone
- Refresh a clone
- Split a clone
- Delete a clone

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, or Backup and Recovery clone admin role. [Learn about BlueXP access roles for all services.](#)



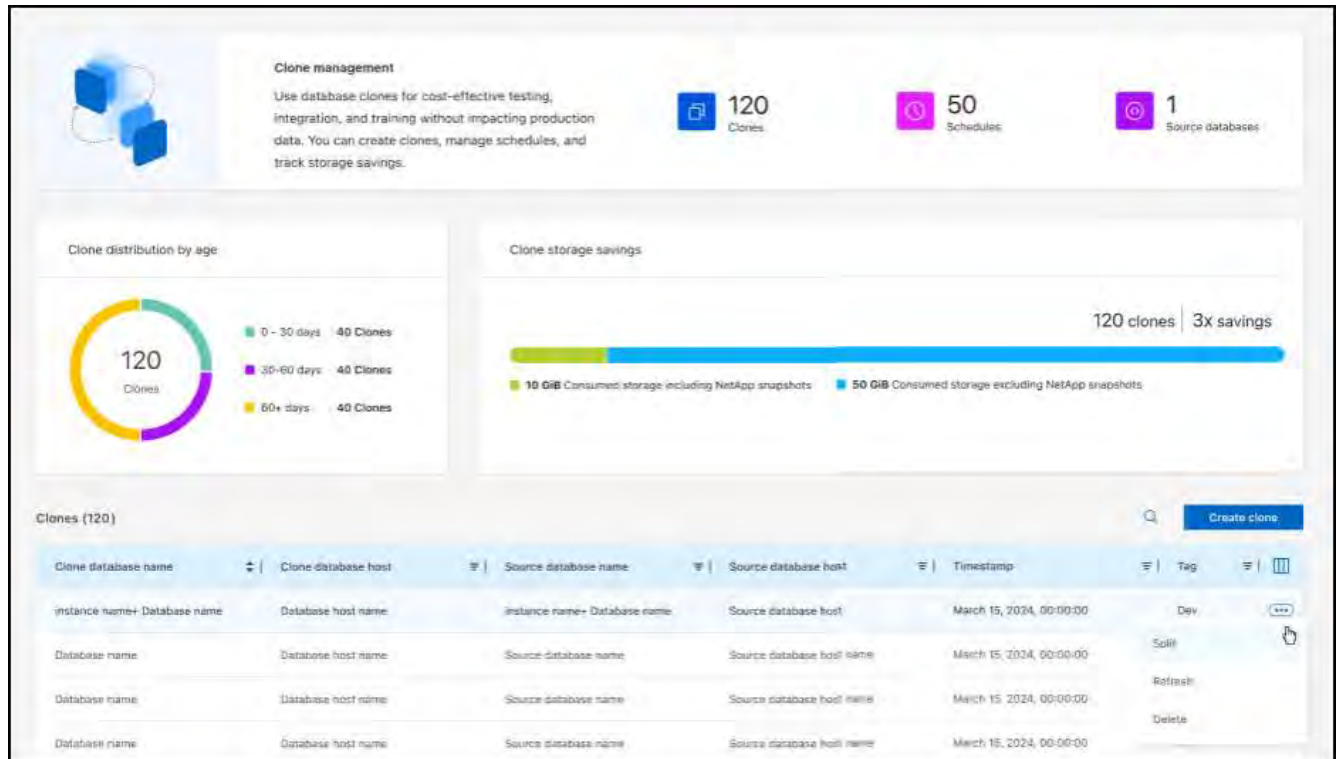
## Create a clone

You can create a clone of your Microsoft SQL Server workloads. A clone is a copy of the source data that is created from a backup. The clone is stored in an object store in your public or private cloud account. You can use the clone to restore your workloads in case of data loss or corruption.

You can create a clone from an existing snapshot or from an instant snapshot. An instant snapshot is a point-in-time copy of the source data that is created from a backup. You can use the clone to restore your workloads in case of data loss or corruption.

### Steps

1. From the BlueXP backup and recovery menu, select **Clone**.



2. Select **Create new clone**.
3. Select the clone type:
  - **Clone and database refresh from existing snapshot:** Choose the snapshot for the clone and configure options for the clone. This is helpful if you want to choose the snapshot for the clone and configure options.
  - **Instant snapshot and clone:** Take a snapshot now of the source data and create a clone from that snapshot. This option is useful if you want to create a clone from the latest data in the source workload.
4. Complete the **Database source** section:
  - **Single clone or bulk clone:** Select whether to create a single clone or multiple clones. If you select **Bulk clone**, you can create multiple clones at once using a protection group that you already created. This option is useful if you want to create multiple clones for different workloads.
  - **Source database host, instance, and name:** Select the source database host, instance, and name for the clone. The source database is the database from which the clone will be created.
5. Complete the **Database target** section:



- **Target database host, instance, and name:** Select the target database host, instance, and name for the clone. The target database is the location where the clone will be created.

Optionally, select **Suffix** from the target name drop-down list and append a suffix to the cloned database name. If you do not specify a suffix, the cloned database name will be the same as the source database name.

- **QoS (max throughput):** Select the quality of service (QoS) maximum throughput in MBps for the clone. The QoS defines the performance characteristics of the clone, such as the maximum throughput and IOPS.

6. Complete the **Mount** section:

- **Auto-assign mount point:** Select this option to automatically assign a mount point for the clone. The mount point is the location where the clone will be mounted in the object store.
- **Define mount point path:** Enter a mount point for the clone. The mount point is the location where the clone will be mounted in the object store. Select the drive letter, enter the data file path, and enter the log file path.

7. Select **Next**.

8. Select the restore point:

- **Existing snapshots:** Select an existing snapshot from the list of snapshots that are available for the workload. This option is useful if you want to create a clone from a specific point in time.
- **Instant snapshot and clone:** Select the latest snapshot from the list of snapshots that are available for the workload. This option is useful if you want to create a clone from the latest data in the source workload.

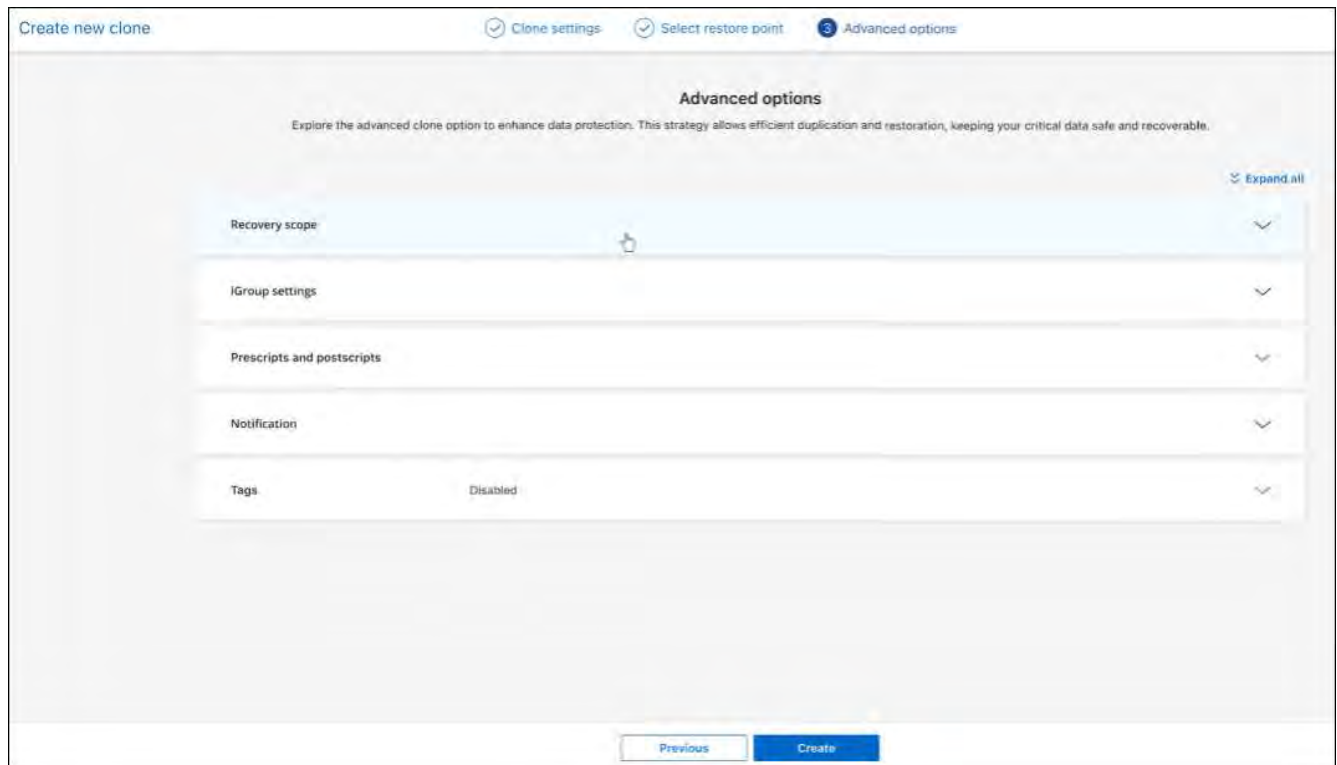
9. If you chose to create **Instant snapshot and clone**, choose the clone storage location:

- **Local storage:** Select this option to create the clone in the local storage of the ONTAP system. The local storage is the storage that is directly attached to the ONTAP system.
- **Secondary storage:** Select this option to create the clone in the secondary storage of the ONTAP system. The secondary storage is the storage that is used for backup and recovery workloads.

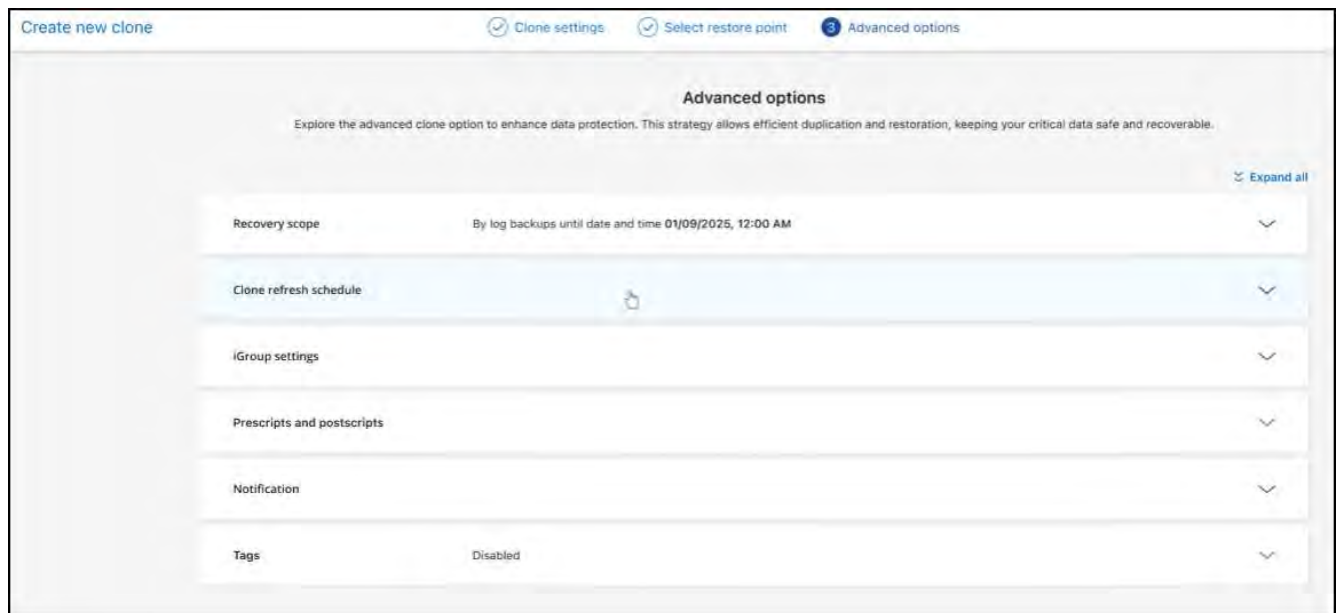
10. Select the destination location for the data and logs.

11. Select **Next**.

12. Complete the **Advanced options** section:



13. If you chose **Instant snapshot and clone**, complete the following options:

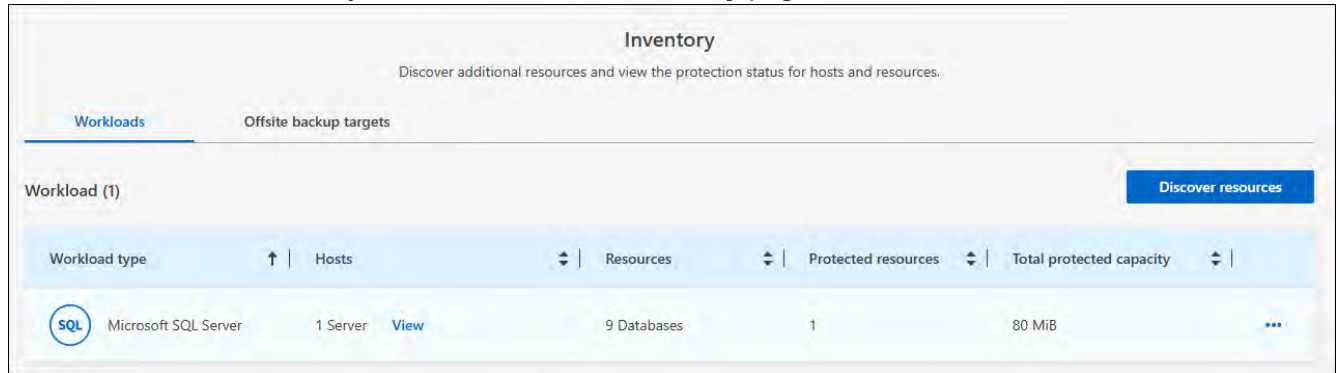


- **Clone refresh schedule and expiration:** If you chose **Instant clone**, enter the date when to begin refreshing the clone. The clone schedule defines when the clone will be created.
  - **Delete clone if schedule expires:** If you want to delete the clone upon the clone expiration date.
  - **Refresh clone every:** Select how often the clone should be refreshed. You can choose to refresh the clone hourly, daily, weekly, monthly, or quarterly. This option is useful if you want to keep the clone up to date with the source workload.
- **Prescripts and postscripts:** Optionally, specify pre- and post-clone scripts to run before and after the clone is created. These scripts can be used to perform additional tasks, such as configuring the clone or sending notifications.

- **Notification:** Optionally, specify email addresses to receive notifications about the clone creation status along with the Job report. You can also specify a webhook URL to receive notifications about the clone creation status. You can specify whether you want success and failure notifications or only one or the other.
- **Tags:** Select one or more labels that will help you later search for the resource group and select **Apply**. For example, if you add "HR" as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

14. Select **Create**.

15. When the clone is created, you can view it in the **Inventory** page.



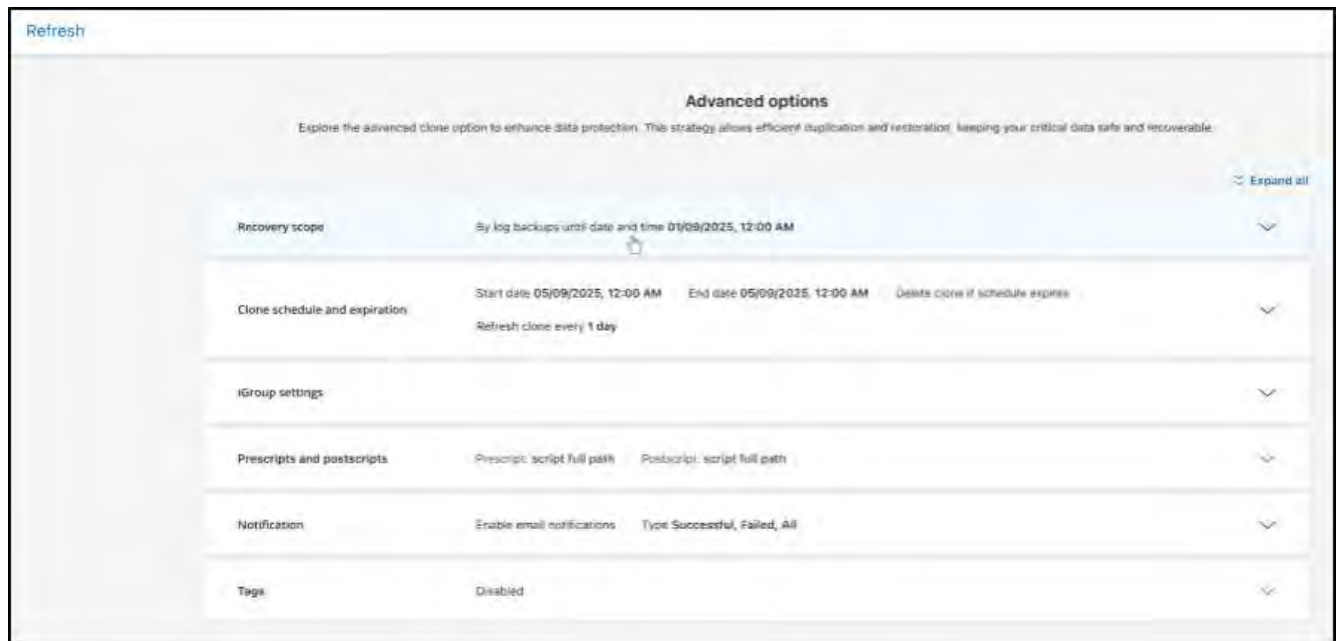
## Refresh a clone

You can refresh a clone of your Microsoft SQL Server workloads. Refreshing a clone updates the clone with the latest data from the source workload. This is useful if you want to keep the clone up to date with the source workload.

You have the option to change the database name, use the latest instant snapshot, or refresh from an existing production snapshot.

### Steps

1. From the BlueXP backup and recovery menu, select **Clone**.
2. Select the clone you want to refresh.
3. Select the Actions icon **...** > **Refresh clone**.



#### 4. Complete the **Advanced settings** section:

- **Recovery scope:** Choose whether to recover all log backups or log backups until a specific point in time. This option is useful if you want to recover the clone to a specific point in time.
- **Clone refresh schedule and expiration:** If you chose **Instant clone**, enter the date when to begin refreshing the clone. The clone schedule defines when the clone will be created.
  - **Delete clone if schedule expires:** If you want to delete the clone upon the clone expiration date.
  - **Refresh clone every:** Select how often the clone should be refreshed. You can choose to refresh the clone hourly, daily, weekly, monthly, or quarterly. This option is useful if you want to keep the clone up to date with the source workload.
- **iGroup settings:** Select the igroup for the clone. The igroup is a logical grouping of initiators that are used to access the clone. You can select an existing igroup or create a new one. Select the igroup from the primary or secondary ONTAP storage system.
- **Prescripts and postscripts:** Optionally, specify pre- and post-clone scripts to run before and after the clone is created. These scripts can be used to perform additional tasks, such as configuring the clone or sending notifications.
- **Notification:** Optionally, specify email addresses to receive notifications about the clone creation status along with the Job report. You can also specify a webhook URL to receive notifications about the clone creation status. You can specify whether you want success and failure notifications or only one or the other.
- **Tags:** Enter one or more labels that will help you later search for the resource group. For example, if you add "HR" as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.

5. In the Refresh confirmation dialog box, to continue, select **Refresh**.

### Skip a clone refresh

You might want to skip a clone refresh if you do not want to update the clone with the latest data from the source workload. Skipping a clone refresh allows you to keep the clone as it is without updating it.

### Steps

1. From the BlueXP backup and recovery menu, select **Clone**.
2. Select the clone you want to skip the refresh for.
3. Select the Actions icon **...** > **Skip refresh**.
4. In the Skip refresh confirmation dialog box, do the following:
  - a. To skip only the next refresh schedule, select **Only skip the next refresh schedule**.
  - b. To continue, select **Skip**.

## Split a clone

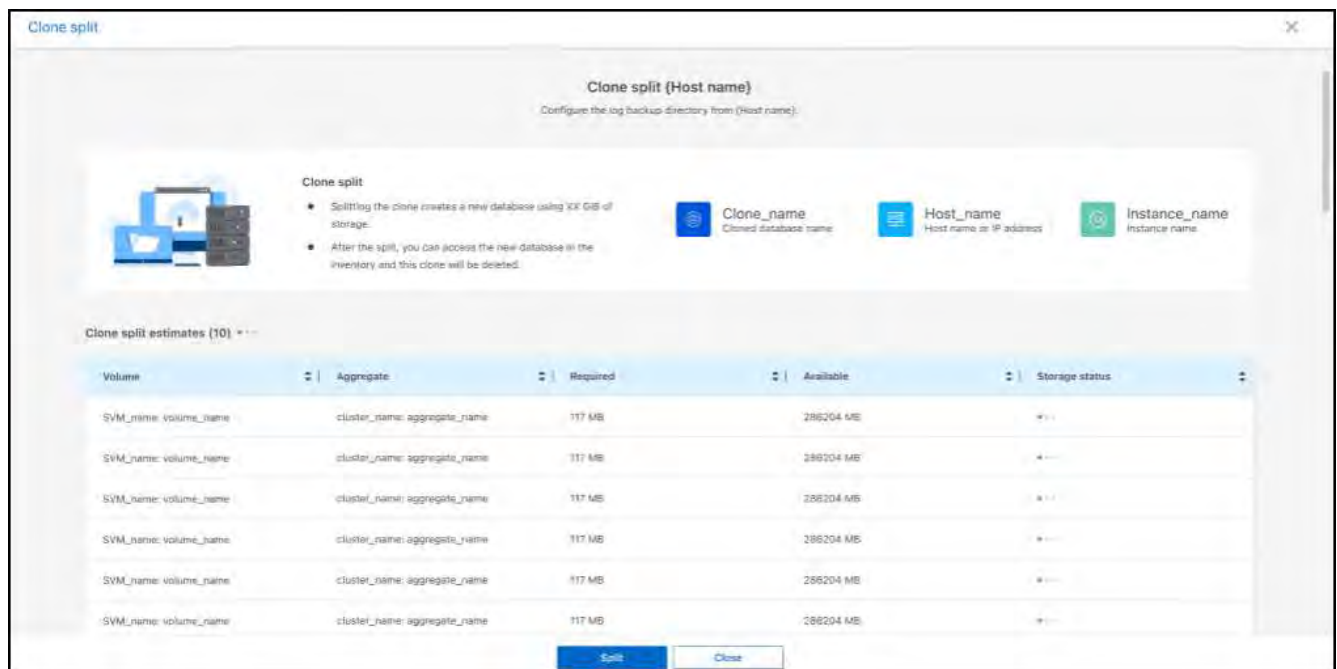
You can split a clone of your Microsoft SQL Server workloads. Splitting a clone creates a new backup from the clone. The new backup can be used to restore the workloads.

You can choose to split a clone as independent or long-term clones. A wizard shows the list of aggregates that are part of the SVM, their sizes, and where the cloned volume resides. BlueXP backup and recovery also indicates whether there is enough space to split the clone. After the clone is split, the clone becomes an independent database for protection.

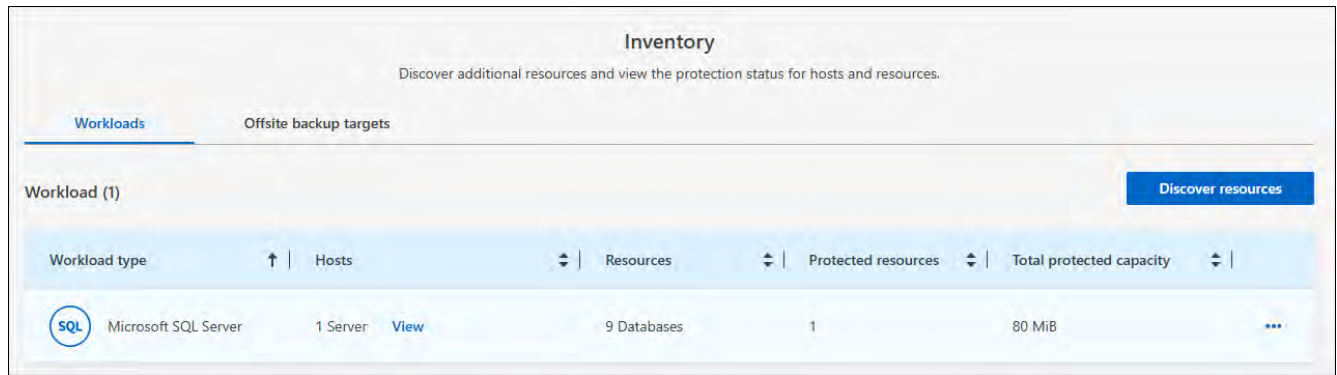
The clone job is not be removed and it can be reused again for other clones.

## Steps

1. From the BlueXP backup and recovery menu, select **Clone**.
2. Select a clone.
3. Select the Actions icon **...** > **Split clone**.



4. Review the split clone details and select **Split**.
5. When the split clone is created, you can view it in the **Inventory** page.



## Delete a clone

You can delete a clone of your Microsoft SQL Server workloads. Deleting a clone removes the clone from the object store and frees up storage space.

If the clone is protected by a policy, the clone is deleted including the job.

### Steps

1. From the BlueXP backup and recovery menu, select **Clone**.
2. Select a clone.
3. Select the Actions icon **...** > **Delete clone**.
4. In the clone Delete confirmation dialog box, review the deletion details.
  - a. To delete the cloned resources from SnapCenter even if the clones or their storage is not accessible, select **Force delete**.
  - b. Select **Delete**.
5. When the clone is deleted, it is removed from the **Inventory** page.

## Manage Microsoft SQL Server inventory with BlueXP backup and recovery

BlueXP backup and recovery enables you to manage your Microsoft SQL Server workload host information, database information, and instances information. You can view, edit, and delete protection settings of your inventory.

You can accomplish the following tasks related to managing your inventory:

- Manage host information
  - Suspend schedules
  - Edit or delete hosts
- Manage instances information
  - Associate credentials with a resource
  - Back up now by starting an on-demand backup
  - Edit protection settings
- Manage database information
  - Protect databases

- Restore databases
- Edit protection settings
- Back up now by starting an on-demand backup
- Configure the log directory (from Inventory > Hosts). If you want to back up logs for your database hosts in the snapshot, first configure the logs in BlueXP backup and recovery. For details, refer to [Configure BlueXP backup and recovery settings](#).

## Manage host information

You can manage host information to ensure that the right hosts are protected. You can view, edit, and delete host information.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, or Backup and Recovery clone admin role. [Learn about BlueXP access roles for all services](#).

- Configure log directory. For details, refer to [Configure BlueXP backup and recovery settings](#).
- Suspend schedules
- Edit a host
- Delete a host

## Manage hosts

You can manage the hosts that are discovered in your working environment. You can manage them separately or as a group.



You can manage only those hosts that show an "Unmanaged" status in the Hosts column. If the status is "Managed", it means that the host is already being managed by BlueXP backup and recovery.


After you manage the hosts in BlueXP backup and recovery, SnapCenter no longer manages the resources on those hosts.

### Required BlueXP role

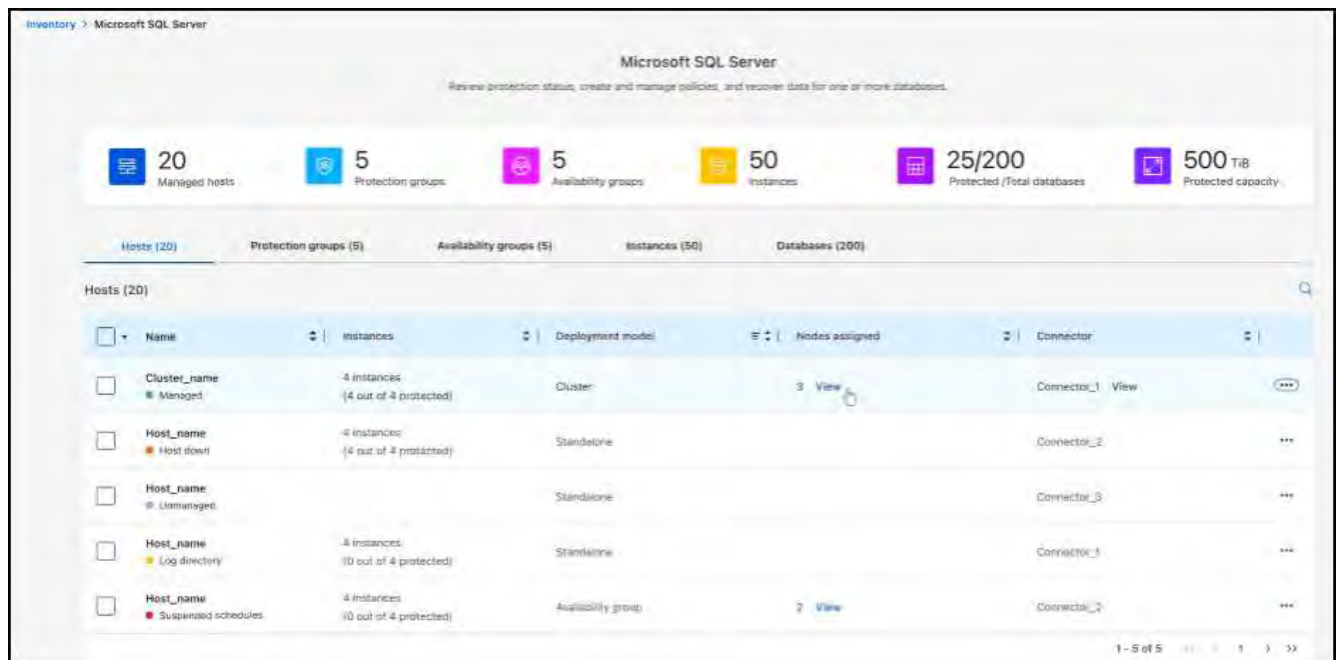
Organization admin, Folder or project admin, or Backup and Recovery super admin. [Learn about BlueXP access roles for all services](#).

## Steps

1. From the menu, select **Inventory**.

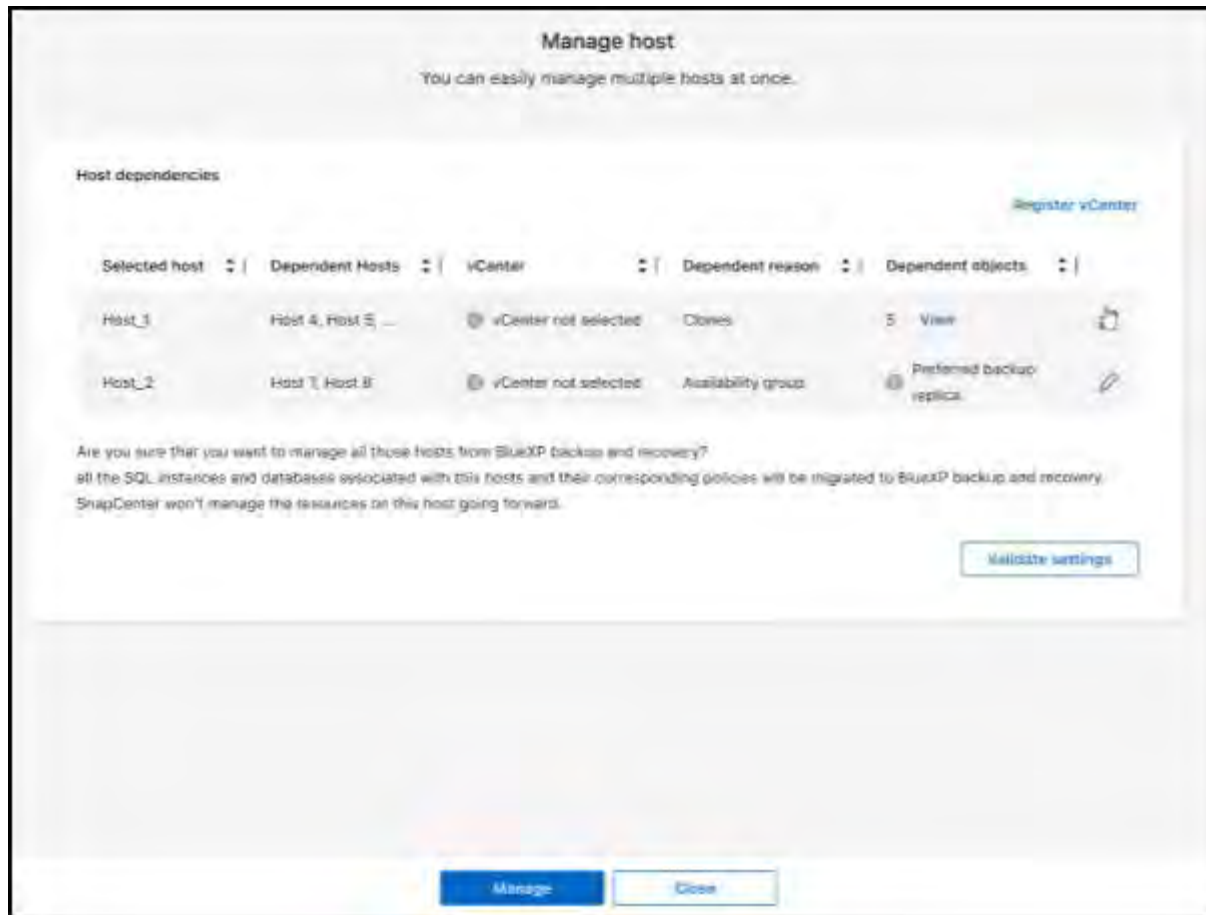


2. Select a workload to view the protection details.
3. Select the Actions icon **...** > **View details**.



4. Select the **Hosts** tab.
5. Select one or more hosts. If you select multiple hosts, a Bulk actions option appears where you can select **Manage (up to 5 hosts)**.
6. Select the Actions icon **...** > **Manage**.





#### 7. Review the host dependencies:

- If the vCenter does not display, select the pencil icon to add or edit the vCenter details.
- If you add a vCenter, you must also register the vCenter by selecting **Register vCenter**.

#### 8. Select **Validate settings** to test your settings.

#### 9. Select **Manage** to manage the host.

### Suspend schedules

You can suspend schedules to stop the backup and restore operations for a host. You might want to do this if you need to perform maintenance activities on the host.

#### Steps


1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the host on which you want to suspend schedules.
3. Select the **Actions** **...** icon, and select **Suspend schedules**.
4. In the confirmation dialog box, select **Suspend**.

### Edit a host

You can change the vCenter server information, host registration credentials, and advanced settings options.

#### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.


2. Select the host that you want to edit.
3. Select the **Actions** , and select **Edit host**.

4. Edit the host information.
5. Select **Done**.

### Delete a host

You can delete the host information to stop service charges.

#### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the host that you want to delete.
3. Select the **Actions** , and select **Delete host**.
4. Review the confirmation information and select **Delete**.

### Manage instances information

You can manage instances information to ensure that resources have the appropriate credentials for protection and you can back up resources in the following ways:

- Protect instances
- Associate credentials
- Disassociate credentials

- Edit protection
- Back up now


### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, or Backup and Recovery clone admin role. [Learn about BlueXP access roles for all services.](#)

### Protect database instances

You can assign a policy to a database instance using policies that govern the schedules and retention of resource protection.

### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** tab.
4. Select the instance.
5. Select the **Actions**  icon, and select **Protect**.
6. Select a policy or create a new one.

For details about creating a policy, refer to [Create a policy](#).

7. Provide information on the scripts that you want to run before and after the backup.
  - **Pre-script:** Enter your script filename and location to run it automatically before the protect action is triggered. This is helpful for performing additional tasks or configurations that need to be executed before the protection workflow.
  - **Post-script:** Enter your script filename and location to run it automatically after the protection action is complete. This is helpful for performing additional tasks or configurations that need to be executed after the protection workflow.
8. Provide information on how you want the snapshot to be verified:
  - **Storage location:** Select the location where the verification snapshot will be stored.
  - **Verification resource:** Select whether the resource that you want to verify is on the local snapshot and on ONTAP secondary storage.
  - **Verification schedule:** Select the frequency of hourly, daily, weekly, monthly, or yearly.


### Associate credentials with a resource

You can associate credentials with a resource so that protection can occur.

For details, see [Configure BlueXP backup and recovery settings, including credentials](#).

### Steps


1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** tab.
4. Select the instance.

5. Select the **Actions**  icon, and select **Associate credentials**.
6. Use existing credentials or create new ones.

### Edit protection settings

You can change the policy, create a new policy, set a schedule, and set retention settings.

### Steps


1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** tab.
4. Select the instance.
5. Select the **Actions**  icon, and select **Edit protection**.

For details about creating a policy, refer to [Create a policy](#).

### Back up now

You can back up your data now to ensure that your data is protected immediately.

### Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** tab.
4. Select the instance.
5. Select the **Actions**  icon, and select **Back up now**.
6. Choose the backup type and set the schedule.

For details about creating an ad hoc backup, refer to [Create a policy](#).

### Manage database information

You can manage database information in the following ways:

- Protect databases
- Restore databases
- View protection details
- Edit protection settings
- Back up now


### Protect databases

You can change the policy, create a new policy, set a schedule, and set retention settings.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about BlueXP access roles for all services](#).

## Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Databases** tab.
4. Select the database.
5. Select the **Actions**  icon, and select **Protect**.

For details about creating a policy, refer to [Create a policy](#).


## Restore databases

You can restore a database to ensure that your data is protected.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery restore admin role. [Learn about BlueXP access roles for all services](#).

## Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Databases** tab.
4. Select the database.
5. Select the **Actions**  icon, and select **Restore**.

For information about restoring workloads, refer to [Restore workloads](#).


## Edit protection settings

You can change the policy, create a new policy, set a schedule, and set retention settings.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about BlueXP access roles for all services](#).

## Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Databases** tab.
4. Select the database.
5. Select the **Actions**  icon, and select **Edit protection**.

For details about creating a policy, refer to [Create a policy](#).

## Back up now

You can back up your Microsoft SQL Server instances and databases now to ensure that your data is protected immediately.

## Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about BlueXP access roles for all services.](#)

## Steps

1. From the BlueXP backup and recovery menu, select **Inventory**.
2. Select the workload that you want to view and select **View**.
3. Select the **Instances** or **Databases** tab.
4. Select the instance or database.
5. Select the **Actions** **...** icon, and select **Back up now**.

## Manage Microsoft SQL Server snapshots with BlueXP backup and recovery

You can manage Microsoft SQL Server snapshots by deleting them from BlueXP backup and recovery.

### Delete a snapshot

You can delete only local snapshots.

## Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin role. [Learn about BlueXP access roles for all services.](#)

## Steps

1. In BlueXP backup and recovery, select **Inventory**.
2. Select the workload and select **View**.
3. Select the **Databases** tab.
4. Select the database that you want to delete a snapshot for.
5. From the Actions menu, select **View protection details**.

The screenshot displays the 'View protection details' page for a Microsoft SQL Server snapshot. At the top, it shows the breadcrumb 'Inventory > SQL Workload > R90128C5619V1\SQL2022'. The main content area is divided into several sections:


- Metadata:** Includes 'CMDCTL\_1 Databases', 'R90 Instance', 'R90 Database host', 'ONPREM Location', 'Ransomware protection' (status: Ransomware protection), and 'Protection health' (status: Healthy).
- Diagram:** A 'Disk to disk' diagram showing data flow from 'ONTAP Primary' to 'ONTAP Secondary'.
- Policy information:** Lists 'Policy name: HourlyDailyPolicy', 'Local schedules: Hourly, Daily', 'Secondary schedules: Hourly, Daily', 'Object store schedules: Disabled', and 'Copy only backup: Disabled'.
- Recovery points (50 / 682):** A table listing snapshot details.

Snapshot name	Size	Recovery point	Location
R90128C56...	119.86 MiB	04 Mar 2025, 11:58 PM	[Icons]
R90128C56...	119.27 MiB	04 Mar 2025, 11:57 PM	[Icons] Delete local snapshot
R90128C51...	118.7 MiB	04 Mar 2025, 11:56 PM	[Icons]

6. Select the local snapshot that you want to delete.



The local snapshot icon in the **Location** column on that row must appear in blue.

7. Select the **Actions**  icon, and select **Delete local snapshot**.

8. In the confirmation dialog box, select **Remove**.

## Create reports for Microsoft SQL Server workloads in BlueXP backup and recovery

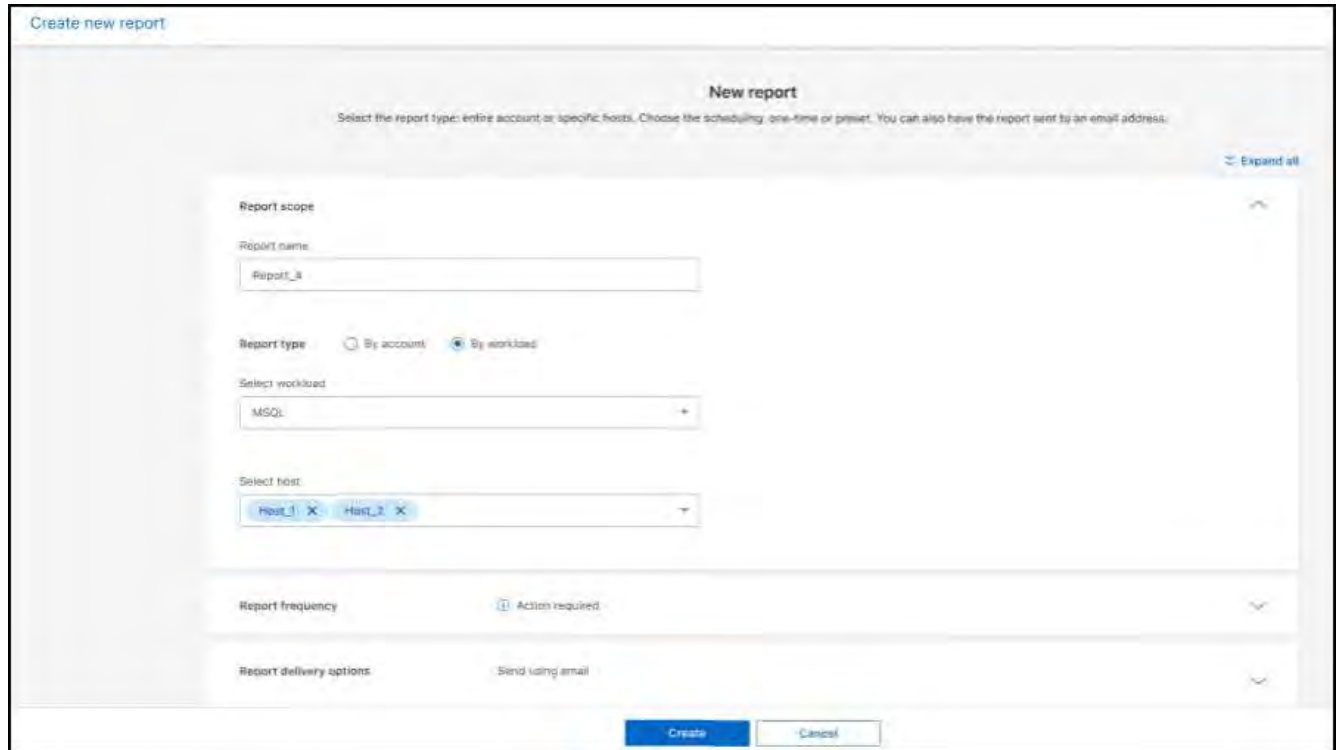
In BlueXP backup and recovery, create reports for Microsoft SQL Server workloads to view the status of your backups, including the number of backups, the number of successful backups, and the number of failed backups. You can also view the details of each backup, including the backup type, the storage system used for the backup, and the time of the backup.

### Create a report

#### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, Backup and Recovery clone admin. Learn about [Backup and recovery roles and privileges](#). Learn about [BlueXP access roles for all services](#).

1. From the BlueXP backup and recovery menu, select the **Reports** tab.
2. Select **Create report**.



3. Enter report scope details:

- **Report name:** Enter a unique name for the report.

- **Report type:** Choose whether you want a report by account or by workload (Microsoft SQL Server).
  - **Select host:** If you selected by workload, select the host for which you want to generate the report.
  - **Select contents:** Choose whether you want the report to include a summary of all backups or details of each backup. (If you chose "By account")
4. Enter reporting range: Choose whether you want the report to include data from the last day, last 7 days, last 30 days, last quarter, or last year.
  5. Enter report delivery details: If you want the report to be delivered by email, check **Send report using email**. Enter the email addresses where you want the report sent.

Configure email notifications in the Settings page. For details about configuring email notifications, see [Configure settings](#).

## Protect virtual machine workloads

### Protect virtual machines workloads in BlueXP backup and recovery overview

Protect your virtual machines workloads with BlueXP backup and recovery. BlueXP backup and recovery provides fast, space-efficient, crash-consistent, and VM-consistent backup and restore operations for VMs, datastores, and VMDKs.

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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

For instructions on protecting virtual machines workloads, see the following topics:

- [Create a policy for VMware workloads](#)
- [Back up VMware datastores to Amazon Web Services](#)
- [Back up VMware datastores to Microsoft Azure](#)
- [Back up VMware datastores to Google Cloud Platform](#)
- [Back up VMware datastores to StorageGRID](#)
- [Restore VMware workloads](#)
- [Manage protection for VMware workloads](#)

### Prerequisites for virtual machines workloads in BlueXP backup and recovery

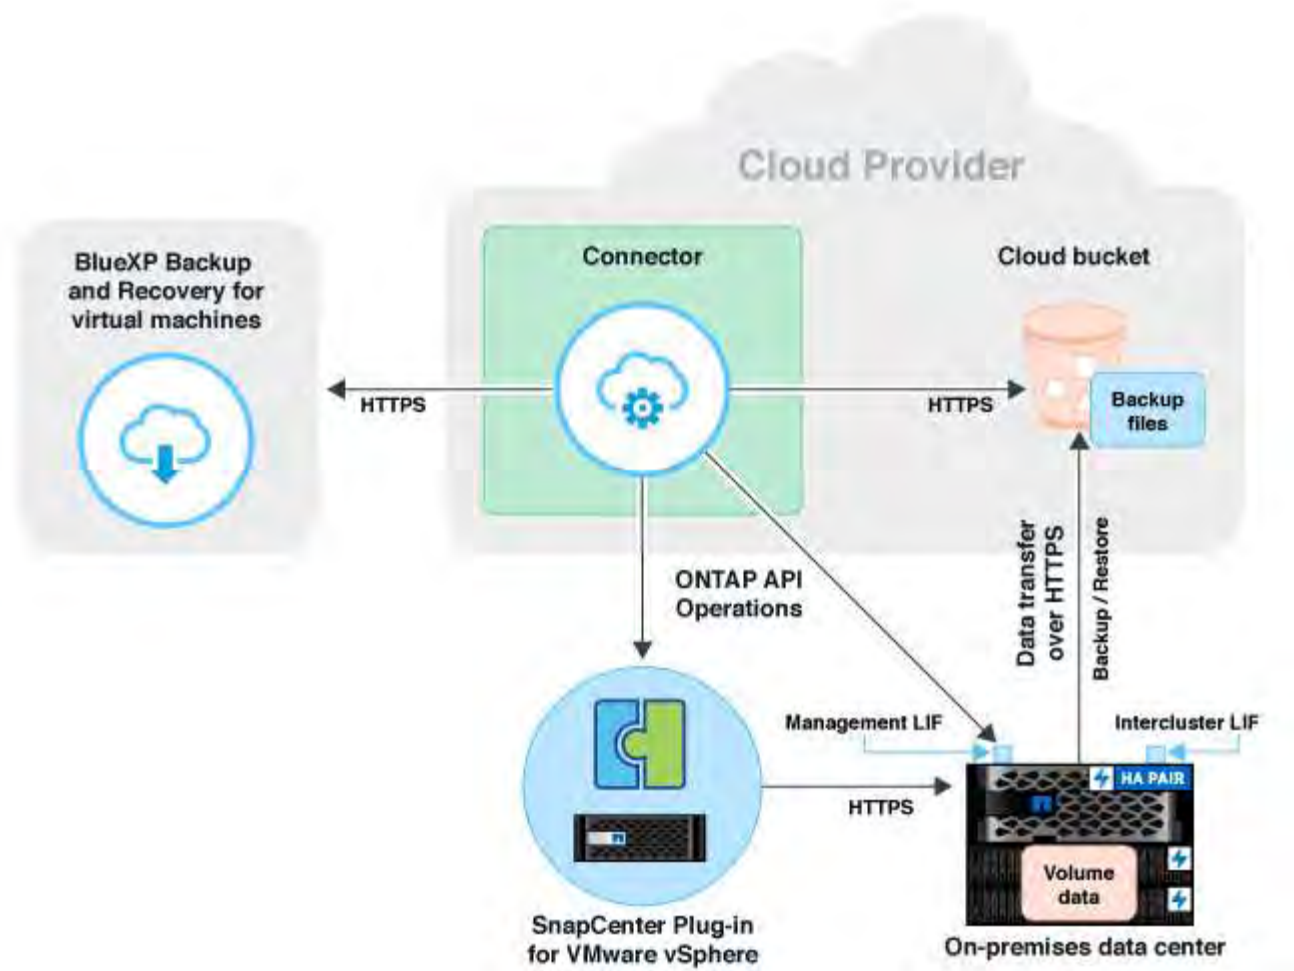
Before you begin protecting your virtual machines workloads with BlueXP backup and recovery, ensure that you meet the following prerequisites:

- 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, the 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 a 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.
- If a volume backing up a datastore is already protected from the Volumes tab (BlueXP Backup and recovery → Volumes), then the same datastore cannot be protected again from the Virtual Machines tab (BlueXP Backup and recovery → Virtual Machines).

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



## Register SnapCenter Plug-in for VMware vSphere host to use with BlueXP backup and recovery

You should register the SnapCenter Plug-in for VMware vSphere host in BlueXP backup and recovery 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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### Steps

1. In BlueXP UI, select **Protection > Backup and recovery > Virtual Machines**.
2. From the **Settings** drop-down, select **SnapCenter Plug-in for VMware vSphere**.
3. Select **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. Select **Register**.

### After you finish

Select **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 in BlueXP backup and recovery

You can create a policy or use one of the following predefined policies that are available in BlueXP backup and recovery.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### 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. Select **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. Select **Create**.



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

## Back up datastores to Amazon Web Services in BlueXP backup and recovery

You can back up and archive one or more datastores with BlueXP backup and recovery 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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### Before you begin

Ensure that you have met all the [virtual machine protection requirements](#) before backing up datastores to the cloud.

### Steps

1. In BlueXP UI, select **Protection > Backup and recovery > Virtual Machines**.
2. Select **...** corresponding to the datastore that you want to back up and click **Activate Backup**.
3. In the Assign Policy page, select the policy and select **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. Select **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. Select **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 select **Activate Backup**.

## Back up datastores to Microsoft Azure with BlueXP backup and recovery

You can back up one or more datastores to Microsoft Azure by integrating the SnapCenter Plug-in for VMware vSphere host with BlueXP backup and recovery. 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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### Before you begin

Ensure that you have met all the [virtual machine protection requirements](#) before backing up datastores to the cloud.

### Steps

1. In BlueXP UI, select **Protection > Backup and recovery > Virtual Machines**.
2. Select **...** corresponding to the datastore that you want to back up and select **Activate Backup**.
3. In the Assign Policy page, select the policy and select **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. Select **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. Select **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 a one-time activity and you will not be allowed to set it up later.

6. Review the details and select **Activate Backup**.

## Back up datastores to Google Cloud Platform with BlueXP backup and recovery

You can back up one or more datastores to Google Cloud Platform by integrating the SnapCenter Plug-in for VMware vSphere host with BlueXP backup and recovery. This will help the VM administrators to easily and quickly back up and archive data for storage efficiency and accelerate cloud transition.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### Before you begin

Ensure that you have met all the [virtual machine protection requirements](#) before backing up datastores to the cloud.

### Steps

1. In BlueXP UI, select **Protection > Backup and recovery > Virtual Machines**.
2. Select **...** corresponding to the datastore that you want to back up and select **Activate Backup**.
3. In the Assign Policy page, select the policy and select **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. Select **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. Select **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 select **Activate Backup**.

## Back up datastores to StorageGRID with BlueXP backup and recovery

You can back up one or more datastores to StorageGRID by integrating the SnapCenter Plug-in for VMware vSphere host with BlueXP backup and recovery. This will help the VM administrators to easily and quickly back up and archive data for storage efficiency and accelerate cloud transition.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

### Before you begin

Ensure that you have met all the [virtual machine protection requirements](#) before backing up datastores to the cloud.

### Steps

1. In BlueXP UI, select **Protection > Backup and recovery > Virtual Machines**.
2. Select **...** corresponding to the datastore that you want to back up and click **Activate Backup**.
3. In the Assign Policy page, select the policy and select **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. Select **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. Select **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 select **Activate Backup**.

## Manage protection of datastores and VMs in BlueXP backup and recovery

You can view policies, datastores, and virtual machines before you back up and restore data with BlueXP backup and recovery. Depending upon the change in database, policies, or resource groups, you can view the updates from the BlueXP UI.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

## 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, select **Protection > Backup and recovery > Virtual Machines**.
2. From the **Settings** drop-down, select **Policies**.
3. Select **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, select **Protection > Backup and recovery > Virtual Machines > Settings > SnapCenter Plug-in for VMware vSphere**.
2. Select 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, select **Protection > Backup and recovery > Virtual Machines**.
2. Select the Actions icon **...** corresponding to the datastore that you want to unprotect and select **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, select **Protection > Backup and recovery > Virtual Machines > Settings > SnapCenter Plug-in for VMware vSphere**.
2. Select the Actions icon **...** and select **Edit**.
3. Modify the details as required.
4. Select **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. Select **Backup and recovery > Virtual Machines**.
2. From the **Settings** drop-down, select **SnapCenter Plug-in for VMware vSphere**.
3. Select the Actions icon **...** corresponding to the SnapCenter Plug-in for VMware vSphere host and select **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. Select **Backup and recovery > Virtual Machines**.
2. Select the Actions icon **...** corresponding to the datastore and select **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. Select **Backup and recovery > Virtual Machines**.
2. From the **Settings** drop-down, select **SnapCenter Plug-in for VMware vSphere**.
3. Select the Actions icon **...** corresponding to the SnapCenter Plug-in for VMware vSphere host and select **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. Select **Backup and recovery > Job Monitoring**.

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

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

## Restore virtual machines data with BlueXP backup and recovery

You can restore virtual machines data from the cloud back to the on-premises vCenter with BlueXP backup and recovery. 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.

**NOTE** To switch to and from BlueXP backup and recovery workloads, refer to [Switch to different BlueXP backup and recovery workloads](#).

**Before you begin**

- Ensure that you have met all the [virtual machine protection requirements](#) before backing up datastores to 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, select **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"> <li>1. Select the Actions icon <b>...</b> corresponding to the datastore that you want to restore and click <b>View Details</b>.</li> <li>2. Select <b>Restore</b> corresponding to the backup you want to restore.</li> <li>3. Select the virtual machine that you want to restore from the backup and select <b>Next</b>.</li> <li>4. Ensure that <b>Original</b> is selected and select <b>Continue</b>.</li> <li>5. If the virtual machine is protected using a policy where archival settings are configured, select the <b>Archival Restore Priority</b> and select <b>Next</b>.  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.</li> <li>6. Review the details and select <b>Restore</b>.</li> </ol>

If you want to restore the virtual machine...	Do this...
to an alternate location from datastore	<ol style="list-style-type: none"> <li>1. Select the Actions icon <b>...</b> corresponding to the datastore that you want to restore and select <b>View Details</b>.</li> <li>2. Select <b>Restore</b> corresponding to the backup you want to restore.</li> <li>3. Select the virtual machine that you want to restore from the backup and select <b>Next</b>.</li> <li>4. Select <b>Alternate</b>.</li> <li>5. Select the alternate vCenter Server, ESXi host, datastore, and network.</li> <li>6. Provide a name for the VM after restore and select <b>Continue</b>.</li> <li>7. If the virtual machine is protected using a policy where archival settings are configured, select the <b>Archival Restore Priority</b> and select <b>Next</b>.  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.</li> <li>8. Review the details and select <b>Restore</b>.</li> </ol>
to the original location from virtual machines	<ol style="list-style-type: none"> <li>1. Select the Actions icon <b>...</b> corresponding to the virtual machine that you want to restore and select <b>Restore</b>.</li> <li>2. Select the backup through which you want to restore the virtual machine.</li> <li>3. Ensure that <b>Original</b> is selected and select <b>Continue</b>.</li> <li>4. If the virtual machine is protected using a policy where archival settings are configured, select the <b>Archival Restore Priority</b> and select <b>Next</b>.  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.</li> <li>5. Review the details and select <b>Restore</b>.</li> </ol>

If you want to restore the virtual machine...	Do this...
to an alternate location from virtual machines	<ol style="list-style-type: none"> <li>1. Select the Actions icon <b>...</b> corresponding to the virtual machine that you want to restore and select <b>Restore</b>.</li> <li>2. Select the backup through which you want to restore the virtual machine.</li> <li>3. Select <b>Alternate</b>.</li> <li>4. Select the alternate vCenter Server, ESXi host, datastore, and network.</li> <li>5. Provide a name for the VM after restore and select <b>Continue</b>.</li> <li>6. If the virtual machine is protected using a policy where archival settings are configured, select the <b>Archival Restore Priority</b> and select <b>Next</b>.  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.</li> <li>7. Review the details and select <b>Restore</b>.</li> </ol>



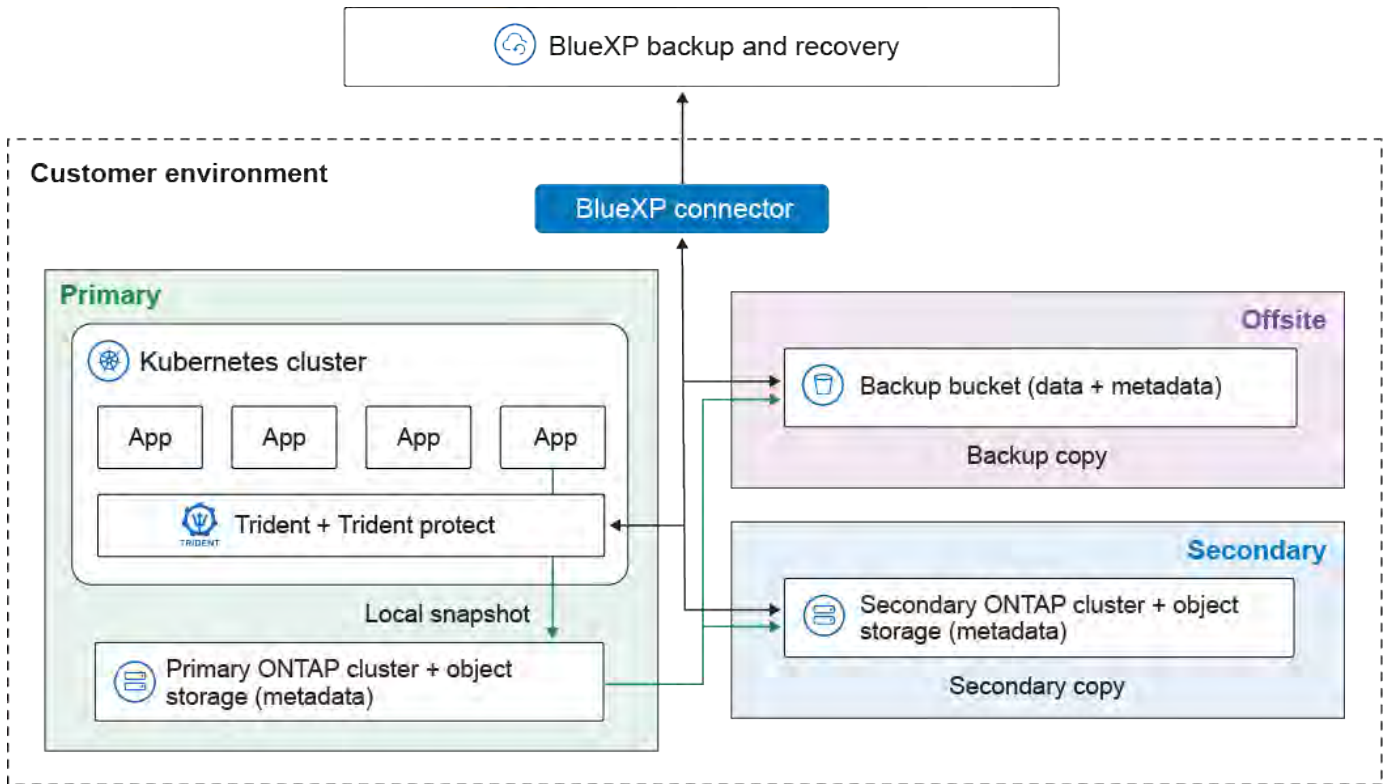
If the restore operation does not complete, do not try the restore process again until the Job Monitor shows that the restore operation has failed. If you try the restore process again before the Job Monitor shows that the restore operation has failed, the restore operation will fail again. When you see the Job Monitor status as "Failed," you can try the restore process again.

## Protect Kubernetes workloads (Preview)

### Manage Kubernetes workloads overview

Managing Kubernetes workloads in BlueXP backup and recovery enables you to discover, manage, and protect your Kubernetes clusters and applications all in one place. You can manage resources and applications hosted on your Kubernetes clusters. You can also create and associate protection policies with your Kubernetes workloads, all using a single interface.

The following diagram shows the components and basic architecture of backup and recovery for Kubernetes workloads and how different copies of your data can be stored in different locations:



BlueXP backup and recovery provides the following benefits for managing Kubernetes workloads:

- A single control plane for protecting applications running across multiple Kubernetes clusters. These applications can include containers or virtual machines running on your Kubernetes clusters.
- Native integration with NetApp SnapMirror, enabling storage offloading capabilities for all backup and recovery workflows.
- Incremental forever backups for Kubernetes applications, translating to lower Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs).



This documentation is provided as a technology preview. During the preview, Kubernetes functionality is not recommended for production workloads. With this preview offering, NetApp reserves the right to modify offering details, contents, and timeline before General Availability.

You can accomplish the following tasks related to managing Kubernetes workloads:

- [Discover Kubernetes workloads.](#)
- [Manage Kubernetes clusters.](#)
- [Add and protect Kubernetes applications.](#)
- [Manage Kubernetes applications.](#)
- [Restore Kubernetes applications.](#)

## Discover Kubernetes workloads in BlueXP backup and recovery

The BlueXP backup and recovery service needs to first discover Kubernetes workloads in order for you to use the service.

### Required BlueXP role

This task requires the data services Backup and recovery super admin role. Learn about [Backup and recovery data services roles and privileges](#). Learn about [BlueXP access roles for all services](#).

## Discover Kubernetes workloads

In the backup and recovery inventory, you can discover Kubernetes workloads that are running in your environment. Discovering a workload adds a Kubernetes cluster to BlueXP backup and recovery, enabling you to then add applications to the cluster and protect the resources hosted by the cluster.

### Steps

1. Do one of the following:
  - If you are discovering Kubernetes workloads for the first time, in BlueXP backup and recovery, select **Discover and Manage** under the Kubernetes workload type.
  - If you have already discovered Kubernetes workloads, in BlueXP backup and recovery, select **Inventory > Workloads** and then select **Discover resources**.

2. Select the **Kubernetes** workload type.
3. Enter a cluster name and choose a connector to use with the cluster.
4. Follow the command line instructions that appear:
  - Create a Trident protect namespace
  - Create a Kubernetes secret
  - Add a Helm repository
  - Install Trident protect and the Trident protect connector

These steps ensure that BlueXP backup and recovery can interact with the cluster.

5. After you complete the steps, select **Discover**.

The cluster is added to the inventory.

6. Select **View** in the associated Kubernetes workload to see the list of applications, clusters, and namespaces for that workload.

## Continue to the BlueXP backup and recovery Dashboard

To display the BlueXP backup and recovery Dashboard, follow these steps.

1. From the top menu, select **Dashboard**.
2. Review the health of data protection. The number of at risk or protected workloads increases based on the newly discovered, protected, and backed up workloads.

[Learn what the Dashboard shows you.](#)

## Add and protect Kubernetes applications

BlueXP backup and recovery enables you to easily discover your Kubernetes clusters, without generating and uploading kubeconfig files. You can connect Kubernetes clusters and install the required software using simple commands copied from the BlueXP user interface.

## Required BlueXP role

Organization admin or SnapCenter admin. [Learn about BlueXP backup and recovery access roles.](#) [Learn about BlueXP access roles for all services.](#)

## Add and protect a new Kubernetes application

The first step in protecting Kubernetes applications is to create an application within BlueXP backup and recovery. When you create an application, you make BlueXP aware of the running application on the Kubernetes cluster.

### Before you begin

Before you can add and protect a Kubernetes application, you need to [discover Kubernetes workloads](#).

### Steps

1. In BlueXP backup and recovery, select **Inventory**.
2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.
3. Select the **Applications** tab.
4. Select **Create application**.
5. Enter a name for the application.
6. Optionally, choose any of the following fields to search for the resources you want to protect:
  - Associated cluster
  - Associated namespaces
  - Resource types
  - Label selectors
7. Optionally, select **Cluster Scoped Resources** to choose any resources that are scoped at the cluster level. If you include them, they are added to the application when you create it.
8. Optionally, select **Search** to find the resources based on your search criteria.



BlueXP does not store the search parameters or results; the parameters are used to search the selected Kubernetes cluster for resources that can be included in the application.

9. BlueXP displays a list of resources that match your search criteria.
10. If the list contains the resources you want to protect, select **Next**.
11. Optionally, in the **Policy** area, choose an existing protection policy to protect the application or create a new policy. If you don't select a policy, the application is created without a protection policy. You can [add a protection policy](#) later.
12. In the **Prescripts and postscripts** area, enable and configure any prescript or postscript execution hooks that you want to run before or after backup operations. To enable prescripts or postscripts, you must have already created at least one [execution hook template](#).
13. Select **Create**.

### Result

The application is created and appears in the list of applications in the **Applications** tab of the Kubernetes inventory. BlueXP enables protection for the application based on your settings, and you can monitor the progress in the **Monitoring** area of backup and recovery.

## Protect an existing Kubernetes application

Enable a protection policy on a Kubernetes application that you have already added.

### Steps

1. In BlueXP backup and recovery, select **Inventory**.
2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.
3. Select the **Applications** tab.
4. In the list of applications, choose an application you want to protect and select the associated Actions menu.
5. Select **Protect**.
6. In the **Policy** area, choose an existing protection policy to protect the application or create a new policy. Refer to [Create a policy](#) for more information about creating protection policies.
7. In the **Prescripts and postscripts** area, enable and configure any prescript or postscript execution hooks that you want to run before or after backup operations. You can configure the type of execution hook, the template it uses, arguments, and label selectors.
8. Select **Done**.

### Result

BlueXP enables protection for the application based on your settings, and you can monitor the progress in the **Monitoring** area of backup and recovery. As soon as you enable protection for an application, BlueXP creates a full backup of the application. Any future incremental backups are created based on the schedule that you define in the protection policy associated with the application.

## Back up a Kubernetes application now

Manually create a backup of a Kubernetes application to establish a baseline for future backups and snapshots, or to ensure the most recent data is protected.

### Steps

1. In BlueXP backup and recovery, select **Inventory**.
2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.
3. Select the **Applications** tab.
4. In the list of applications, choose an application you want to back up and select the associated Actions menu.
5. Select **Backup now**.
6. Ensure the correct application name is selected.
7. Select **Back up**.

### Result

BlueXP creates a backup of the application and displays the progress in the **Monitoring** area of backup and recovery. The backup is created based on the protection policy associated with the application.

## Restore Kubernetes applications

BlueXP backup and recovery enables you to restore applications that you have protected with a protection policy. To restore an application, an application needs to have at least



one restore point available. A restore point consists of either the local snapshot or the backup to the object store (or both). You can restore an application using the local, secondary, or object store archive.

#### Required BlueXP role

Organization admin or SnapCenter admin. [Learn about BlueXP backup and recovery access roles.](#) [Learn about BlueXP access roles for all services.](#)

#### Steps

1. In BlueXP backup and recovery, select **Inventory**.
2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.
3. Select the **Applications** tab.
4. In the list of applications, choose an application you want to restore and select the associated Actions menu.
5. Select **View and restore**.

The list of restore points appears.

6. Open the Actions menu for the restore point you want to use, and select **Restore**.

#### General settings

1. Choose the source to restore from (local or object store).
2. Choose the destination cluster from the **Cluster** list.
3. Choose the restore destination namespace.

You can restore to the original namespace or restore to a new namespace.

4. Select **Next**.

#### Resource selection

1. Choose whether you want to restore all resources associated with the application or use a filter to select specific resources to restore:

### Restore all resources

- a. Select **Restore all resources**.
- b. Select **Next**.

### Restore specific resources

- a. Select **Selective resources**.
- b. Choose the behavior of the resource filter. If you choose **Include**, the resources you select are restored. If you choose **Exclude**, the resources you select are not restored.
- c. Select **Add rules** to add rules that define filters for selecting resources. You need at least one rule to filter resources.

Each rule can filter on criteria such as the resource namespace, labels, group, version, and kind.

- d. Select **Save** to save each rule.
- e. When you have added all the rules you need, select **Search** to see the resources available in the backup archive that match your filter criteria.



The resources shown are the resources that currently exist on the cluster.

- f. When satisfied with the results, select **Next**.

### Destination settings

1. Choose to restore either to the default storage class or to a different storage class.
2. Optionally, if you chose to restore to a different storage class, select a destination storage class to match each source storage class.
3. Select **Restore**.

### Manage Kubernetes clusters

BlueXP backup and recovery enables you to discover and manage your Kubernetes clusters so that you can protect resources hosted by the clusters.

#### Required BlueXP role

Organization admin or SnapCenter admin. [Learn about BlueXP backup and recovery access roles](#). [Learn about BlueXP access roles for all services](#).



To discover Kubernetes clusters, refer to [Discover Kubernetes workloads](#).

### Edit Kubernetes cluster information

You can edit a cluster if you need to change its name.

#### Steps

1. In BlueXP backup and recovery, select **Inventory > Clusters**.
2. In the list of clusters, choose a cluster you want to edit and select the associated Actions menu.

3. Select **Edit cluster**.
4. Make any required changes to the cluster name. The cluster name needs to match the name that you used with the Helm command during the discovery process.
5. Select **Done**.

## Remove a Kubernetes cluster

If you no longer need to protect the resources hosted by a Kubernetes cluster, you can remove it from BlueXP backup and recovery. Removing a cluster does not delete the cluster or its resources; it only removes the cluster from the BlueXP inventory. Before you can remove a cluster, you need to disable protection and delete the associated applications from BlueXP backup and recovery.

### Steps

1. In BlueXP backup and recovery, select **Inventory > Clusters**.
2. In the list of clusters, choose a cluster you want to edit and select the associated Actions menu.
3. Select **Remove cluster**.
4. Review the information in the confirmation dialog box, and select **Remove**.

## Manage Kubernetes applications

BlueXP backup and recovery enables you to unprotect and delete your Kubernetes applications and associated resources.

### Required BlueXP role

Organization admin or SnapCenter admin. [Learn about BlueXP backup and recovery access roles](#). [Learn about BlueXP access roles for all services](#).

## Unprotect a Kubernetes application

You can unprotect an application if you no longer want to protect it. When you unprotect an application, BlueXP backup and recovery stops protecting the application but keeps all associated backups and snapshots.

### Steps

1. In BlueXP backup and recovery, select **Inventory**.
2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.
3. Select the **Applications** tab.
4. In the list of applications, choose an application you want to unprotect and select the associated Actions menu.
5. Select **Unprotect**.
6. Read the notice, and when ready, select **Unprotect**.

## Delete a Kubernetes application

You can delete an application if you no longer need it. When you delete an application, BlueXP backup and recovery stops protecting the application and deletes all associated backups and snapshots.

### Steps

1. In BlueXP backup and recovery, select **Inventory**.

2. Choose a Kubernetes instance, and select **View** to view the resources associated with that instance.
3. Select the **Applications** tab.
4. In the list of applications, choose an application you want to delete and select the associated Actions menu.
5. Select **Delete**.
6. Enable **Delete snapshots and backups** to remove all snapshots and backups of the application.



You will no longer be able to restore the application using these snapshots and backups.

7. Confirm the action and select **Delete**.

## Manage BlueXP backup and recovery execution hook templates for Kubernetes workloads

An execution hook is a custom action that you can configure to run in conjunction with a data protection operation of a managed Kubernetes application. For example, if you have a database app, you can use an execution hook to pause all database transactions before a snapshot, and resume transactions after the snapshot is complete. This ensures application-consistent snapshots. When you create an execution hook template, you can specify the type of hook, the script to run, and any filters that determine which containers the hook applies to. You can then use the template to associate execution hooks with your applications.

### Required BlueXP role

Organization admin or SnapCenter admin. [Learn about BlueXP backup and recovery access roles.](#) [Learn about BlueXP access roles for all services.](#)

### Types of execution hooks

BlueXP backup and recovery supports the following types of execution hooks, based on when they can be run:

- Pre-snapshot
- Post-snapshot
- Pre-backup
- Post-backup
- Post-restore

### Order of execution

When a data protection operation is run, execution hook events take place in the following order:

1. Any applicable custom pre-operation execution hooks are run on the appropriate containers. You can create and run as many custom pre-operation hooks as you need, but the order of execution of these hooks before the operation is neither guaranteed nor configurable.
2. Filesystem freezes occur, if applicable.
3. The data protection operation is performed.

4. Frozen filesystems are unfrozen, if applicable.
5. Any applicable custom post-operation execution hooks are run on the appropriate containers. You can create and run as many custom post-operation hooks as you need, but the order of execution of these hooks after the operation is neither guaranteed nor configurable.

If you create multiple execution hooks of the same type (for example, pre-snapshot), the order of execution of those hooks is not guaranteed. However, the order of execution of hooks of different types is guaranteed. For example, the following is the order of execution of a configuration that has all of the different types of hooks:

1. Pre-snapshot hooks executed
2. Post-snapshot hooks executed
3. Pre-backup hooks executed
4. Post-backup hooks executed



You should always test your execution hook scripts before enabling them in a production environment. You can use the 'kubectl exec' command to conveniently test the scripts. After you enable the execution hooks in a production environment, test the resulting snapshots and backups to ensure they are consistent. You can do this by cloning the app to a temporary namespace, restoring the snapshot or backup, and then testing the app.



If a pre-snapshot execution hook adds, changes, or removes Kubernetes resources, those changes are included in the snapshot or backup and in any subsequent restore operation.

### Important notes about custom execution hooks

Consider the following when planning execution hooks for your apps.

- An execution hook must use a script to perform actions. Many execution hooks can reference the same script.
- Execution hooks need to be written in the format of executable shell scripts.
- Script size is limited to 96KB.
- Execution hook settings and any matching criteria are used to determine which hooks are applicable to a snapshot, backup, or restore operation.



Because execution hooks often reduce or completely disable the functionality of the application they are running against, you should always try to minimize the time your custom execution hooks take to run. If you start a backup or snapshot operation with associated execution hooks but then cancel it, the hooks are still allowed to run if the backup or snapshot operation has already begun. This means that the logic used in a post-backup execution hook cannot assume that the backup was completed.

### Execution hook filters

When you add or edit an execution hook for an application, you can add filters to the execution hook to manage which containers the hook will match. Filters are useful for applications that use the same container image on all containers, but might use each image for a different purpose (such as Elasticsearch). Filters allow you to create scenarios where execution hooks run on some but not necessarily all identical containers. If you create multiple filters for a single execution hook, they are combined with a logical AND operator. You can have up to 10 active filters per execution hook.

Each filter you add to an execution hook uses a regular expression to match containers in your cluster. When a hook matches a container, the hook will run its associated script on that container. Regular expressions for filters use the Regular Expression 2 (RE2) syntax, which does not support creating a filter that excludes containers from the list of matches. For information on the syntax that BlueXP backup and recovery supports for regular expressions in execution hook filters, see [Regular Expression 2 \(RE2\) syntax support](#).



If you add a namespace filter to an execution hook that runs after a restore or clone operation and the restore or clone source and destination are in different namespaces, the namespace filter is only applied to the destination namespace.

## Execution hook examples

Visit the [NetApp Verda GitHub project](#) to download real execution hooks for popular apps such as Apache Cassandra and Elasticsearch. You can also see examples and get ideas for structuring your own custom execution hooks.

## Create an execution hook template

You can create a custom execution hook template that you can use to perform actions before or after a data protection operation on an application.

### Steps

1. In BlueXP, go to **Protection > Backup and recovery**.
2. Select the **Settings** tab.
3. Expand the **Execution hook template** section.
4. Select **Create execution hook template**.
5. Enter a name for the execution hook.
6. Optionally, choose a type of hook. For example, a post-restore hook is run after the restore operation is complete.
7. In the **Script** text box, enter the executable shell script that you want to run as part of the execution hook template. Optionally, you can select **Upload script** to upload a script file instead.
8. Select **Create**.

The template is created and appears in the list of templates in the **Execution hook template** section.

## Monitor jobs in BlueXP backup and recovery

With BlueXP backup and recovery, 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.

### Required BlueXP role

Organization admin, Folder or project admin, Backup and Recovery super admin, Backup and Recovery backup admin, Backup and Recovery restore admin, Backup and Recovery clone admin, or Backup and

Recovery viewer role. Learn about [Backup and recovery roles and privileges](#). Learn about [BlueXP access roles for all services](#).

## 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, and virtual machines. 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 backup and recovery menu, select the **Monitoring** tab.
2. To show additional columns (Working Environment, SVM, User Name, Workload, Policy Name, Snapshot Label), select the plus sign.

### 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. From the BlueXP backup and recovery menu, select the **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.

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

- ONTAP volumes (Cloud Volumes ONTAP and on-premises ONTAP volumes)
- Microsoft SQL Server
- 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. From the BlueXP backup and recovery menu, select the **Monitoring** tab.
2. Select the name of the job.
3. Select the Actions menu  and select **View Details**.
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. From the BlueXP backup and recovery menu, select the **Monitoring** tab.
2. To download a CSV file for all jobs, select the Download 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. From the BlueXP backup and recovery menu, select the **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 break 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

Operation type	Event	Alert level	Email sent
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 organization and 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.](#)

## 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.

### Steps

1. Connect to the Linux system that the Connector is running on.

Connector location	Procedure
Cloud deployment	Follow the instructions for <a href="#">connecting to the Connector Linux virtual machine</a> depending on the cloud provider you're using.
Manual installation	Log in to the Linux system.

2. Enter the command to restart the service.

<b>Connector location</b>	<b>Docker command</b>	<b>Podman command</b>
Cloud deployment	<code>docker restart cloudmanager_cbs</code>	<code>podman restart cloudmanager_cbs</code>
Manual installation with internet access	<code>docker restart cloudmanager_cbs</code>	<code>podman restart cloudmanager_cbs</code>
Manual installation without internet access	<code>docker restart ds_cloudmanager_cbs_1</code>	<code>podman restart ds_cloudmanager_cbs_1</code>

# Automate with BlueXP backup and recovery REST 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

## API reference

Documentation for each BlueXP backup and recovery API is available from [BlueXP automation for BlueXP backup and recovery](#).

## 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.



If you are using a service account, you should use the service access token instead of a user token. The value for "client\_id" ("Mu0V1ywgYtel6w1MbD15fKfVIUrNXGWC") is a fixed value and cannot be changed. In this case, follow the instructions here: [Create a service access token](#).

### 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": "JxaVHn9cGkX92aPVCkhat3zxxxxxwsC9qMl_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-
OOoAR4ZS", "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/occ
mui", "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://local
host:1337", "http://localhost:1337/occmui", "https://localhost", "https://l
ocalhost/occmui", "https://localhost:1337", "https://localhost:1337/occmui
"], "createDate": "1652120369286", "agent": { "useDockerInfra": true, "network"
: "default", "name": "cbsgcpdevcntsg-
asia", "agentId": "imEdsEW4HyYTFbt8ZcNKTKDF05jMIe6Zclients", "provider": "gc
p", "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 `iZwFFeVCZjWnzG1w8RgD0QQNANZvpP7Iclients`, and user token `Bearer eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6Iks5rSx1PVFUzUWpZek1E...y6nyhBjwkeMwHc4ValobjUmju2x0xUH48g` in this command.

```

curl --location --request POST
'https://api.bluexp.netapp.com/account/account-
DpTFcxN3/providers/cloudmanager_cbs/api/v3/backup/working-
environment/VsaWorkingEnvironment-99hPYEgk' \
--header 'x-agent-id: iZwFFeVCZjWnzG1w8RgD0QQNANZvpP7Iclients' \
--header 'Accept: application/json' \
--header 'Content-Type: application/json' \
--header 'Authorization: Bearer
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6Ikk5rSx1PVFUzUWpZek1E...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
eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6Iks5rSXlPVFUzUWpZek1E...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
    }
  ]
}
```



# Reference

## Policies in SnapCenter compared to those in BlueXP backup and recovery

There are some differences between policies used in SnapCenter and those used in BlueXP backup and recovery that might impact what you see after importing resources and policies from SnapCenter.

### Schedule tiers

SnapCenter uses the following schedule tiers:

- **Hourly:** Multiple hours and minutes with any hours (0-23) and any minutes (0-60).
- **Daily:** Includes an option to repeat every so many days, for example, every 3 days.
- **Weekly:** Sunday to Monday, with an option to perform a snapshot on Day 1 of the week or on multiple days of the week.
- **Monthly:** Months January to December, with an option to perform on specific days of the month, for example, the 7th of every month and even on multiple days of the month.

BlueXP backup and recovery uses the following schedule tiers, which are slightly different:

- **Hourly:** Performs snapshots only on 15-minute intervals, for example, 1 hour or 15-minute intervals less than 60.
- **Daily:** Hours of the day (0-23) with start time for example at 10:00 AM with an option to perform every so many hours.
- **Weekly:** Day of the week (Sunday to Monday) with an option to perform on 1 day or multiple days. This is the same as SnapCenter.
- **Monthly:** Dates of the month (0-30) with a starting time on multiple dates of the month.
- **Yearly:** Monthly. This matches SnapCenter's monthly.

### Multiple policies in SnapCenter with the same schedule tier

You can assign multiple policies with the same schedule tier to a resource in SnapCenter. However, BlueXP backup and recovery does not support multiple policies on a resource that uses the same schedule tier.

**Example:** If you use three policies (for Data, Log, and Log of snapshots) in SnapCenter, after migration from SnapCenter, BlueXP backup and recovery uses a single policy instead of all three.

### Imported SnapCenter daily schedules

BlueXP backup and recovery adjusts the SnapCenter schedules as follows:

- If the SnapCenter schedule is set to less than or equal to 7 days, BlueXP backup and recovery sets the schedule to weekly. Some snapshots will be skipped during the week.

**Example:** If you have a SnapCenter daily policy with a repeating interval of every 3 days starting on Monday, BlueXP backup and recovery sets the schedule to weekly on Monday, Thursday, and Sunday.

Some days will be skipped because it is not exactly every 3 days.

- If the SnapCenter schedule is set to greater than 7 days, BlueXP backup and recovery sets the schedule to monthly. Some snapshots will be skipped during the month.

**Example:** If you have a SnapCenter daily policy with a repeating interval of every 10 days starting on the 2nd of the month, BlueXP backup and recovery (post migration) sets the schedule to monthly on the 2nd, 12th, and 22nd day of the month. Some days will be skipped the next month.

## Imported SnapCenter hourly schedules

SnapCenter hourly policies with repeating intervals greater than one hour are converted to a daily policy in BlueXP backup and recovery.

Any hourly policy with repeating intervals that are not a factor of 24 (for example 5, 7, etc) will skip some snapshots in a day.

**Example:** If you have a SnapCenter hourly policy with a repeating interval every 5 hours starting at 1:00 AM, BlueXP backup and recovery (after migration) will set the schedule to daily with 5-hour intervals at 1:00 AM, 6:00 AM, 11:00 AM, 4:00 PM, and 9:00 PM. Some hours will be skipped, after 9:00 PM it should be 2:00 AM to repeat after every 5 hours, but it will be always 1:00 AM.

## Log retention from SnapCenter policies

If you have a resource in SnapCenter with multiple policies, BlueXP backup and recovery uses the following priority order to assign the log retention value:

- For "Full backup with log backup policy" plus "log-only" policies in SnapCenter, BlueXP backup and recovery uses the log-only policy retention value.
- For "Full backup with log only" and "Full and Log" policies in SnapCenter, BlueXP backup and recovery uses the log-only retention value.
- For "Full backup and log" plus "Full backup" in SnapCenter, BlueXP backup and recovery uses the "Full backup and log" retention value.
- If you have only a full backup in SnapCenter, BlueXP backup and recovery does not enable the log backup.

## Log backup retention

With SnapCenter, you can have multiple retention values across multiple policies attached to a resource. However, BlueXP backup and recovery supports only a single retention value for all policies attached to a resource.

## Retention count from SnapCenter policies

If you have a resource with secondary protection enabled in SnapCenter with multiple source volumes, multiple destination volumes, and multiple SnapMirror relationships, BlueXP backup and recovery uses only the first policy's retention count.

**Example:** If you have a SnapCenter policy with a retention count of 5 and another policy with a retention count of 10, BlueXP backup and recovery uses the retention count of 5.

## SnapMirror labels from SnapCenter policies

SnapMirror labels for every policy in SnapCenter remain intact post migration even though the tier is changed.

**Example:** An hourly policy from SnapCenter might change to daily in BlueXP backup and recovery. However, the SnapMirror labels remain the same after migration.

## BlueXP backup and recovery identity and access management to features

BlueXP backup and recovery employs identity and access management (IAM) to govern the access that each user has to specific features and actions.

The service uses the following roles that are specific to BlueXP backup and recovery.

- **Backup and recovery super admin:** Perform any actions in BlueXP backup and recovery.
- **Backup admin:** Perform backups to local snapshots, replicate to secondary storage, and back up to object storage actions in BlueXP backup and recovery.
- **Restore admin:** Restore workloads using BlueXP backup and recovery.
- **Clone admin:** Clone applications and data using BlueXP backup and recovery.
- **Backup and recovery viewer:** View information in BlueXP backup and recovery, but not perform any actions.

For details about all BlueXP access roles, see [the BlueXP setup and administration documentation](#).

The following table indicates the actions that each BlueXP backup and recovery role can perform.

Feature and action	Backup and recovery super admin	Backup admin	Restore admin	Clone admin	Viewer
Add, edit, or delete hosts	Yes	No	No	No	No
Install plugins	Yes	No	No	No	No
Add credentials (host, instance, vCenter)	Yes	No	No	No	No
View dashboard and all tabs	Yes	Yes	Yes	Yes	Yes
Start free trial	Yes	No	No	No	No
Initiate discovery of workloads	No	Yes	Yes	Yes	No
View license information	Yes	Yes	Yes	Yes	Yes
Activate license	Yes	No	No	No	No

Feature and action	Backup and recovery super admin	Backup admin	Restore admin	Clone admin	Viewer
View hosts	Yes	Yes	Yes	Yes	Yes
<b>Schedules:</b>					
Activate schedules	Yes	Yes	Yes	Yes	No
Suspend schedules	Yes	Yes	Yes	Yes	No
<b>Policies and protection:</b>					
View protection plans	Yes	Yes	Yes	Yes	Yes
Create, modify, or delete protection	Yes	Yes	No	No	No
Restore workloads	Yes	No	Yes	No	No
Create clone, split clone, or delete clone	Yes	No	No	Yes	No
Create, modify, or delete policy	Yes	Yes	No	No	No
<b>Reports:</b>					
View reports	Yes	Yes	Yes	Yes	Yes
Create reports	Yes	Yes	Yes	Yes	No
Delete reports	Yes	No	No	No	No
<b>Import from SnapCenter and manage host:</b>					
View imported SnapCenter data	Yes	Yes	Yes	Yes	Yes
Import data from SnapCenter	Yes	Yes	No	No	No
Manage (migrate) host	Yes	Yes	No	No	No
<b>Configure settings:</b>					
Configure log directory	Yes	Yes	Yes	No	No
Associate or remove instance credentials	Yes	Yes	Yes	No	No
<b>Buckets:</b>					

Feature and action	Backup and recovery super admin	Backup admin	Restore admin	Clone admin	Viewer
View buckets	Yes	Yes	Yes	Yes	Yes
Create, edit, or delete bucket	Yes	Yes	No	No	No

## Supported AWS archive storage tiers with BlueXP backup and recovery

BlueXP backup and recovery supports two S3 archival storage classes and most regions.

**NOTE** To switch to and from BlueXP backup and recovery UI versions, refer to [Switch to the previous BlueXP backup and recovery UI](#).

### 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. You can set this to "0" or to 1-999 days. If you set it to "0" days, you cannot change it later to 1-999 days.

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 on this page 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.](#)

### Restore 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
<b>S3 Glacier</b>	Fastest retrieval, highest cost	Slower retrieval, lower cost	Slowest retrieval, lowest cost
<b>S3 Glacier Deep Archive</b>		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
<b>S3 Glacier</b>	3-5 minutes	3-5 hours	5-12 hours
<b>S3 Glacier Deep Archive</b>		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](#).

## Supported Azure archive access tiers with BlueXP backup and recovery

BlueXP backup and recovery supports one Azure archival access tier and most regions.

**NOTE** To switch to and from BlueXP backup and recovery UI versions, refer to [Switch to the previous BlueXP backup and recovery UI](#).

## 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 section on this page 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.](#)

## Restore 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](#).

## Supported Google archive storage tiers with BlueXP backup and recovery

BlueXP backup and recovery supports one Google archival storage class and most regions.

**NOTE** To switch to and from BlueXP backup and recovery UI versions, refer to [Switch to the previous BlueXP backup and recovery UI](#).

## 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 on this page 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.](#)

## Restore 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.



# 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 © 2025 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.



# **BlueXP classification documentation**

## **BlueXP classification**

NetApp  
August 13, 2025

# Table of Contents

- BlueXP classification documentation . . . . . 1
- Release notes . . . . . 2
  - What’s new in BlueXP classification . . . . . 2
    - 14 July 2025 . . . . . 2
    - 10 June 2025 . . . . . 2
    - 12 May 2025 . . . . . 3
    - 14 April 2025 . . . . . 4
    - 10 March 2025 . . . . . 4
    - 19 February 2025 . . . . . 4
    - 22 January 2025 . . . . . 5
    - 16 December 2024 . . . . . 6
    - 4 November 2024 . . . . . 6
    - 10 October 2024 . . . . . 6
    - 2 September 2024 . . . . . 7
    - 05 August 2024 . . . . . 7
    - 01 July 2024 . . . . . 7
    - 05 June 2024 . . . . . 8
    - 15 May 2024 . . . . . 8
    - 01 April 2024 . . . . . 8
    - 04 March 2024 . . . . . 9
    - 10 January 2024 . . . . . 9
    - 14 December 2023 . . . . . 10
    - 06 November 2023 . . . . . 10
    - 04 October 2023 . . . . . 10
    - 05 September 2023 . . . . . 10
    - 17 July 2023 . . . . . 11
    - 06 June 2023 . . . . . 11
    - 03 April 2023 . . . . . 12
    - 07 March 2023 . . . . . 12
    - 05 February 2023 . . . . . 13
    - 09 January 2023 . . . . . 14
  - Known limitations in BlueXP classification . . . . . 15
    - BlueXP classification disabled options . . . . . 15
    - BlueXP classification scanning . . . . . 15
- Get started . . . . . 17
  - Learn about BlueXP classification . . . . . 17
    - Features . . . . . 17
    - Supported working environments and data sources . . . . . 18
    - Cost . . . . . 18
    - The BlueXP classification instance . . . . . 19
    - How BlueXP classification scanning works . . . . . 20
    - What’s the difference between Mapping and Classification scans . . . . . 21
    - Information that BlueXP classification categorizes . . . . . 21

Networking overview .....	22
User roles in BlueXP classification .....	22
Access BlueXP classification .....	22
Deploy BlueXP classification .....	23
Which BlueXP classification deployment should you use? .....	23
Deploy BlueXP classification in the cloud using BlueXP .....	23
Install BlueXP classification on a host that has internet access .....	33
Install BlueXP classification on a Linux host with no internet access .....	43
Check that your Linux host is ready to install BlueXP classification .....	52
Activate scanning on your data sources .....	57
Scan data sources overview with BlueXP classification .....	57
Scan Azure NetApp Files volumes with BlueXP classification .....	61
Scan Amazon FSx for ONTAP volumes with BlueXP classification .....	64
Scan Cloud Volumes ONTAP and on-premises ONTAP volumes with BlueXP classification .....	69
Scan database schemas with BlueXP classification .....	73
Scan file shares with BlueXP classification .....	75
Scan StorageGRID data with BlueXP classification .....	79
Integrate your Active Directory with BlueXP classification .....	81
Supported data sources .....	82
Connect to your Active Directory server .....	82
Manage your Active Directory integration .....	84
Use BlueXP classification .....	85
View governance details about the data stored in your organization with BlueXP classification .....	85
Review the Governance dashboard .....	85
Create the Data Discovery Assessment Report .....	88
Create the Data Mapping Overview Report .....	88
View compliance details about the private data stored in your organization with BlueXP classification .....	91
View files that contain personal data .....	92
View files that contain sensitive personal data .....	94
View files by categories .....	96
View files by file types .....	96
Categories of private data in BlueXP classification .....	97
Types of personal data .....	97
Types of sensitive personal data .....	102
Types of categories .....	103
Types of files .....	104
Accuracy of information found .....	104
Create a custom classification in BlueXP classification .....	105
Create a custom classification .....	105
Investigate the data stored in your organization with BlueXP classification .....	107
Filter data in the Data Investigation page .....	107
View file metadata .....	110
View users' permissions for files and directories .....	111
Check for duplicate files in your storage systems .....	112
Create the Data Investigation Report .....	113

Create a saved search based on selected filters . . . . .	116
Manage saved searches with BlueXP classification . . . . .	117
View saved searches results in the Investigation page . . . . .	117
Create custom saved searches . . . . .	117
Edit saved searches . . . . .	119
Delete saved searches . . . . .	121
Default searches . . . . .	121
Change the BlueXP classification scan settings for your repositories . . . . .	121
View the scan status for your repositories . . . . .	122
Change the type of scanning for a repository . . . . .	123
Prioritize scans . . . . .	124
Stop scanning for a repository . . . . .	124
Pause and resume scanning for a repository . . . . .	125
View BlueXP classification compliance reports . . . . .	126
Select the working environments for reports . . . . .	127
Data Subject Access Request Report . . . . .	127
Health Insurance Portability and Accountability Act (HIPAA) Report . . . . .	129
Payment Card Industry Data Security Standard (PCI DSS) Report . . . . .	130
Privacy Risk Assessment Report . . . . .	131
Manage BlueXP classification . . . . .	134
Exclude specific directories from BlueXP classification scans . . . . .	134
Supported data sources . . . . .	134
Define the directories to exclude from scanning . . . . .	134
Examples . . . . .	135
Escaping special characters in folder names . . . . .	136
View the current exclusion list . . . . .	137
Define additional group IDs as open to organization in BlueXP classification . . . . .	137
Add the "open to organization" permission to group IDs . . . . .	137
View the current list of group IDs . . . . .	138
Remove data sources from BlueXP classification . . . . .	138
Deactivate compliance scans for a working environment . . . . .	138
Remove a database from BlueXP classification . . . . .	138
Remove a group of file shares from BlueXP classification . . . . .	139
Uninstall BlueXP classification . . . . .	139
Uninstall BlueXP classification from a cloud deployment . . . . .	139
Uninstall BlueXP classification from an on-premises deployment . . . . .	140
Deprecated features . . . . .	141
BlueXP classification deprecated features . . . . .	141
Supported data sources . . . . .	141
Compliance features . . . . .	141
Features to manage your data . . . . .	142
Deploy BlueXP classification deprecations . . . . .	142
Install BlueXP classification on multiple hosts for large configurations with no internet access . . . . .	143
Scan data deprecations . . . . .	144
Scan Amazon S3 buckets with BlueXP classification . . . . .	144

Scan OneDrive accounts with BlueXP classification . . . . .	151
Scan SharePoint accounts with BlueXP classification . . . . .	155
Scan Google Drive accounts with BlueXP classification . . . . .	159
Scan StorageGRID data with BlueXP classification . . . . .	161
Manage data deprecations . . . . .	164
View governance details about your data using the BlueXP classification Governance dashboard . . . . .	164
Organize your private data with BlueXP classification . . . . .	166
Manage your private data with BlueXP classification . . . . .	174
Add personal data identifiers to your BlueXP classification scans . . . . .	185
View the status of your compliance actions in BlueXP classification . . . . .	200
Audit the history of BlueXP classification actions . . . . .	201
Reducing the BlueXP classification scan speed . . . . .	202
Remove a OneDrive, SharePoint, or Google Drive account from BlueXP classification . . . . .	203
Reference . . . . .	204
Supported BlueXP classification instance types . . . . .	204
AWS instance types . . . . .	204
Azure instance types . . . . .	204
GCP instance types . . . . .	205
Metadata collected from data sources in BlueXP classification . . . . .	205
Last access time timestamp . . . . .	205
Log in to the BlueXP classification system . . . . .	206
BlueXP classification APIs . . . . .	207
Overview . . . . .	207
Accessing the Swagger API reference . . . . .	208
Example using the APIs . . . . .	208
Knowledge and support . . . . .	218
Register for BlueXP support . . . . .	218
Support registration overview . . . . .	218
Register BlueXP for NetApp support . . . . .	218
Associate NSS credentials for Cloud Volumes ONTAP support . . . . .	220
Get help for BlueXP classification . . . . .	222
Get support for a cloud provider file service . . . . .	222
Use self-support options . . . . .	222
Create a case with NetApp support . . . . .	222
Manage your support cases (Preview) . . . . .	225
Frequently asked questions about BlueXP classification . . . . .	228
BlueXP classification service . . . . .	228
How does BlueXP classification work? . . . . .	228
Does BlueXP classification have a REST API, and does it work with third-party tools? . . . . .	228
Is BlueXP classification available through the cloud marketplaces? . . . . .	228
BlueXP classification scanning and analytics . . . . .	228
How often does BlueXP classification scan my data? . . . . .	228
Does scan performance vary? . . . . .	229
Can I search my data using BlueXP classification? . . . . .	229
BlueXP classification management and privacy . . . . .	229

How do I enable or disable BlueXP classification? .....	229
Can the service exclude scanning data in certain directories? .....	230
Are snapshots that reside on ONTAP volumes scanned? .....	230
What happens if data tiering is enabled on your ONTAP volumes? .....	230
Types of source systems and data types .....	230
Are there any restrictions when deployed in a Government region? .....	230
What data sources can I scan if I install BlueXP classification in a site without internet access? .....	230
Which file types are supported? .....	230
What kinds of data and metadata does BlueXP classification capture? .....	231
Can I limit BlueXP classification information to specific users? .....	231
Can anyone access the private data sent between my browser and BlueXP classification? .....	231
How is sensitive data handled? .....	231
Where is the data stored? .....	232
How is the data accessed? .....	232
Licenses and costs .....	232
How much does BlueXP classification cost? .....	232
Connector deployment .....	232
What is the Connector? .....	232
Where does the Connector need to be installed? .....	232
Does BlueXP classification require access to credentials? .....	232
Does communication between the service and the Connector use HTTP? .....	232
BlueXP classification deployment .....	233
What deployment models does BlueXP classification support? .....	233
What type of instance or VM is required for BlueXP classification? .....	233
Can I deploy the BlueXP classification on my own host? .....	233
What about secure sites without internet access? .....	233
Legal notices .....	234
Copyright .....	234
Trademarks .....	234
Patents .....	234
Privacy policy .....	234
Open source .....	234



# BlueXP classification documentation

# Release notes

## What's new in BlueXP classification

Learn what's new in BlueXP classification.

### 14 July 2025

#### Version 1.45

This BlueXP classification release includes code changes that optimize resource utilization and:

#### Improved workflow to add file shares for scanning

The workflow to add files shares to a file share group has been simplified. The process also now differentiates CIFS protocol support based on authentication type (Kerberos or NTLM).

For more information, see [Scan file shares](#).

#### Enhanced file owner information

You can now view more information about file owners for files captured in the Investigation tab. When viewing metadata for a file in the Investigation tab, locate the file owner then select **View details** to see the username, email, and SAM account name. You can also view other items owned by this user. This feature is only available for working environments with Active Directory.

For more information, see [Investigate the data stored in your organization](#).

### 10 June 2025

#### Version 1.44

This BlueXP classification release includes:

#### Improved update times for the Governance dashboard

Update times for individual components of the Governance dashboard have been improved. The following table displays the frequency of updates for each component.

Component	Update times
Age of Data	24 hours
Categories	24 hours
Data Overview	5 minutes
Duplicate Files	2 hours
File Types	24 hours
Non-Business Data	2 hours
Open Permissions	24 hours
Saved Searches	2 hours
Sensitive Data and Wide Permissions	24 hours

Component	Update times
Size of Data	24 hours
Stale Data	2 hours
Top Data Repositories by Sensitivity Level	2 hours

You can view the time of the last update and manually update the Duplicate Files, Non-Business Data, Saved Searches, Stale Data, and Top Data Repositories by Sensitivity Level components. For more information about the Governance dashboard, see [View governance details about the data stored in your organization](#).

**Performance and security improvements**

Enhancements have been made to improve BlueXP classification’s performance, memory consumption, and security.

**Bug fixes**

Redis has been upgraded to improve the reliability of BlueXP classification. BlueXP classification now uses Elasticsearch to improve the accuracy of file count reporting during scans.

**12 May 2025**

**Version 1.43**

This BlueXP classification release includes:

**Prioritize classification scans**

BlueXP classification supports the ability to prioritize Map & Classify scans in addition to Mapping-only scans, enabling you to select which scans are completed first. Prioritization of Map & Classify scans is supported during and before the scans begin. If you choose to prioritize a scan while it’s in progress, both the mapping and classification scans are prioritized.

For more information, see [Prioritize scans](#).

**Support for Canadian personally identifiable information (PII) data categories**

BlueXP classification scans identify Canadian PII data categories. These categories include banking information, passport numbers, social insurance numbers, driver’s license numbers and health card numbers for all Canadian provinces and territories.

For more information, see [Personal data categories](#).

**Custom classification (preview)**

BlueXP classification supports custom classifications for Map & Classify scans. With custom classifications, you can tailor BlueXP scans to capture data specific to your organization using regular expressions. This feature is currently in preview.

For more information, see [Add custom classifications](#).

**Saved searches tab**

The **Policies** tab has been renamed **Saved searches**. The functionality is unchanged.

**Send scan events to BlueXP timeline**

BlueXP classification supports sending classification events (when a scan is initiated and when it ends) to the [BlueXP timeline](#).

## Security updates

- The Keras package has been updated, mitigating vulnerabilities (BDSA-2025-0107 and BDSA-2025-1984).
- The Docker containers configuration has been updated. The container no longer has access to the host's network interfaces for crafting raw network packets. By reducing unnecessary access, the update mitigates potential security risks.

## Performance enhancements

Code enhancements have been implemented to reduce RAM usage and improve the overall performance of BlueXP classification.

## Bug fixes

Bugs that caused StorageGRID scans to fail, the investigation page filter options to not load, and the Data Discovery Assessment to not download for high volume assessments have been fixed.

## 14 April 2025

### Version 1.42

This BlueXP classification release includes:

#### Bulk scanning for working environments

BlueXP classification supports bulk operations for working environments. You can choose to enable Mapping scans, enable Map & Classify scans, disable scans, or create a custom configuration across volumes in working environment. If you make a selection for an individual volume, it overrides the bulk selection. To perform a bulk operation, navigate to the **Configuration** page and make your selection.

#### Download investigation report locally

BlueXP classification supports the ability to download data investigation reports locally to view in the browser. If you choose the local option, the data investigation is only available in the CSV format and only displays the first 10,000 rows of data.

For more information, see [Investigate the data stored in your organization with BlueXP classification](#).

## 10 March 2025

### Version 1.41

This BlueXP classification release includes general improvements and bug fixes. It also includes:

#### Scan status

BlueXP classification tracks the real time progress of the *initial* mapping and classification scans on a volume. Separate progressive bars track the mapping and classification scans, presenting a percentage of total files scanned. You can also hover over a progress bar to view the number of files scanned and the total files. Tracking the status of your scans creates deeper insights into the scan progress, enabling you to better plan your scans and understand resource allocation.

To view the status of your scans, navigate to **Configuration** in BlueXP classification then select the **Working Environment configuration**. Progress is displayed in line for each volume.

## 19 February 2025

## Version 1.40

This BlueXP classification release includes the following updates.

### Support for RHEL 9.5

This release provides support for Red Hat Enterprise Linux v9.5 in addition to previously supported versions. This is applicable to any manual on-premises installation of BlueXP classification, including dark site deployments.

The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater: Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4, and 9.5.

### Prioritize mapping-only scans

When conducting Mapping-only scans, you can prioritize the most important scans. This feature helps when you have many working environments and want to ensure high priority scans are completed first.

By default, scans are queued based on the order in which they are initiated. With the ability to prioritize scans, you can move scans to the front of the queue. Multiple scans can be prioritized. Priority is designated in a first-in, first-out order, meaning the first scan you prioritize moves to the front of the queue; the second scan you prioritize becomes second in the queue, and so forth.

Priority is granted on a one-time basis. Automatic rescans of mapping data occur in the default order.

Prioritization is limited to [mapping-only scans](#); it's not available for map and classify scans.

For more information, see [Prioritize scans](#).

### Retry all scans

BlueXP classification supports the ability to batch retry all failed scans.

You can reattempt scans in a batch operation with the **Retry all** function. If classification scans are failing due to a temporary issue such as a network outage, you can retry all scans at the same time with one button instead of retrying them individually. Scans can be retried as many times as needed.

To retry all scans:

1. From the BlueXP classification menu, select **Configuration**.
2. To retry all failed scans, select **Retry all scans**.

### Improved categorization model accuracy

The accuracy of the machine learning model for [predefined categories](#) has improved by 11%.

## 22 January 2025

### Version 1.39

This BlueXP classification release updates the export process for the Data Investigation report. This export update is useful for performing additional analyses on your data, creating additional visualizations on the data, or sharing the results of your data investigation with others.

Previously, the Data Investigation report export was limited to 10,000 rows. With this release, the limit has been removed so that you can export all of your data. This change enables you to export more data from your Data Investigation reports, providing you with more flexibility in your data analysis.

You can choose the working environment, volumes, destination folder, and either JSON or CSV format. The exported filename includes a timestamp to help you identify when the data was exported.

The supported working environments include:

- Cloud Volumes ONTAP
- FSx for ONTAP
- ONTAP
- Share group

Exporting data from the Data Investigation report has the following limitations:

- The maximum number of records to download is 500 million. per type (files, directories, and tables)
- One million records are expected to take about 35 minutes to export.

For details about data investigation and the report, see [Investigate data stored in your organization](#).

## 16 December 2024

### Version 1.38

This BlueXP classification release includes general improvements and bug fixes.

## 4 November 2024

### Version 1.37

This BlueXP classification release includes the following updates.

#### Support for RHEL 8.10

This release provides support for Red Hat Enterprise Linux v8.10 in addition to previously supported versions. This is applicable to any manual on-premises installation of BlueXP classification, including dark site deployments.

The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater: Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, and 9.4.

Learn more about [BlueXP classification](#).

#### Support for NFS v4.1

This release provides support for NFS v4.1 in addition to previously supported versions.

Learn more about [BlueXP classification](#).

## 10 October 2024

### Version 1.36

#### Support for RHEL 9.4

This release provides support for Red Hat Enterprise Linux v9.4 in addition to previously supported versions. This is applicable to any manual on-premises installation of BlueXP classification, including dark site

deployments.

The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater: Red Hat Enterprise Linux version 8.8, 9.0, 9.1, 9.2, 9.3, and 9.4.

Learn more about [BlueXP classification deployments overview](#).

### **Improved scan performance**

This release provides improved scan performance.

## **2 September 2024**

### **Version 1.35**

#### **Scan StorageGRID data**

BlueXP classification supports scanning data in StorageGRID.

For details, refer to [Scan StorageGRID data](#).

## **05 August 2024**

### **Version 1.34**

This BlueXP classification release includes the following update.

#### **Change from CentOS to Ubuntu**

BlueXP classification has updated its Linux operating system for Microsoft Azure and Google Cloud Platform (GCP) from CentOS 7.9 to Ubuntu 22.04.

For deployment details, refer to [Install on a Linux host with internet access and prepare the Linux host system](#).

## **01 July 2024**

### **Version 1.33**

#### **Ubuntu supported**

This release supports the Ubuntu 24.04 Linux platform.

#### **Mapping scans gather metadata**

The following metadata is extracted from files during mapping scans and is displayed on the Governance, Compliance, and Investigation dashboards:

- Working environment
- Working environment type
- Storage repository
- File type
- Used capacity
- Number of files
- File size

- File creation
- File last access
- File last modified
- File discovered time
- Permissions extraction

### **Additional data in dashboards**

This release updates which data appears in the Governance, Compliance, and Investigation dashboards during mapping scans.

For details, see [What's the difference between mapping and classification scans](#).

## **05 June 2024**

### **Version 1.32**

#### **New Mapping status column in the Configuration page**

This release now shows a new Mapping status column in the Configuration page. The new column helps you identify if the mapping is running, queued, paused or more.

For explanations of the statuses, see [Change scan settings](#).

## **15 May 2024**

### **Version 1.31**

#### **Classification is available as a core service within BlueXP**

BlueXP classification is now available as a core capability within BlueXP at no additional charge for up to 500 TiB of scanned data per connector. No Classification license or paid subscription is required. As we focus BlueXP classification functionality on scanning NetApp storage systems with this new version, some legacy functionality will only be available to customers who had previously paid for a license. The use of those legacy features will expire when the paid contract reaches its end date.



BlueXP classification does not impose a limit on the amount of data it can scan. Each Connector supports scanning and displaying 500 TiB of data. To scan more than 500 TiB of data, [install another Connector](#) then [deploy another classification instance](#).

The BlueXP UI displays data from a single connector. For tips on viewing data from multiple Connectors, see [Work with multiple Connectors](#).

[Learn more about the deprecated features](#).

## **01 April 2024**

### **Version 1.30**

#### **Support added for RHEL v8.8 and v9.3 BlueXP classification**

This release provides support for Red Hat Enterprise Linux v8.8 and v9.3 in addition to previously supported 9.x, which requires Podman, rather than the Docker engine. This is applicable to any manual on-premises installation of BlueXP classification.



The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater: Red Hat Enterprise Linux version 8.8, 9.0, 9.1, 9.2, and 9.3.

Learn more about [BlueXP classification deployments overview](#).

BlueXP classification is supported if you install the Connector on a RHEL 8 or 9 host that resides on-premises. It's not supported if the RHEL 8 or 9 host resides in AWS, Azure, or Google Cloud.

#### **Option to activate audit log collection removed**

The option to activate audit log collection has been disabled.

#### **Scan speed improved**

Scan performance on secondary scanner nodes has been improved. You can add more scanner nodes if you need additional processing power for your scans. For details, refer to [Install BlueXP classification on a host that has internet access](#).

#### **Automatic upgrades**

If you deployed BlueXP classification on a system with internet access, the system upgrades automatically. Previously, the upgrade occurred after a specific time elapsed since the last user activity. With this release, BlueXP classification upgrades automatically if the local time is between 1:00 AM and 5:00 AM. If the local time is outside of these hours, the upgrade occurs after a specific time elapses since the last user activity. For details, refer to [Install on a Linux host with internet access](#).

If you deployed BlueXP classification without internet access, you'll need to upgrade manually. For details, refer to [Install BlueXP classification on a Linux host with no internet access](#).

## **04 March 2024**

### **Version 1.29**

#### **Now you can exclude scanning data that resides in certain data source directories**

If you want BlueXP classification to exclude scanning data that resides in certain data source directories, you can add these directory names to a configuration file that BlueXP classification processes. This feature enables you to avoid scanning directories that are unnecessary, or that would result in returning false positive personal data results.

[Learn more](#).

#### **Extra Large instance support is now qualified**

If you need BlueXP classification to scan more than 250 million files, you can use an Extra Large instance in your cloud deployment or on-premises installation. This type of system can scan up to 500 million files.

[Learn more](#).

## **10 January 2024**

### **Version 1.27**

#### **Investigation page results display the total size in addition to total number of items**

The filtered results in the Investigation page display the total size of the items in addition to the total number of files. This can help when moving files, deleting files, and more.

#### **Configure additional Group IDs as "Open to Organization"**

Now you can configure Group IDs in NFS to be considered as "Open to Organization" directly from BlueXP classification if the group had not initially been set with that permission. Any files and folders that have these group IDs attached will show as "Open to Organization" in the Investigation Details page. See how to [add additional Group IDs as "open to organization"](#).

## 14 December 2023

### Version 1.26.6

This release included some minor enhancements.

The release also removed the following options:

- The option to activate audit log collection has been disabled.
- During the Directories investigation, the option to calculate the number of personal identifiable information (PII) data by Directories is not available. Refer to [Investigate the data stored in your organization](#).
- The option to integrate data using Azure Information Protection (AIP) labels has been disabled. Refer to [Organize your private data](#).

## 06 November 2023

### Version 1.26.3

The following issues have been fixed in this release

- Fixed an inconsistency when presenting the number of files scanned by the system in dashboards.
- Improved the scanning behavior by handling and reporting on files and directories with special characters in the name and metadata.

## 04 October 2023

### Version 1.26

#### Support for on-premises installations of BlueXP classification on RHEL version 9

Red Hat Enterprise Linux versions 8 and 9 do not support the Docker engine; which was required for the BlueXP classification installation. We now support BlueXP classification installation on RHEL 9.0, 9.1, and 9.2 using Podman version 4 or greater as the container infrastructure. If your environment requires using the newest versions of RHEL, now you can install BlueXP classification (version 1.26 or greater) when using Podman.

At this time we don't supported dark site installations or distributed scanning environments (using a master and remote scanner nodes) when using RHEL 9.x.

## 05 September 2023

### Version 1.25

#### Small and medium deployments temporarily unavailable

When you deploy an instance of BlueXP classification in AWS, the option to select **Deploy > Configuration** and choose a small or medium-sized instance is unavailable at this time. You can still deploy the instance using the large instance size by selecting **Deploy > Deploy**.

## **Apply tags on up to 100,000 items from the Investigation Results page**

In the past you could only apply tags to a single page at a time in the Investigation Results page (20 items). Now you can select **all** items in the Investigation Results pages and apply tags to all the items - up to 100,000 items at a time. [See how](#).

## **Identify duplicated files with a minimum file size of 1 MB**

BlueXP classification used to identify duplicated files only when files were 50 MB or larger. Now duplicated files starting with 1 MB can be identified. You can use the Investigation page filters "File Size" along with "Duplicates" to see which files of a certain size are duplicated in your environment.

## **17 July 2023**

### **Version 1.24**

#### **Two new types of German personal data are identified by BlueXP classification**

BlueXP classification can identify and categorize files that contain the following types of data:

- German ID (Personalausweisnummer)
- German Social Security Number (Sozialversicherungsnummer)

[See all the types of personal data that BlueXP classification can identify in your data.](#)

#### **BlueXP classification is fully supported in Restricted mode and Private mode**

BlueXP classification is now fully supported in sites with no internet access (Private mode) and with limited outbound internet access (Restricted mode). [Learn more about BlueXP deployment modes for the Connector.](#)

#### **Ability to skip versions when upgrading a Private mode installation of BlueXP classification**

Now you can upgrade to a newer version of BlueXP classification even if it is not sequential. This means that the current limitation of upgrading BlueXP classification by one version at a time is no longer required. This feature is relevant starting from version 1.24 onwards.

#### **The BlueXP classification API is now available**

The BlueXP classification API enables you to perform actions, create queries, and export information about the data you are scanning. The interactive documentation is available using Swagger. The documentation is separated into multiple categories, including Investigation, Compliance, Governance, and Configuration. Each category is a reference to the tabs in the BlueXP classification UI.

[Learn more about the BlueXP classification APIs.](#)

## **06 June 2023**

### **Version 1.23**

#### **Japanese is now supported when searching for data subject names**

Japanese names can now be entered when searching for a subject's name in response to a Data Subject Access Request (DSAR). You can generate a [Data Subject Access Request report](#) with the resulting information. You can also enter Japanese names in the "[Data Subject](#)" filter in the [Data Investigation page](#) to identify files that contain the subject's name.

#### **Ubuntu is now a supported Linux distribution on which you can install BlueXP classification**

Ubuntu 22.04 has been qualified as a supported operating system for BlueXP classification. You can install BlueXP classification on a Ubuntu Linux host in your network, or on a Linux host in the cloud when using

version 1.23 of the installer. [See how to install BlueXP classification on a host with Ubuntu installed.](#)

### **Red Hat Enterprise Linux 8.6 and 8.7 are no longer supported with new BlueXP classification installations**

These versions are not supported with new deployments because Red Hat no longer supports Docker, which is a prerequisite. If you have an existing BlueXP classification machine running on RHEL 8.6 or 8.7, NetApp will continue to support your configuration.

### **BlueXP classification can be configured as an FPolicy Collector to receive FPolicy events from ONTAP systems**

You can enable file access audit logs to be collected on your BlueXP classification system for file access events detected on volumes in your working environments. BlueXP classification can capture the following types of FPolicy events and the users who performed the actions on your files: Create, Read, Write, Delete, Rename, Change owner/permissions, and Change SACL/DAACL.

### **Data Sense BYOL licenses are now supported in dark sites**

Now you can upload your Data Sense BYOL license into the BlueXP digital wallet in a dark site so that you are notified when your license is getting low.

## **03 April 2023**

### **Version 1.22**

#### **New Data Discovery Assessment Report**

The Data Discovery Assessment Report provides a high-level analysis of your scanned environment to highlight the system's findings and to show areas of concern and potential remediation steps. The goal of this report is to raise awareness of data governance concerns, data security exposures, and data compliance gaps of your data set. [See how to generate and use the Data Discovery Assessment Report.](#)

#### **Ability to deploy BlueXP classification on smaller instances in the cloud**

When deploying BlueXP classification from a BlueXP Connector in an AWS environment, now you can select from two smaller instance types than what is available with the default instance. If you are scanning a small environment this can help you save on cloud costs. However, there are some restrictions when using the smaller instance. [See the available instance types and limitations.](#)

#### **Standalone script is now available to qualify your Linux system prior to BlueXP classification installation**

If you would like to verify that your Linux system meets all prerequisites independently of running the BlueXP classification installation, there is a separate script you can download that only tests for the prerequisites. [See how to check if your Linux host is ready to install BlueXP classification.](#)

## **07 March 2023**

### **Version 1.21**

#### **New functionality to add your own custom categories from the BlueXP classification UI**

BlueXP classification now enables you to add your own custom categories so that BlueXP classification will identify the files that fit into those categories. BlueXP classification has many [predefined categories](#), so this feature enables you to add custom categories to identify where information that is unique to your organization are found in your data.

[Learn more.](#)

#### **Now you can add custom keywords from the BlueXP classification UI**

BlueXP classification has had the ability to add custom keywords that BlueXP classification will identify in future scans for a while. However, you needed to log into the BlueXP classification Linux host and use a command line interface to add the keywords. In this release, the ability to add custom keywords is in the BlueXP classification UI, making it very easy to add and edit these keywords.

[Learn more about adding custom keywords from the BlueXP classification UI.](#)

### **Ability to have BlueXP classification not scan files when the "last access time" will be changed**

By default, if BlueXP classification doesn't have adequate "write" permissions, the system won't scan files in your volumes because BlueXP classification can't revert the "last access time" to the original timestamp. However, if you don't care if the last access time is reset to the original time in your files, you can override this behavior in the Configuration page so that BlueXP classification will scan the volumes regardless of permissions.

In conjunction with this capability, a new filter named "Scan Analysis Event" has been added so you can view the files that were not classified because BlueXP classification couldn't revert last accessed time, or the files that were classified even though BlueXP classification couldn't revert last accessed time.

[Learn more about the "Last access time timestamp" and the permissions BlueXP classification requires.](#)

### **Three new types of personal data are identified by BlueXP classification**

BlueXP classification can identify and categorize files that contain the following types of data:

- Botswana Identity Card (Omang) Number
- Botswana Passport Number
- Singapore National Registration Identity Card (NRIC)

[See all the types of personal data that BlueXP classification can identify in your data.](#)

### **Updated functionality for directories**

- The "Light CSV Report" option for Data Investigation Reports now includes information from directories.
- The "Last Accessed" time filter now shows the last accessed time for both files and directories.

### **Installation enhancements**

- The BlueXP classification installer for sites without internet access (dark sites) now performs a pre-check to make sure your system and networking requirements are in place for a successful installation.
- Installation audit log files are saved now; they are written to `/ops/netapp/install_logs`.

## **05 February 2023**

### **Version 1.20**

#### **Ability to send Policy-based notification emails to any email address**

In earlier versions of BlueXP classification you could send email alerts to the BlueXP users in your account when certain critical Policies return results. This feature enables you to get notifications to protect your data when you're not online. Now you can also send email alerts from Policies to any other users - up to 20 email addresses - who are not in your BlueXP account.

[Learn more about sending email alerts based on Policy results.](#)

#### **Now you can add personal patterns from the BlueXP classification UI**

BlueXP classification has had the ability to add custom "personal data" that BlueXP classification will identify in future scans for a while. However, you needed to log into the BlueXP classification Linux host and use a command line to add the custom patterns. In this release, the ability to add personal patterns using a regex is in the BlueXP classification UI, making it very easy to add and edit these custom patterns.

[Learn more about adding custom patterns from the BlueXP classification UI.](#)

### **Ability to move 15 million files using BlueXP classification**

In the past you could have BlueXP classification move a maximum of 100,000 source files to any NFS share. Now you can move up to 15 million files at a time. [Learn more about moving source files using BlueXP classification.](#)

### **Ability to see the number of users who have access to SharePoint Online files**

The filter "Number of users with access" now supports files stored in SharePoint Online repositories. In the past only files on CIFS shares were supported. Note that SharePoint groups that are not active directory based will not be counted in this filter at this time.

### **New "Partial Success" status has been added to the Action Status panel**

The new "Partial Success" status indicates that a BlueXP classification action is finished and some items failed and some items succeeded, for example, when you are moving or deleting 100 files. Additionally, the "Finished" status has been renamed to "Success". In the past, the "Finished" status might list actions that succeeded and that failed. Now the "Success" status means that all actions succeeded on all items. [See how to view the Actions Status panel.](#)

## **09 January 2023**

### **Version 1.19**

#### **Ability to view a chart of files that contain sensitive data and that are overly permissive**

The Governance dashboard has added a new *Sensitive Data and Wide Permissions* area that provides a heatmap of files that contain sensitive data (including both sensitive and sensitive personal data) and that are overly permissive. This can help you to see where you may have some risks with sensitive data. [Learn more.](#)

#### **Three new filters are available in the Data Investigation page**

New filters are available to refine the results that display in the Data Investigation page:

- The "Number of users with access" filter shows which files and folders are open to a certain number of users. You can choose a number range to refine the results - for example, to see which files are accessible by 51-100 users.
- The "Created Time", "Discovered Time", "Last Modified", and "Last Accessed" filters now allow you to create a custom date range instead of just selecting a pre-defined range of days. For example, you can look for files with a "Created Time" "older than 6 months", or with a "Last Modified" date within the "last 10 days".
- The "File Path" filter now enables you to specify paths that you want to exclude from the filtered query results. If you enter paths to both include and exclude certain data, BlueXP classification finds all files in the included paths first, then it removes files from excluded paths, and then it displays the results.

[See the list of all the filters you can use to investigate your data.](#)

#### **BlueXP classification can identify the Japanese Individual Number**

BlueXP classification can identify and categorize files that contain the Japanese Individual Number (also known as My Number). This includes both the Personal and Corporate My Number. [See all the types of](#)

personal data that BlueXP classification can identify in your data.

## Known limitations in BlueXP classification

Known limitations identify functions that are not supported or do not interoperate correctly in this release. Review these limitations carefully.

### BlueXP classification disabled options

The December 2023 (version 1.26.6) release removed the following options:

- The option to activate audit log collection has been disabled.
- During the Directories investigation, the option to calculate the number of personally identifiable information (PII) data by Directories is not available.
- The option to integrate data using Azure Information Protection (AIP) labels has been disabled.

### BlueXP classification scanning

The following limitations occur with BlueXP classifications scans.

#### BlueXP classification scans only one share under a volume

If you have multiple file shares under a single volume, BlueXP classification scans the share with the highest hierarchy. For example, if you have shares like the following:

- /A
- /A/B
- /C
- /D/E

In this configuration, only the data in /A is scanned. The data in /C and /D is not scanned.

#### Workaround

There is a workaround to make sure you are scanning data from all the shares in your volume. Follow these steps:

1. In the working environment, add the volume to be scanned.
2. After BlueXP classification has completed scanning the volume, go to the *Data Investigation* page and create a filter to see which share is being scanned:

Filter the data by "Working Environment Name" and "Directory Type = Share" to see which share is being scanned.

3. Get the complete list of shares that exist in the volume so you can see which shares are not being scanned.
4. [Add the remaining shares to a share group.](#)

Add all the shares individually, for example:

/C

/D

5. Perform these steps for each volume in the working environment that has multiple shares.

### **Last accessed timestamp**

When BlueXP classification conducts a scan of a directory, the scan impacts the directory's **Last accessed** field. When you view the **Last accessed** field, that metadata reflects either the date and time of the scan or the last time a user accessed the directory.



# Get started

## Learn about BlueXP classification

BlueXP classification (Cloud Data Sense) is a data governance service for BlueXP that scans your corporate on-premises and cloud data sources to map and classify data, and to identify private information. This can help reduce your security and compliance risk, decrease storage costs, and assist with your data migration projects.



Beginning with version 1.31, BlueXP classification is available as a core capability with BlueXP. There's no additional charge. No Classification license or subscription is required. If you've been using legacy version 1.30 or earlier, that version is available until your subscription expires. [See a list of deprecated features.](#)

### Features

BlueXP classification uses artificial intelligence (AI), natural language processing (NLP), and machine learning (ML) to understand the content that it scans in order to extract entities and categorize the content accordingly. This allows BlueXP classification to provide the following areas of functionality.

[Learn more about the use cases for BlueXP classification.](#)

#### Maintain compliance

BlueXP classification provides several tools that can help with your compliance efforts. You can use BlueXP classification to:

- Identify Personal Identifiable Information (PII).
- Identify a wide scope of sensitive personal information as required by GDPR, CCPA, PCI, and HIPAA privacy regulations.
- Respond to Data Subject Access Requests (DSAR) based on name or email address.

#### Strengthen security

BlueXP classification can identify data that is potentially at risk for being accessed for criminal purposes. You can use BlueXP classification to:

- Identify all the files and directories (shares and folders) with open permissions that are exposed to your entire organization or to the public.
- Identify sensitive data that resides outside of the initial, dedicated location.
- Comply with data retention policies.
- Use *Policies* to automatically detect new security issues so security staff can take action immediately.

#### Optimize storage usage

BlueXP classification provides tools that can help with your storage total cost of ownership (TCO). You can use BlueXP classification to:

- Increase storage efficiency by identifying duplicate or non-business-related data.
- Save storage costs by identifying inactive data that you can tier to less expensive object storage. [Learn more about tiering from Cloud Volumes ONTAP systems.](#) [Learn more about tiering from on-premises](#)

ONTAP systems.

## Supported working environments and data sources

BlueXP classification can scan and analyze structured and unstructured data from the following types of working environments and data sources:

### Working environments

- Amazon FSx for ONTAP
- Azure NetApp Files
- Cloud Volumes ONTAP (deployed in AWS, Azure, or GCP)
- On-premises ONTAP clusters
- StorageGRID

### Data sources

- NetApp file shares
- Databases:
  - Amazon Relational Database Service (Amazon RDS)
  - MongoDB
  - MySQL
  - Oracle
  - PostgreSQL
  - SAP HANA
  - SQL Server (MSSQL)

BlueXP classification supports NFS versions 3.x, 4.0, and 4.1, and CIFS versions 1.x, 2.0, 2.1, and 3.0.

## Cost

BlueXP classification is free to use. No Classification license or paid subscription is required.

### Infrastructure costs

- Installing BlueXP classification in the cloud requires deploying a cloud instance, which results in charges from the cloud provider where it is deployed. See [the type of instance that is deployed for each cloud provider](#). There is no cost if you install BlueXP classification on an on-premises system.
- BlueXP classification requires that you have deployed a BlueXP Connector. In many cases you already have a Connector because of other storage and services you are using in BlueXP. The Connector instance results in charges from the cloud provider where it is deployed. See the [type of instance that is deployed for each cloud provider](#). There is no cost if you install the Connector on an on-premises system.

### Data transfer costs

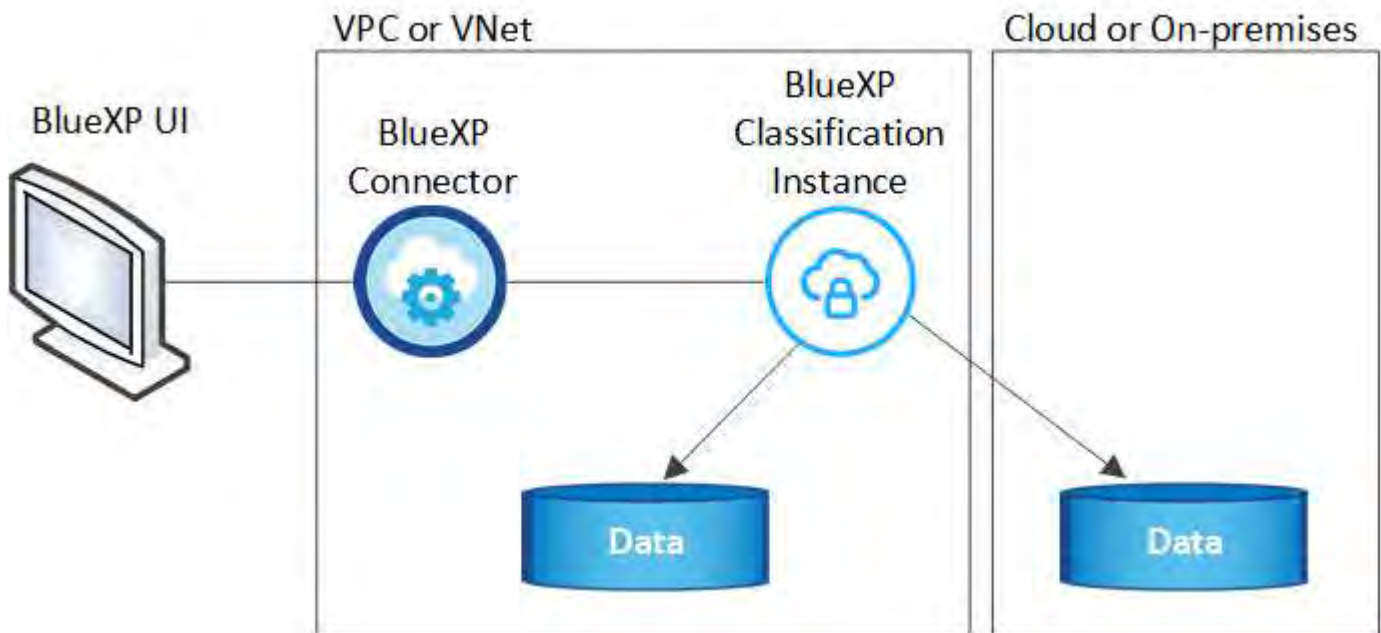
Data transfer costs depend on your setup. If the BlueXP classification instance and data source are in the same Availability Zone and region, then there are no data transfer costs. But if the data source, such as a Cloud Volumes ONTAP system, is in a *different* Availability Zone or region, then you'll be charged by your

cloud provider for data transfer costs. See these links for more details:

- [AWS: Amazon Elastic Compute Cloud \(Amazon EC2\) Pricing](#)
- [Microsoft Azure: Bandwidth Pricing Details](#)
- [Google Cloud: Storage Transfer Service pricing](#)

## The BlueXP classification instance

When you deploy BlueXP classification in the cloud, BlueXP deploys the instance in the same subnet as the Connector. [Learn more about Connectors.](#)



Note the following about the default instance:

- In AWS, BlueXP classification runs on an [m6i.4xlarge instance](#) with a 500 GiB GP2 disk. The operating system image is Amazon Linux 2. When deployed in AWS, you can choose a smaller instance size if you are scanning a small amount of data.
- In Azure, BlueXP classification runs on a [Standard\\_D16s\\_v3 VM](#) with a 500 GiB disk. The operating system image is Ubuntu 22.04.
- In GCP, BlueXP classification runs on an [n2-standard-16 VM](#) with a 500 GiB Standard persistent disk. The operating system image is Ubuntu 22.04.
- In regions where the default instance isn't available, BlueXP classification runs on an alternate instance. [See the alternate instance types.](#)
- The instance is named *CloudCompliance* with a generated hash (UUID) concatenated to it. For example: *CloudCompliance-16bb6564-38ad-4080-9a92-36f5fd2f71c7*
- Only one BlueXP classification instance is deployed per Connector.

You can also deploy BlueXP classification on a Linux host on your premises or on a host in your preferred cloud provider. The software functions exactly the same way regardless of which installation method you choose. Upgrades of BlueXP classification software are automated as long as the instance has internet access.



The instance should remain running at all times because BlueXP classification continuously scans the data.

## Deploy on different instance types

Review the following specifications for instance types:

System size	Specs	Limitations
Extra Large	32 CPUs, 128 GB RAM, 1 TiB SSD	Can scan up to 500 million files.
Large (default)	16 CPUs, 64 GB RAM, 500 GiB SSD	Can scan up to 250 million files.

When deploying BlueXP classification in Azure or GCP, email [ng-contact-data-sense@netapp.com](mailto:ng-contact-data-sense@netapp.com) for assistance if you want to use a smaller instance type.


## How BlueXP classification scanning works

At a high-level, BlueXP classification scanning works like this:

1. You deploy an instance of BlueXP classification in BlueXP.
2. You enable high-level mapping (called *Mapping only* scans) or deep-level scanning (called *Map & Classify* scans) on one or more data sources.
3. BlueXP classification scans the data using an AI learning process.
4. You use the provided dashboards and reporting tools to help in your compliance and governance efforts.

After you enable BlueXP classification and select the repositories that you want to scan (these are the volumes, database schemas, or other user data), it immediately starts scanning the data to identify personal and sensitive data. You should focus on scanning live production data in most cases instead of backups, mirrors, or DR sites. Then BlueXP classification maps your organizational data, categorizes each file, and identifies and extracts entities and predefined patterns in the data. The result of the scan is an index of personal information, sensitive personal information, data categories, and file types.

BlueXP classification connects to the data like any other client by mounting NFS and CIFS volumes. NFS volumes are automatically accessed as read-only, while you need to provide Active Directory credentials to scan CIFS volumes.



After the initial scan, BlueXP classification continuously scans your data in a round-robin fashion to detect incremental changes. This is why it's important to keep the instance running.

You can enable and disable scans at the volume level or the database schema level.



BlueXP classification does not impose a limit on the amount of data it can scan. Each Connector supports scanning and displaying 500 TiB of data. To scan more than 500 TiB of data, [install another Connector](#) then [deploy another classification instance](#).

The BlueXP UI displays data from a single connector. For tips on viewing data from multiple Connectors, see [Work with multiple Connectors](#).

## What's the difference between Mapping and Classification scans

You can conduct two types of scans in BlueXP classification:

- **Mapping-only scans** provide only a high-level overview of your data and are performed on selected data sources. Mapping-only scans take less time than map and classify scans because they do not access files to see the data inside. You might want to do this initially to identify areas of research and then perform a Map & Classify scan on those areas.
- **Map & Classify scans** provide deep-level scanning of your data.

For details about the differences between Mapping and Classification scans, see [What's the difference between Mapping and Classification scans?](#)

## Information that BlueXP classification categorizes

BlueXP classification collects, indexes, and assigns categories to the following data:

- **Standard metadata** about files: the file type, its size, creation and modification dates, and so on.
- **Personal data:** Personally identifiable information (PII) such as email addresses, identification numbers, or credit card numbers, which BlueXP classification identifies using specific words, strings, and patterns in the files. [Learn more about personal data](#).
- **Sensitive personal data:** Special types of sensitive personal information (SPII), such as health data, ethnic origin, or political opinions, as defined by General Data Protection Regulation (GDPR) and other privacy regulations. [Learn more about sensitive personal data](#).

- **Categories:** BlueXP classification takes the data that it scanned and divides it into different types of categories. Categories are topics based on AI analysis of the content and metadata of each file. [Learn more about categories.](#)
- **Types:** BlueXP classification takes the data that it scanned and breaks it down by file type. [Learn more about types.](#)
- **Name entity recognition:** BlueXP classification uses AI to extract people's natural names from documents. [Learn about responding to Data Subject Access Requests.](#)

## Networking overview

BlueXP classification deploys a single server, or cluster, wherever you choose — in the cloud or on premises. The servers connect via standard protocols to the data sources and index the findings in an Elasticsearch cluster, which is also deployed on the same servers. This enables support for multi-cloud, cross-cloud, private cloud, and on-premises environments.

BlueXP deploys the BlueXP classification instance with a security group that enables inbound HTTP connections from the Connector instance.

When you use BlueXP in SaaS mode, the connection to BlueXP is served over HTTPS, and the private data sent between your browser and the BlueXP classification instance are secured with end-to-end encryption using TLS 1.2, which means NetApp and third parties can't read it.

Outbound rules are completely open. Internet access is needed to install and upgrade the BlueXP classification software and to send usage metrics.

If you have strict networking requirements, [learn about the endpoints that BlueXP classification contacts.](#)

## User roles in BlueXP classification

The role each user has been assigned provides different capabilities within BlueXP and within BlueXP classification. For details, refer to [BlueXP IAM roles](#) (when using BlueXP in standard mode).

## Access BlueXP classification

You can access the BlueXP classification service through NetApp BlueXP.

To sign in to BlueXP, you can use your NetApp Support Site credentials or you can sign up for a NetApp cloud login using your email and a password. [Learn more about logging in to BlueXP.](#)

Specific tasks require specific BlueXP user roles. [Learn about BlueXP access roles for all services.](#)

### Before you begin

- [You should add a BlueXP Connector.](#)
- [Understand which BlueXP classification deployment style suits your workload.](#)

### Steps

1. In a web browser, navigate to the [BlueXP console](#).

The NetApp BlueXP login page appears.

2. Sign in to BlueXP.

3. From the BlueXP left navigation menu, select **Governance > Classification**.
4. If this is your first time accessing BlueXP classification, the landing page appears.

Select **Deploy Classification On-Premises or Cloud** to begin deploying your classification instance. For more information, see [Which BlueXP classification deployment should you use?](#)

[A screenshot of selecting the button to activate BlueXP classification.]

Otherwise, the BlueXP classification Dashboard appears.

## Deploy BlueXP classification

### Which BlueXP classification deployment should you use?

You can deploy BlueXP classification in different ways. Learn which method meets your needs.

BlueXP classification can be deployed in the following ways:

- [Deploy in the cloud using BlueXP](#). BlueXP will deploy the BlueXP classification instance in the same cloud provider network as the BlueXP Connector.
- [Install on a Linux host with internet access](#). Install BlueXP classification on a Linux host in your network, or on a Linux host in the cloud, that has internet access. This type of installation may be a good option if you prefer to scan on-premises ONTAP systems using a BlueXP classification instance that's also located on premises—but this is not a requirement.
- [Install on a Linux host in an on-premises site without internet access](#), also known as *private mode*. This type of installation, which uses an installation script, has no connectivity to the BlueXP SaaS layer.

Both the installation on a Linux host with internet access and the on-premises installation on a Linux host without internet access use an installation script. The script starts by checking if the system and environment meet the prerequisites. If the prerequisites are met, the installation starts. If you would like to verify the prerequisites independently of running the BlueXP classification installation, there is a separate software package you can download that only tests for the prerequisites.

Refer to [Check that your Linux host is ready to install BlueXP classification](#).

### Deploy BlueXP classification in the cloud using BlueXP

Complete a few steps to deploy BlueXP classification in the cloud. BlueXP will deploy the BlueXP classification instance in the same cloud provider network as the BlueXP Connector.

Note that you can also [install BlueXP classification on a Linux host that has internet access](#). This type of installation may be a good option if you prefer to scan on-premises ONTAP systems using a BlueXP classification instance that's also located on premises — but this is not a requirement. The software functions exactly the same way regardless of which installation method you choose.

#### Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.



**1**

### Create a Connector

If you don't already have a Connector, create a Connector now. See [creating a Connector in AWS](#), [creating a Connector in Azure](#), or [creating a Connector in GCP](#).

You can also [install the Connector on-premises](#) on a Linux host in your network or on a Linux host in the cloud.

**2**

### Review prerequisites

Ensure that your environment can meet the prerequisites. This includes outbound internet access for the instance, connectivity between the Connector and BlueXP classification over port 443, and more. [See the complete list](#).

**3**

### Deploy BlueXP classification

Launch the installation wizard to deploy the BlueXP classification instance in the cloud.

## Create a Connector

If you don't already have a Connector, create a Connector in your cloud provider. See [creating a Connector in AWS](#) or [creating a Connector in Azure](#), or [creating a Connector in GCP](#). In most cases you will probably have a Connector set up before you attempt to activate BlueXP classification because most [BlueXP features require a Connector](#), but there are cases where you'll need to set one up now.

There are some scenarios where you have to use a Connector that's deployed in a specific cloud provider:

- When scanning data in Cloud Volumes ONTAP in AWS or Amazon FSx for ONTAP buckets, you use a Connector in AWS.
- When scanning data in Cloud Volumes ONTAP in Azure or in Azure NetApp Files, you use a Connector in Azure.
  - For Azure NetApp Files, it must be deployed in the same region as the volumes you wish to scan.
- When scanning data in Cloud Volumes ONTAP in GCP, you use a Connector in GCP.

On-prem ONTAP systems, NetApp file shares, and databases can be scanned when using any of these cloud Connectors.

Note that you can also [install the Connector on-premises](#) on a Linux host in your network or in the cloud. Some users planning to install BlueXP classification on-prem may also choose to install the Connector on-prem.

As you can see, there may be some situations where you need to use [multiple Connectors](#).



BlueXP classification does not impose a limit on the amount of data it can scan. Each Connector supports scanning and displaying 500 TiB of data. To scan more than 500 TiB of data, [install another Connector](#) then [deploy another classification instance](#). The BlueXP UI displays data from a single connector. For tips on viewing data from multiple Connectors, see [Work with multiple Connectors](#).

## Government region support

BlueXP classification is supported when the Connector is deployed in a Government region (AWS GovCloud,



Azure Gov, or Azure DoD). When deployed in this manner, BlueXP classification has the following restrictions:

[See more information about deploying the Connector in a Government region.](#)

### **Review prerequisites**

Review the following prerequisites to make sure that you have a supported configuration before you deploy BlueXP classification in the cloud. When you deploy BlueXP classification in the cloud, it's located in the same subnet as the Connector.

### **Enable outbound internet access from BlueXP classification**

BlueXP classification requires outbound internet access. If your virtual or physical network uses a proxy server for internet access, ensure that the BlueXP classification instance has outbound internet access to contact the following endpoints. The proxy must be non-transparent. Transparent proxies are not currently supported.

Review the appropriate table below depending on whether you are deploying BlueXP classification in AWS, Azure, or GCP.

### Required endpoints for AWS

Endpoints	Purpose
<a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a>	Communication with the BlueXP service, which includes NetApp accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://auth0.com">https://auth0.com</a>	Communication with the BlueXP website for centralized user authentication.
<a href="https://cloud-compliance-support-netapp.s3.us-west-2.amazonaws.com">https://cloud-compliance-support-netapp.s3.us-west-2.amazonaws.com</a> <a href="https://hub.docker.com">https://hub.docker.com</a> <a href="https://auth.docker.io">https://auth.docker.io</a> <a href="https://registry-1.docker.io">https://registry-1.docker.io</a> <a href="https://index.docker.io/">https://index.docker.io/</a> <a href="https://dseasb33srrn.cloudfront.net/">https://dseasb33srrn.cloudfront.net/</a> <a href="https://production.cloudflare.docker.com/">https://production.cloudflare.docker.com/</a>	Provides access to software images, manifests, and templates.
<a href="https://kinesis.us-east-1.amazonaws.com">https://kinesis.us-east-1.amazonaws.com</a>	Enables NetApp to stream data from audit records.
<a href="https://cognito-idp.us-east-1.amazonaws.com">https://cognito-idp.us-east-1.amazonaws.com</a> <a href="https://cognito-identity.us-east-1.amazonaws.com">https://cognito-identity.us-east-1.amazonaws.com</a> <a href="https://user-feedback-store-prod.s3.us-west-2.amazonaws.com">https://user-feedback-store-prod.s3.us-west-2.amazonaws.com</a> <a href="https://customer-data-production.s3.us-west-2.amazonaws.com">https://customer-data-production.s3.us-west-2.amazonaws.com</a>	Enables BlueXP classification to access and download manifests and templates, and to send logs and metrics.

### Required endpoints for Azure

Endpoints	Purpose
<a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a>	Communication with the BlueXP service, which includes NetApp accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://auth0.com">https://auth0.com</a>	Communication with the BlueXP website for centralized user authentication.
<a href="https://support.compliance.api.bluexp.netapp.com/">https://support.compliance.api.bluexp.netapp.com/</a> <a href="https://hub.docker.com">https://hub.docker.com</a> <a href="https://auth.docker.io">https://auth.docker.io</a> <a href="https://registry-1.docker.io">https://registry-1.docker.io</a> <a href="https://index.docker.io/">https://index.docker.io/</a> <a href="https://dseasb33srrn.cloudfront.net/">https://dseasb33srrn.cloudfront.net/</a> <a href="https://production.cloudflare.docker.com/">https://production.cloudflare.docker.com/</a>	Provides access to software images, manifests, templates, and to send logs and metrics.
<a href="https://support.compliance.api.bluexp.netapp.com/">https://support.compliance.api.bluexp.netapp.com/</a>	Enables NetApp to stream data from audit records.

### Required endpoints for GCP

Endpoints	Purpose
<a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a>	Communication with the BlueXP service, which includes NetApp accounts.

Endpoints	Purpose
https://netapp-cloud-account.auth0.com https://auth0.com	Communication with the BlueXP website for centralized user authentication.
https://support.compliance.api.bluexp.netapp.com/ https://hub.docker.com https://auth.docker.io https://registry-1.docker.io https://index.docker.io/ https://dseasb33srrn.cloudfront.net/ https://production.cloudflare.docker.com/	Provides access to software images, manifests, templates, and to send logs and metrics.
https://support.compliance.api.bluexp.netapp.com/	Enables NetApp to stream data from audit records.

### Ensure that BlueXP has the required permissions

Ensure that BlueXP has permissions to deploy resources and create security groups for the BlueXP classification instance. You can find the latest BlueXP permissions in [the policies provided by NetApp](#).

### Ensure that the BlueXP Connector can access BlueXP classification

Ensure connectivity between the Connector and the BlueXP classification instance. The security group for the Connector must allow inbound and outbound traffic over port 443 to and from the BlueXP classification instance. This connection enables deployment of the BlueXP classification instance and enables you to view information in the Compliance and Governance tabs. BlueXP classification is supported in Government regions in AWS and Azure.

Additional inbound and outbound security group rules are required for AWS and AWS GovCloud deployments. See [Rules for the Connector in AWS](#) for details.

Additional inbound and outbound security group rules are required for Azure and Azure Government deployments. See [Rules for the Connector in Azure](#) for details.

### Ensure that you can keep BlueXP classification running

The BlueXP classification instance needs to stay on to continuously scan your data.

### Ensure web browser connectivity to BlueXP classification

After BlueXP classification is enabled, ensure that users access the BlueXP interface from a host that has a connection to the BlueXP classification instance.

The BlueXP classification instance uses a private IP address to ensure that the indexed data isn't accessible to the internet. As a result, the web browser that you use to access BlueXP must have a connection to that private IP address. That connection can come from a direct connection to your cloud provider (for example, a VPN), or from a host that's inside the same network as the BlueXP classification instance.

### Check your vCPU limits

Ensure that your cloud provider's vCPU limit allows for the deployment of an instance with the necessary number of cores. You'll need to verify the vCPU limit for the relevant instance family in the region where BlueXP is running. See [the required instance types](#).

See the following links for more details on vCPU limits:

- [AWS documentation: Amazon EC2 service quotas](#)
- [Azure documentation: Virtual machine vCPU quotas](#)
- [Google Cloud documentation: Resource quotas](#)

### **Deploy BlueXP classification in the cloud**

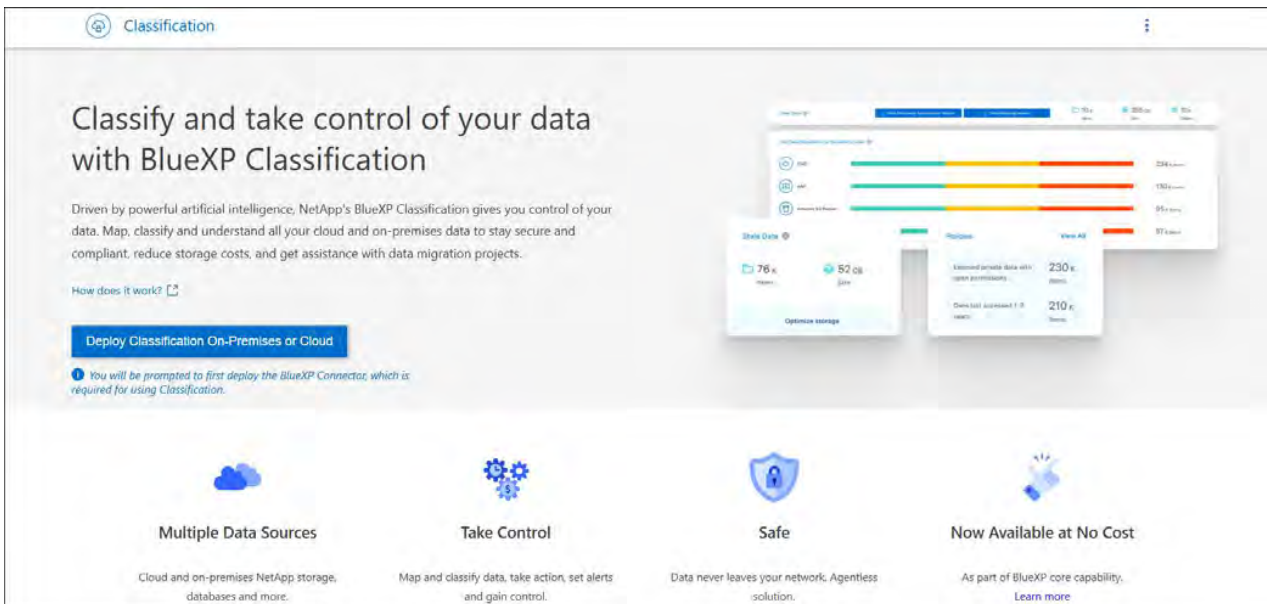
Follow these steps to deploy an instance of BlueXP classification in the cloud. The Connector will deploy the instance in the cloud, and then install BlueXP classification software on that instance.

In regions where the default instance type isn't available, BlueXP classification runs on an [alternate instance type](#).

## Deploy in AWS

### Steps

1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. Select **Deploy Classification On-Premises or Cloud**.



3. From the *Installation* page, select **Deploy > Deploy** to use the "Large" instance size and start the cloud deployment wizard.
4. The wizard displays progress as it goes through the deployment steps. It will stop and prompt for input if it runs into any issues.



5. When the instance is deployed and BlueXP classification is installed, select **Continue to configuration** to go to the *Configuration* page.

## Deploy in Azure

### Steps

1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. Select **Deploy Classification On-Premises or Cloud**.

**Classification**

## Classify and take control of your data with BlueXP Classification

Driven by powerful artificial intelligence, NetApp's BlueXP Classification gives you control of your data. Map, classify and understand all your cloud and on-premises data to stay secure and compliant, reduce storage costs, and get assistance with data migration projects.

How does it work? [🔗](#)

[Deploy Classification On-Premises or Cloud](#)

**!** You will be prompted to first deploy the BlueXP Connector, which is required for using Classification.

- Multiple Data Sources**  
Cloud and on-premises NetApp storage, databases and more.
- Take Control**  
Map and classify data, take action, set alerts and gain control.
- Safe**  
Data never leaves your network. Agentless solution.
- Now Available at No Cost**  
As part of BlueXP core capability. [Learn more](#)

3. Select **Deploy** to start the cloud deployment wizard.

## Install your Data Sense instance

Select your preferred deployment location:

[Learn more about deploying Data Sense](#)

### Cloud Environment

- I want BlueXP to deploy the instance and install Data Sense** [Deploy](#)
- I deployed an instance and I'm ready to install Data Sense** [Deploy](#)

### On Premise

- I prepared a local machine and I'm ready to install Data Sense** [Deploy](#)

4. The wizard displays progress as it goes through the deployment steps. It will stop and prompt for input if it runs into any issues.

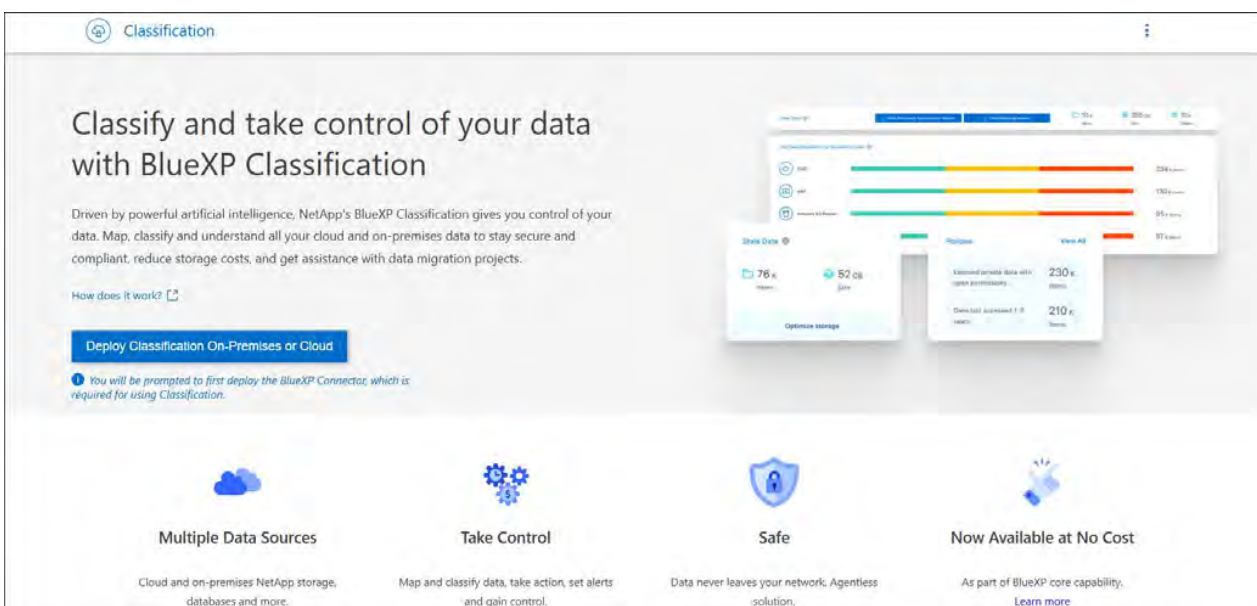


5. When the instance is deployed and BlueXP classification is installed, select **Continue to configuration** to go to the *Configuration* page.

## Deploy in Google Cloud


### Steps

1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. Select **Deploy Classification On-Premises or Cloud**.



3. Select **Deploy** to start the cloud deployment wizard.





4. The wizard displays progress as it goes through the deployment steps. It will stop and prompt for input if it runs into any issues.



5. When the instance is deployed and BlueXP classification is installed, select **Continue to configuration** to go to the *Configuration* page.

## Result

BlueXP deploys the BlueXP classification instance in your cloud provider.

Upgrades to the BlueXP Connector and BlueXP classification software is automated as long as the instances have internet connectivity.

## What's Next

From the Configuration page you can select the data sources that you want to scan.



## Install BlueXP classification on a host that has internet access

Complete a few steps to install BlueXP classification on a Linux host in your network, or on a Linux host in the cloud, that has internet access. You'll need to deploy the Linux host manually in your network or in the cloud as part of this installation.

The on-premises installation may be a good option if you prefer to scan on-premises ONTAP systems using a BlueXP classification instance that's also located on premises — but this is not a requirement. The software functions exactly the same way regardless of which installation method you choose.

The BlueXP classification installation script starts by checking if the system and environment meet the required prerequisites. If the prerequisites are all met, then the installation starts. If you would like to verify the prerequisites independently of running the BlueXP classification installation, there is a separate software package you can download that only tests for the prerequisites. [See how to check if your Linux host is ready to install BlueXP classification.](#)

The typical installation on a Linux host *in your premises* has the following components and connections.

The typical installation on a Linux host *in the cloud* has the following components and connections.



For legacy versions 1.30 and earlier, if you need to install BlueXP classification on multiple hosts, refer to [Install BlueXP classification on multiple hosts with no internet access.](#)

You can also [install BlueXP classification in an on-premises site that doesn't have internet access.](#)

### Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

#### Create a Connector

If you don't already have a Connector, [deploy the Connector on-premises](#) on a Linux host in your network, or on a Linux host in the cloud.

You can also create a Connector with your cloud provider. See [creating a Connector in AWS](#), [creating a Connector in Azure](#), or [creating a Connector in GCP](#).

2

#### Review prerequisites

Ensure that your environment can meet the prerequisites. This includes outbound internet access for the instance, connectivity between the Connector and BlueXP classification over port 443, and more. [See the complete list.](#)

You also need a Linux system that meets the [following requirements](#).

3

#### Download and deploy BlueXP classification

Download the Cloud BlueXP classification software from the NetApp Support Site and copy the installer file to the Linux host you plan to use. Then launch the installation wizard and follow the prompts to deploy the BlueXP classification instance.

## Create a Connector

A BlueXP Connector is required before you can install and use BlueXP classification. In most cases you'll probably have a Connector set up before you attempt to activate BlueXP classification because most [BlueXP features require a Connector](#), but there are cases where you'll need to set one up now.

To create one in your cloud provider environment, see [creating a Connector in AWS](#), [creating a Connector in Azure](#), or [creating a Connector in GCP](#).

There are some scenarios where you have to use a Connector that's deployed in a specific cloud provider:

- When scanning data in Cloud Volumes ONTAP in AWS or Amazon FSx for ONTAP, you use a Connector in AWS.
- When scanning data in Cloud Volumes ONTAP in Azure or in Azure NetApp Files, you use a Connector in Azure.

For Azure NetApp Files, it must be deployed in the same region as the volumes you wish to scan.

- When scanning data in Cloud Volumes ONTAP in GCP, you use a Connector in GCP.

On-prem ONTAP systems, NetApp file shares and database accounts can be scanned using any of these cloud Connectors.

Note that you can also [deploy the Connector on-premises](#) on a Linux host in your network or on a Linux host in the cloud. Some users planning to install BlueXP classification on-prem may also choose to install the Connector on-prem.

You'll need the IP address or host name of the Connector system when installing BlueXP classification. You'll have this information if you installed the Connector in your premises. If the Connector is deployed in the cloud, you can find this information from the BlueXP console: click the Help icon, select **Support**, and click **BlueXP Connector**.

## Prepare the Linux host system

BlueXP classification software must run on a host that meets specific operating system requirements, RAM requirements, software requirements, and so on. The Linux host can be in your network, or in the cloud.

Ensure that you can keep BlueXP classification running. The BlueXP classification machine needs to stay on to continuously scan your data.

- BlueXP classification is not supported on a host that is shared with other applications—the host must be a dedicated host.
- When building the host system in your premises, you can choose among these system sizes depending on the size of the dataset that you plan to have BlueXP classification scan.

System size	CPU	RAM (swap memory must be disabled)	Disk
<b>Extra Large</b>	32 CPUs	128 GB RAM	<ul style="list-style-type: none"> <li>• 1 TiB SSD on /, or 100 GiB available on /opt</li> <li>• 895 GiB available on /var/lib/docker</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>
<b>Large</b>	16 CPUs	64 GB RAM	<ul style="list-style-type: none"> <li>• 500 GiB SSD on /, or 100 GiB available on /opt</li> <li>• 395 GiB available on /var/lib/docker or for Podman /var/lib/containers or for Podman /var/lib/containers</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>

- When deploying a compute instance in the cloud for your BlueXP classification installation, we recommend a system that meets the "Large" system requirements above:
  - **Amazon Elastic Compute Cloud (Amazon EC2) instance type:** We recommend "m6i.4xlarge". [See additional AWS instance types.](#)
  - **Azure VM size:** We recommend "Standard\_D16s\_v3". [See additional Azure instance types.](#)
  - **GCP machine type:** We recommend "n2-standard-16". [See additional GCP instance types.](#)
- **UNIX folder permissions:** The following minimum UNIX permissions are required:

Folder	Minimum Permissions
/tmp	rwXrwxrwt
/opt	rwXr-xr-x
/var/lib/docker	rwX-----
/usr/lib/systemd/system	rwXr-xr-x

- **Operating system:**
  - The following operating systems require using the Docker container engine:
    - Red Hat Enterprise Linux version 7.8 and 7.9
    - Ubuntu 22.04 (requires BlueXP classification version 1.23 or greater)
    - Ubuntu 24.04 (requires BlueXP classification version 1.23 or greater)

- The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater:
  - Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4, and 9.5
- Advanced Vector Extensions (AVX2) must be enabled on the host system.
- **Red Hat Subscription Management:** The host must be registered with Red Hat Subscription Management. If it's not registered, the system can't access repositories to update required 3rd-party software during installation.
- **Additional software:** You must install the following software on the host before you install BlueXP classification:
  - Depending on the OS you are using, you'll need to install one of the container engines:
    - Docker Engine version 19.3.1 or greater. [View installation instructions.](#)
    - Podman version 4 or greater. To install Podman, enter `(sudo yum install podman netavark -y)`.
- Python version 3.6 or greater. [View installation instructions.](#)
  - **NTP considerations:** NetApp recommends configuring the BlueXP classification system to use a Network Time Protocol (NTP) service. The time must be synchronized between the BlueXP classification system and the BlueXP Connector system.
- **Firewalld considerations:** If you are planning to use `firewalld`, we recommend that you enable it before installing BlueXP classification. Run the following commands to configure `firewalld` so that it is compatible with BlueXP classification:

```
firewall-cmd --permanent --add-service=http
firewall-cmd --permanent --add-service=https
firewall-cmd --permanent --add-port=80/tcp
firewall-cmd --permanent --add-port=8080/tcp
firewall-cmd --permanent --add-port=443/tcp
firewall-cmd --reload
```

If you're planning to use additional BlueXP classification hosts as scanner nodes, add these rules to your primary system at this time:

```
firewall-cmd --permanent --add-port=2377/tcp
firewall-cmd --permanent --add-port=7946/udp
firewall-cmd --permanent --add-port=7946/tcp
firewall-cmd --permanent --add-port=4789/udp
```

Note that you must restart Docker or Podman whenever you enable or update `firewalld` settings.



The IP address of the BlueXP classification host system can't be changed after installation.

## Enable outbound internet access from BlueXP classification

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

Endpoints	Purpose
https://api.bluexp.netapp.com	Communication with the BlueXP service, which includes NetApp accounts.
https://netapp-cloud-account.auth0.com https://auth0.com	Communication with the BlueXP website for centralized user authentication.
https://support.compliance.api.bluexp.netapp.com/ https://hub.docker.com https://auth.docker.io https://registry-1.docker.io https://index.docker.io/ https://dseasb33srrn.cloudfront.net/ https://production.cloudflare.docker.com/	Provides access to software images, manifests, templates, and to send logs and metrics.
https://support.compliance.api.bluexp.netapp.com/	Enables NetApp to stream data from audit records.
https://github.com/docker https://download.docker.com	Provides prerequisite packages for docker installation.
http://packages.ubuntu.com/ http://archive.ubuntu.com	Provides prerequisite packages for Ubuntu installation.

## Verify that all required ports are enabled


You must ensure that all required ports are open for communication between the Connector, BlueXP classification, Active Directory, and your data sources.

Connection Type	Ports	Description
Connector <> BlueXP classification	8080 (TCP), 443 (TCP), and 80. 9000	<p>The firewall or routing rules for the Connector must allow inbound and outbound traffic over port 443 to and from the BlueXP classification instance.</p> <p>Make sure port 8080 is open so you can see the installation progress in BlueXP.</p> <p>If a firewall is used on the Linux host, port 9000 is required for internal processes within an Ubuntu server.</p>

Connection Type	Ports	Description
Connector <> ONTAP cluster (NAS)	443 (TCP)	<p>BlueXP discovers ONTAP clusters using HTTPS. If you use custom firewall policies, they must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• The Connector host must allow outbound HTTPS access through port 443. If the Connector is in the cloud, all outbound communication is allowed by the predefined firewall or routing rules.</li> <li>• The ONTAP cluster must allow inbound HTTPS access through port 443. The default "mgmt" firewall policy allows inbound HTTPS access from all IP addresses. If you modified this default policy, or if you created your own firewall policy, you must associate the HTTPS protocol with that policy and enable access from the Connector host.</li> </ul>
BlueXP classification <> ONTAP cluster	<ul style="list-style-type: none"> <li>• For NFS - 111 (TCP\UDP) and 2049 (TCP\UDP)</li> <li>• For CIFS - 139 (TCP\UDP) and 445 (TCP\UDP)</li> </ul>	<p>BlueXP classification needs a network connection to each Cloud Volumes ONTAP subnet or on-prem ONTAP system. Firewalls or routing rules for Cloud Volumes ONTAP must allow inbound connections from the BlueXP classification instance.</p> <p>Make sure these ports are open to the BlueXP classification instance:</p> <ul style="list-style-type: none"> <li>• For NFS - 111 and 2049</li> <li>• For CIFS - 139 and 445</li> </ul> <p>NFS volume export policies must allow access from the BlueXP classification instance.</p>
BlueXP classification <> Active Directory	389 (TCP & UDP), 636 (TCP), 3268 (TCP), and 3269 (TCP)	<p>You must have an Active Directory already set up for the users in your company. Additionally, BlueXP classification needs Active Directory credentials to scan CIFS volumes.</p> <p>You must have the information for the Active Directory:</p> <ul style="list-style-type: none"> <li>• DNS Server IP Address, or multiple IP Addresses</li> <li>• User Name and Password for the server</li> <li>• Domain Name (Active Directory Name)</li> <li>• Whether you are using secure LDAP (LDAPS) or not</li> <li>• LDAP Server Port (typically 389 for LDAP, and 636 for secure LDAP)</li> </ul>

## Install BlueXP classification on the Linux host

For typical configurations you'll install the software on a single host system. [See those steps here.](#)



See [Preparing the Linux host system](#) and [Reviewing prerequisites](#) for the full list of requirements before you deploy BlueXP classification.

Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.



BlueXP classification is currently unable to scan S3 buckets, Azure NetApp Files, or FSx for ONTAP when the software is installed on premises. In these cases you'll need to deploy a separate Connector and instance of BlueXP classification in the cloud and [switch between Connectors](#) for your different data sources.

### Single-host installation for typical configurations

Review the requirements and follow these steps when installing BlueXP classification software on a single on-premises host.

[Watch this video](#) to see how to install BlueXP classification.

Note that all installation activities are logged when installing BlueXP classification. If you run into any issues during installation, you can view the contents of the installation audit log. It is written to `/opt/netapp/install_logs/`. [See more details here.](#)

### Before you begin

- Verify that your Linux system meets the [host requirements](#).
- Verify that the system has the two prerequisite software packages installed (Docker Engine or Podman, and Python 3).
- Make sure you have root privileges on the Linux system.
- If you're using a proxy for access to the internet:

- You'll need the proxy server information (IP address or host name, connection port, connection scheme: https or http, user name and password).
  - If the proxy is performing TLS interception, you'll need to know the path on the BlueXP classification Linux system where the TLS CA certificates are stored.
  - The proxy must be non-transparent. BlueXP classification does not currently support transparent proxies.
  - The user must be a local user. Domain users are not supported.
- Verify that your offline environment meets the required [permissions and connectivity](#).

## Steps

1. Download the BlueXP classification software from the [NetApp Support Site](#). The file you should select is named **DATASENSE-INSTALLER-<version>.tar.gz**.
2. Copy the installer file to the Linux host you plan to use (using `scp` or some other method).
3. Unzip the installer file on the host machine, for example:

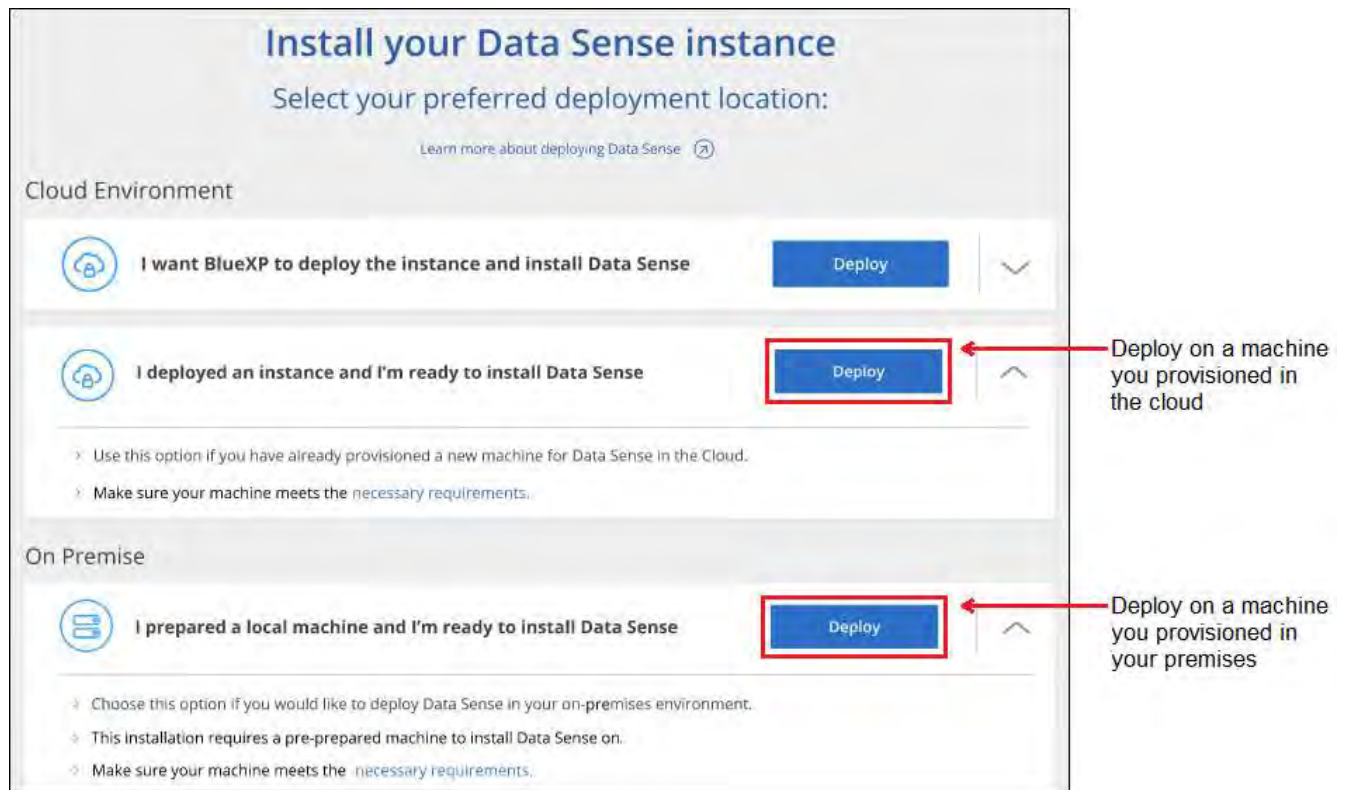
```
tar -xzf DATASENSE-INSTALLER-V1.25.0.tar.gz
```

4. In BlueXP, select **Governance > Classification**.
5. Select **Deploy Classification On-Premises or Cloud**.

The screenshot shows the 'Classification' page in the BlueXP interface. The main heading is 'Classify and take control of your data with BlueXP Classification'. Below this, there is a sub-heading 'Driven by powerful artificial intelligence, NetApp's BlueXP Classification gives you control of your data. Map, classify and understand all your cloud and on-premises data to stay secure and compliant, reduce storage costs, and get assistance with data migration projects.' A 'How does it work?' link is present. A prominent blue button reads 'Deploy Classification On-Premises or Cloud'. Below the button, a note states: 'You will be prompted to first deploy the BlueXP Connector, which is required for using Classification.' The page also features four key benefits with icons: 'Multiple Data Sources' (Cloud and on-premises NetApp storage, databases and more.), 'Take Control' (Map and classify data, take action, set alerts and gain control.), 'Safe' (Data never leaves your network. Agentless solution.), and 'Now Available at No Cost' (As part of BlueXP core capability. Learn more).

6. Depending on whether you are installing BlueXP classification on an instance you prepared in the cloud or on an instance you prepared in your premises, click the appropriate **Deploy** button to start the BlueXP classification installation.





7. The *Deploy Data Sense On Premises* dialog is displayed. Copy the provided command (for example: `sudo ./install.sh -a 12345 -c 27AG75 -t 2198qq`) and paste it in a text file so you can use it later. Then click **Close** to dismiss the dialog.
8. On the host machine, enter the command you copied and then follow a series of prompts, or you can provide the full command including all required parameters as command line arguments.

Note that the installer performs a pre-check to make sure your system and networking requirements are in place for a successful installation. [Watch this video](#) to understand the pre-check messages and implications.

Enter parameters as prompted:	Enter the full command:
<p>1. Paste the command you copied from step 7:  <code>sudo ./install.sh -a &lt;account_id&gt;  -c &lt;client_id&gt; -t &lt;user_token&gt;</code></p> <p>If you are installing on a cloud instance (not on your premises), add <code>--manual-cloud -install &lt;cloud_provider&gt;</code>.</p> <p>2. Enter the IP address or host name of the BlueXP classification host machine so it can be accessed by the Connector system.</p> <p>3. Enter the IP address or host name of the BlueXP Connector host machine so it can be accessed by the BlueXP classification system.</p> <p>4. Enter proxy details as prompted. If your BlueXP Connector already uses a proxy, there is no need to enter this information again here since BlueXP classification will automatically use the proxy used by the Connector.</p>	<p>Alternatively, you can create the whole command in advance, providing the necessary host and proxy parameters:</p> <pre>sudo ./install.sh -a &lt;account_id&gt; -c &lt;client_id&gt; -t &lt;user_token&gt; --host &lt;ds_host&gt; --manager-host &lt;cm_host&gt; --manual-cloud-install &lt;cloud_provider&gt; --proxy-host &lt;proxy_host&gt; --proxy-port &lt;proxy_port&gt; --proxy-scheme &lt;proxy_scheme&gt; --proxy -user &lt;proxy_user&gt; --proxy-password &lt;proxy_password&gt; --cacert-folder-path &lt;ca_cert_dir&gt;</pre>

Variable values:

- *account\_id* = NetApp Account ID
- *client\_id* = Connector Client ID (add the suffix "clients" to the client ID if it not already there)
- *user\_token* = JWT user access token
- *ds\_host* = IP address or host name of the BlueXP classification Linux system.
- *cm\_host* = IP address or host name of the BlueXP Connector system.
- *cloud\_provider* = When installing on a cloud instance, enter "AWS", "Azure", or "Gcp" depending on cloud provider.
- *proxy\_host* = IP or host name of the proxy server if the host is behind a proxy server.
- *proxy\_port* = Port to connect to the proxy server (default 80).
- *proxy\_scheme* = Connection scheme: https or http (default http).
- *proxy\_user* = Authenticated user to connect to the proxy server, if basic authentication is required. The user must be a local user - domain users are not supported.
- *proxy\_password* = Password for the user name that you specified.
- *ca\_cert\_dir* = Path on the BlueXP classification Linux system containing additional TLS CA certificate bundles. Only required if the proxy is performing TLS interception.

## Result

The BlueXP classification installer installs packages, registers the installation, and installs BlueXP classification. Installation can take 10 to 20 minutes.

If there is connectivity over port 8080 between the host machine and the Connector instance, you'll see the installation progress in the BlueXP classification tab in BlueXP.

## What's Next

From the Configuration page you can select the data sources that you want to scan.

## Install BlueXP classification on a Linux host with no internet access

Complete a few steps to install BlueXP classification on a Linux host in an on-premises site that doesn't have internet access - also known as *private mode*. This type of installation, which uses an installation script, has no connectivity to the BlueXP SaaS layer.

[Learn about the different deployment modes for the BlueXP Connector and BlueXP classification.](#)

You can also [deploy BlueXP classification in an on-premises site that has internet access.](#)

The BlueXP classification installation script starts by checking if the system and environment meet the required prerequisites. If the prerequisites are all met, then the installation starts. If you would like to verify the prerequisites independently of running the BlueXP classification installation, there is a separate software package you can download that only tests for the prerequisites. [See how to check if your Linux host is ready to install BlueXP classification.](#)



For legacy versions 1.30 and earlier, if you need to install BlueXP classification on multiple hosts, refer to [Install BlueXP classification on multiple hosts with no internet access.](#)

## Supported data sources

When installed private mode (sometimes called an "offline" or "dark" site), BlueXP classification can only scan data from data sources that are also local to the on-premises site. At this time, BlueXP classification can scan the following **local** data sources:

- On-premises ONTAP systems
- Database schemas

There is no support currently for scanning Cloud Volumes ONTAP, Azure NetApp Files, or FSx for ONTAP accounts when BlueXP classification is deployed in private mode.

## Limitations

Most BlueXP classification features work when it is deployed in a site with no internet access. However, certain features that require internet access are not supported, for example:

- Setting BlueXP roles for different users (for example, Account Admin or Compliance Viewer)
- Copying and synchronizing source files using BlueXP copy and sync
- Automated software upgrades from BlueXP

Both the BlueXP Connector and BlueXP classification will require periodic manual upgrades to enable new features. You can see the BlueXP classification version at the bottom of the BlueXP classification UI pages. Check the [BlueXP classification Release Notes](#) to see the new features in each release and whether you want those features. Then you can follow the steps to [upgrade the BlueXP Connector](#) and [upgrade your BlueXP classification software.](#)

## Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

### Install the BlueXP Connector

If you don't already have a Connector installed in private mode, [deploy the Connector](#) on a Linux host now.

2

### Review BlueXP classification prerequisites

Ensure that your Linux system meets the [host requirements](#), that it has all required software installed, and that your offline environment meets the required [permissions and connectivity](#).

3

### Download and deploy BlueXP classification

Download the BlueXP classification software from the NetApp Support Site and copy the installer file to the Linux host you plan to use. Then launch the installation wizard and follow the prompts to deploy the BlueXP classification instance.

## Install the BlueXP Connector

If you don't already have a BlueXP Connector installed in private mode, [deploy the Connector](#) on a Linux host in your offline site.

## Prepare the Linux host system

BlueXP classification software must run on a host that meets specific operating system requirements, RAM requirements, software requirements, and so on.

- BlueXP classification is not supported on a host that is shared with other applications—the host must be a dedicated host.
- When building the host system in your premises, you can choose among these system sizes depending on the size of the dataset that you plan to have BlueXP classification scan.

System size	CPU	RAM (swap memory must be disabled)	Disk
Extra Large	32 CPUs	128 GB RAM	<ul style="list-style-type: none"><li>• 1 TiB SSD on /, or 100 GiB available on /opt</li><li>• 895 GiB available on /var/lib/docker</li><li>• 5 GiB on /tmp</li><li>• <b>For Podman, 5 GB on /tmp</b></li><li>• <b>For Podman, 30 GB on /var/tmp</b></li></ul>

System size	CPU	RAM (swap memory must be disabled)	Disk
Large	16 CPUs	64 GB RAM	<ul style="list-style-type: none"> <li>• 500 GiB SSD on /, or 100 GiB available on /opt</li> <li>• 395 GiB available on /var/lib/docker or for Podman /var/lib/containers or for Podman /var/lib/containers</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>

- When deploying a compute instance in the cloud for your BlueXP classification installation, we recommend a system that meets the "Large" system requirements above:
  - **Amazon Elastic Compute Cloud (Amazon EC2) instance type:** We recommend "m6i.4xlarge". [See additional AWS instance types.](#)
  - **Azure VM size:** We recommend "Standard\_D16s\_v3". [See additional Azure instance types.](#)
  - **GCP machine type:** We recommend "n2-standard-16". [See additional GCP instance types.](#)
- **UNIX folder permissions:** The following minimum UNIX permissions are required:

Folder	Minimum Permissions
/tmp	rwXrwxrwt
/opt	rwXr-Xr-X
/var/lib/docker	rwX-----
/usr/lib/systemd/system	rwXr-Xr-X

- **Operating system:**
  - The following operating systems require using the Docker container engine:
    - Red Hat Enterprise Linux version 7.8 and 7.9
    - Ubuntu 22.04 (requires BlueXP classification version 1.23 or greater)
    - Ubuntu 24.04 (requires BlueXP classification version 1.23 or greater)
  - The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater:
    - Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4, and 9.5
  - Advanced Vector Extensions (AVX2) must be enabled on the host system.
- **Red Hat Subscription Management:** The host must be registered with Red Hat Subscription Management. If it's not registered, the system can't access repositories to update required 3rd-party software during installation.
- **Additional software:** You must install the following software on the host before you install BlueXP classification:

- Depending on the OS you are using, you'll need to install one of the container engines:
  - Docker Engine version 19.3.1 or greater. [View installation instructions.](#)
  - Podman version 4 or greater. To install Podman, enter `(sudo yum install podman netavark -y)`.
- Python version 3.6 or greater. [View installation instructions.](#)
  - **NTP considerations:** NetApp recommends configuring the BlueXP classification system to use a Network Time Protocol (NTP) service. The time must be synchronized between the BlueXP classification system and the BlueXP Connector system.
- **Firewalld considerations:** If you are planning to use `firewalld`, we recommend that you enable it before installing BlueXP classification. Run the following commands to configure `firewalld` so that it is compatible with BlueXP classification:

```
firewall-cmd --permanent --add-service=http
firewall-cmd --permanent --add-service=https
firewall-cmd --permanent --add-port=80/tcp
firewall-cmd --permanent --add-port=8080/tcp
firewall-cmd --permanent --add-port=443/tcp
firewall-cmd --reload
```

Note that you must restart Docker or Podman whenever you enable or update `firewalld` settings.



The IP address of the BlueXP classification host system can't be changed after installation.

## Verify BlueXP and BlueXP classification prerequisites

Review the following prerequisites to make sure that you have a supported configuration before you deploy BlueXP classification.

- Ensure that the Connector has permissions to deploy resources and create security groups for the BlueXP classification instance. You can find the latest BlueXP permissions in [the policies provided by NetApp.](#)
- Ensure that you can keep BlueXP classification running. The BlueXP classification instance needs to stay on to continuously scan your data.
- Ensure web browser connectivity to BlueXP classification. After BlueXP classification is enabled, ensure that users access the BlueXP interface from a host that has a connection to the BlueXP classification instance.

The BlueXP classification instance uses a private IP address to ensure that the indexed data isn't accessible to others. As a result, the web browser that you use to access BlueXP must have a connection to that private IP address. That connection can come from a host that's inside the same network as the BlueXP classification instance.

## Verify that all required ports are enabled


You must ensure that all required ports are open for communication between the Connector, BlueXP classification, Active Directory, and your data sources.

Connection Type	Ports	Description
Connector <> BlueXP classification	8080 (TCP), 6000 (TCP), 443 (TCP), and 80. 9000	<p>The security group for the Connector must allow inbound and outbound traffic over ports 6000 and 443 to and from the BlueXP classification instance.</p> <ul style="list-style-type: none"> <li>• Port 6000 is required so that the BlueXP classification BYOL license works in a dark site.</li> <li>• Port 8080 should be open so you can see the installation progress in BlueXP.</li> <li>• If a firewall is used on the Linux host, port 9000 is required for internal processes within an Ubuntu server.</li> </ul>
Connector <> ONTAP cluster (NAS)	443 (TCP)	<p>BlueXP discovers ONTAP clusters using HTTPS. If you use custom firewall policies, they must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• The Connector host must allow outbound HTTPS access through port 443. If the Connector is in the cloud, all outbound communication is allowed by the predefined security group.</li> <li>• The ONTAP cluster must allow inbound HTTPS access through port 443. The default "mgmt" firewall policy allows inbound HTTPS access from all IP addresses. If you modified this default policy, or if you created your own firewall policy, you must associate the HTTPS protocol with that policy and enable access from the Connector host.</li> </ul>
BlueXP classification <> ONTAP cluster	<ul style="list-style-type: none"> <li>• For NFS - 111 (TCP\UDP) and 2049 (TCP\UDP)</li> <li>• For CIFS - 139 (TCP\UDP) and 445 (TCP\UDP)</li> </ul>	<p>BlueXP classification needs a network connection to each Cloud Volumes ONTAP subnet or on-prem ONTAP system. Security groups for Cloud Volumes ONTAP must allow inbound connections from the BlueXP classification instance.</p> <p>Make sure these ports are open to the BlueXP classification instance:</p> <ul style="list-style-type: none"> <li>• For NFS - 111 and 2049</li> <li>• For CIFS - 139 and 445</li> </ul> <p>NFS volume export policies must allow access from the BlueXP classification instance.</p>

Connection Type	Ports	Description
BlueXP classification <> Active Directory	389 (TCP & UDP), 636 (TCP), 3268 (TCP), and 3269 (TCP)	<p>You must have an Active Directory already set up for the users in your company. Additionally, BlueXP classification needs Active Directory credentials to scan CIFS volumes.</p> <p>You must have the information for the Active Directory:</p> <ul style="list-style-type: none"> <li>• DNS Server IP Address, or multiple IP Addresses</li> <li>• User Name and Password for the server</li> <li>• Domain Name (Active Directory Name)</li> <li>• Whether you are using secure LDAP (LDAPS) or not</li> <li>• LDAP Server Port (typically 389 for LDAP, and 636 for secure LDAP)</li> </ul>
If a firewall used on Linux host	9000	Needed for internal processes within an Ubuntu server.

### Install BlueXP classification on the on-premises Linux host

For typical configurations you'll install the software on a single host system.



### Single-host installation for typical configurations

Follow these steps when installing BlueXP classification software on a single on-premises host in an offline environment.

Note that all installation activities are logged when installing BlueXP classification. If you run into any issues during installation, you can view the contents of the installation audit log. It is written to



/opt/netapp/install\_logs/. [See more details here.](#)

## Before you begin

- Verify that your Linux system meets the [host requirements](#).
- Verify that you have installed the two prerequisite software packages (Docker Engine or Podman, and Python 3).
- Make sure you have root privileges on the Linux system.
- Verify that your offline environment meets the required [permissions and connectivity](#).

## Steps

1. On an internet-configured system, download the BlueXP classification software from the [NetApp Support Site](#). The file you should select is named **DataSense-offline-bundle-<version>.tar.gz**.
2. Copy the installer bundle to the Linux host you plan to use in private mode.
3. Unzip the installer bundle on the host machine, for example:

```
tar -xzf DataSense-offline-bundle-v1.25.0.tar.gz
```

This extracts required software and the actual installation file **cc\_onprem\_installer.tar.gz**.

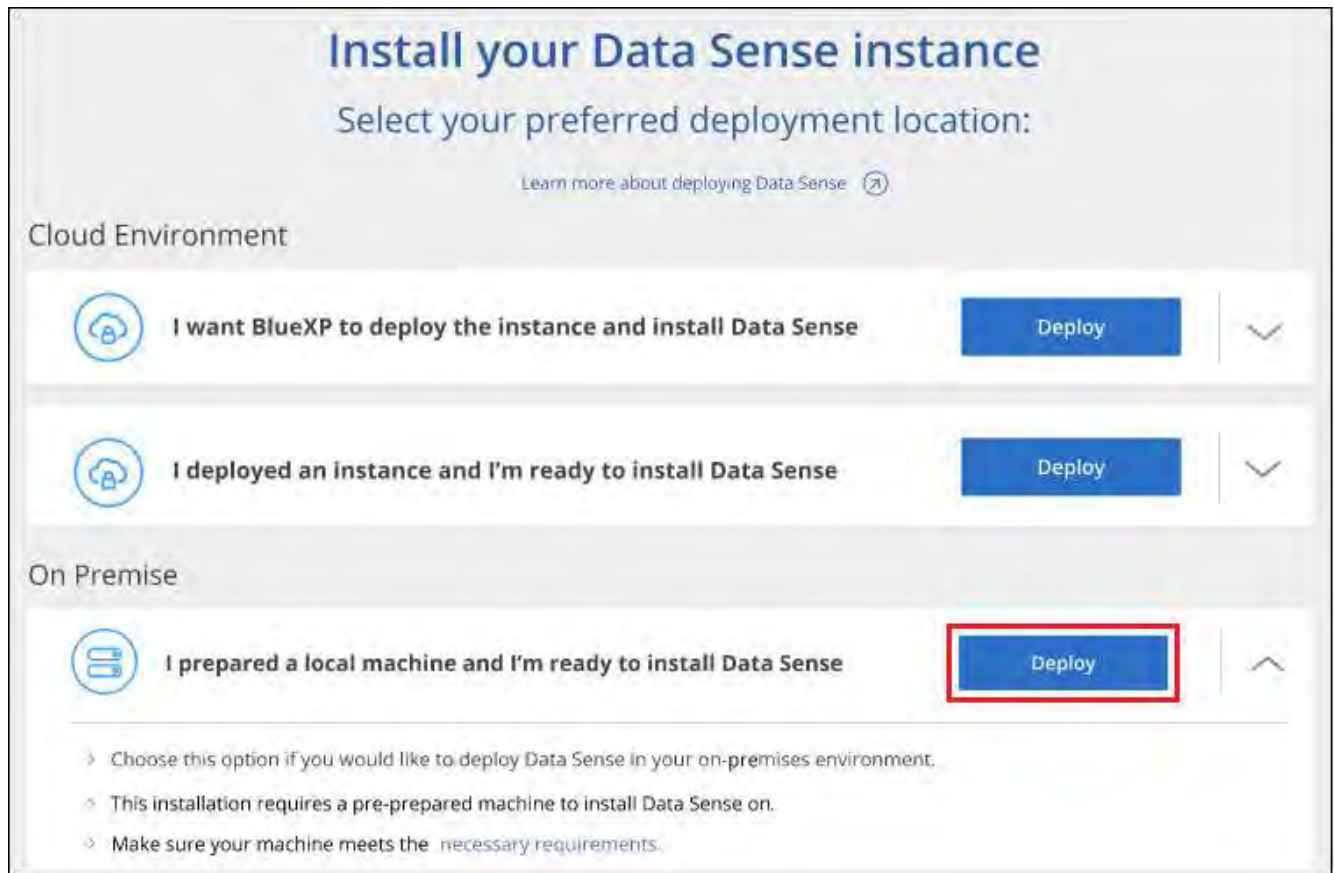
4. Unzip the installation file on the host machine, for example:

```
tar -xzf cc_onprem_installer.tar.gz
```

5. Launch BlueXP and select **Governance > Classification**.
6. Select **Deploy Classification On-Premises or Cloud**.

The screenshot shows the 'Classification' page in the BlueXP interface. The main heading is 'Classify and take control of your data with BlueXP Classification'. Below this, there is a brief description of the service and a 'Deploy Classification On-Premises or Cloud' button. A note indicates that the user will be prompted to first deploy the BlueXP Connector. The page also features four key benefits: 'Multiple Data Sources' (Cloud and on-premises NetApp storage, databases and more), 'Take Control' (Map and classify data, take action, set alerts and gain control), 'Safe' (Data never leaves your network. Agentless solution), and 'Now Available at No Cost' (As part of BlueXP core capability. Learn more). The background of the page shows a dashboard with various charts and data points.

7. Click **Deploy** to start the on-prem installation.



8. The *Deploy Data Sense On Premises* dialog is displayed. Copy the provided command (for example: `sudo ./install.sh -a 12345 -c 27AG75 -t 2198qq --darksite`) and paste it in a text file so you can use it later. Then click **Close** to dismiss the dialog.
9. On the host machine, enter the command you copied and then follow a series of prompts, or you can provide the full command including all required parameters as command line arguments.

Note that the installer performs a pre-check to make sure your system and networking requirements are in place for a successful installation.

Enter parameters as prompted:	Enter the full command:
<ol style="list-style-type: none"> <li>1. Paste the information you copied from step 8:  <code>sudo ./install.sh -a &lt;account_id&gt;  -c &lt;client_id&gt; -t &lt;user_token&gt;  --darksite</code> </li> <li>2. Enter the IP address or host name of the BlueXP classification host machine so it can be accessed by the Connector system.</li> <li>3. Enter the IP address or host name of the BlueXP Connector host machine so it can be accessed by the BlueXP classification system.</li> </ol>	<p>Alternatively, you can create the whole command in advance, providing the necessary host parameters:</p> <pre>sudo ./install.sh -a &lt;account_id&gt; -c &lt;client_id&gt; -t &lt;user_token&gt; --host &lt;ds_host&gt; --manager-host &lt;cm_host&gt; --no-proxy --darksite</pre>

Variable values:

- *account\_id* = NetApp Account ID

- *client\_id* = Connector Client ID (add the suffix "clients" to the client ID if it not already there)
- *user\_token* = JWT user access token
- *ds\_host* = IP address or host name of the BlueXP classification system.
- *cm\_host* = IP address or host name of the BlueXP Connector system.

## Result

The BlueXP classification installer installs packages, registers the installation, and installs BlueXP classification. Installation can take 10 to 20 minutes.

If there is connectivity over port 8080 between the host machine and the Connector instance, you'll see the installation progress in the BlueXP classification tab in BlueXP.

## What's Next

From the Configuration page you can select the local [on-prem ONTAP clusters](#) and [databases](#) that you want to scan.

## Upgrade BlueXP classification software

Since BlueXP classification software is updated with new features on a regular basis, you should get into a routine to check for new versions periodically to make sure you're using the newest software and features. You'll need to upgrade BlueXP classification software manually because there's no internet connectivity to perform the upgrade automatically.

### Before you begin

- We recommend that your BlueXP Connector software is upgraded to the newest available version. [See the Connector upgrade steps](#).
- Starting with BlueXP classification version 1.24 you can perform upgrades to any future version of software.

If your BlueXP classification software is running a version prior to 1.24, you can upgrade only one major version at a time. For example, if you have version 1.21.x installed, you can upgrade only to 1.22.x. If you are a few major versions behind, you'll need to upgrade the software multiple times.

### Steps

1. On an internet-configured system, download the BlueXP classification software from the [NetApp Support Site](#). The file you should select is named **DataSense-offline-bundle-<version>.tar.gz**.
2. Copy the software bundle to the Linux host where BlueXP classification is installed in the dark site.
3. Unzip the software bundle on the host machine, for example:

```
tar -xvf DataSense-offline-bundle-v1.25.0.tar.gz
```

This extracts the installation file **cc\_onprem\_installer.tar.gz**.

4. Unzip the installation file on the host machine, for example:

```
tar -xzf cc_onprem_installer.tar.gz
```

This extracts the upgrade script **start\_darksite\_upgrade.sh** and any required third-party software.

5. Run the upgrade script on the host machine, for example:

```
start_darksite_upgrade.sh
```

## Result

The BlueXP classification software is upgraded on your host. The update can take 5 to 10 minutes.

You can verify that the software has been updated by checking the version at the bottom of the BlueXP classification UI pages.

## Check that your Linux host is ready to install BlueXP classification

Before installing BlueXP classification manually on a Linux host, optionally run a script on the host to verify that all the prerequisites are in place for installing BlueXP classification. You can run this script on a Linux host in your network, or on a Linux host in the cloud. The host can be connected to the internet, or the host can reside in a site that doesn't have internet access (a *dark site*).

There is also a prerequisite test script that is part of the BlueXP classification installation script. The script described here is specifically designed for users who want to verify the Linux host independently of running the BlueXP classification installation script.

## Getting Started

You'll perform the following tasks.

1. Optionally, install a BlueXP Connector if you don't already have one installed. You can run the test script without having a Connector installed, but the script checks for connectivity between the Connector and the BlueXP classification host machine - so it is recommended that you have a Connector.
2. Prepare the host machine and verify that it meets all the requirements.
3. Enable outbound internet access from the BlueXP classification host machine.
4. Verify that all required ports are enabled on all systems.
5. Download and run the Prerequisite test script.

## Create a Connector

A BlueXP Connector is required before you can install and use BlueXP classification. You can, however, run the Prerequisites script without a Connector.

You can [install the Connector on-premises](#) on a Linux host in your network or on a Linux host in the cloud. Some users planning to install BlueXP classification on-prem may also choose to install the Connector on-prem.

To create a Connector in your cloud provider environment, see [creating a Connector in AWS](#), [creating a Connector in Azure](#), or [creating a Connector in GCP](#).

You'll need the IP address or host name of the Connector system when running the Prerequisites script. You'll

have this information if you installed the Connector in your premises. If the Connector is deployed in the cloud, you can find this information from the BlueXP console: click the Help icon, select **Support**, and click **BlueXP Connector**.

### Verify host requirements

BlueXP classification software must run on a host that meets specific operating system requirements, RAM requirements, software requirements, and so on.

- BlueXP classification is not supported on a host that is shared with other applications—the host must be a dedicated host.
- When building the host system in your premises, you can choose among these system sizes depending on the size of the dataset that you plan to have BlueXP classification scan.

System size	CPU	RAM (swap memory must be disabled)	Disk
<b>Extra Large</b>	32 CPUs	128 GB RAM	<ul style="list-style-type: none"> <li>• 1 TiB SSD on /, or 100 GiB available on /opt</li> <li>• 895 GiB available on /var/lib/docker</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>
<b>Large</b>	16 CPUs	64 GB RAM	<ul style="list-style-type: none"> <li>• 500 GiB SSD on /, or 100 GiB available on /opt</li> <li>• 395 GiB available on /var/lib/docker or for Podman /var/lib/containers or for Podman /var/lib/containers</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>

- When deploying a compute instance in the cloud for your BlueXP classification installation, we recommend a system that meets the "Large" system requirements above:
  - **Amazon Elastic Compute Cloud (Amazon EC2) instance type:** We recommend "m6i.4xlarge". [See additional AWS instance types.](#)
  - **Azure VM size:** We recommend "Standard\_D16s\_v3". [See additional Azure instance types.](#)
  - **GCP machine type:** We recommend "n2-standard-16". [See additional GCP instance types.](#)
- **UNIX folder permissions:** The following minimum UNIX permissions are required:

Folder	Minimum Permissions
/tmp	rwXrwxrwt
/opt	rwXr-Xr-X
/var/lib/docker	rwX-----
/usr/lib/systemd/system	rwXr-Xr-X

- **Operating system:**

- The following operating systems require using the Docker container engine:
  - Red Hat Enterprise Linux version 7.8 and 7.9
  - Ubuntu 22.04 (requires BlueXP classification version 1.23 or greater)
  - Ubuntu 24.04 (requires BlueXP classification version 1.23 or greater)
- The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater:
  - Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4, and 9.5
- Advanced Vector Extensions (AVX2) must be enabled on the host system.

- **Red Hat Subscription Management:** The host must be registered with Red Hat Subscription Management. If it's not registered, the system can't access repositories to update required 3rd-party software during installation.

- **Additional software:** You must install the following software on the host before you install BlueXP classification:

- Depending on the OS you are using, you'll need to install one of the container engines:
  - Docker Engine version 19.3.1 or greater. [View installation instructions.](#)
  - Podman version 4 or greater. To install Podman, enter `(sudo yum install podman netavark -y)`.

- Python version 3.6 or greater. [View installation instructions.](#)

- **NTP considerations:** NetApp recommends configuring the BlueXP classification system to use a Network Time Protocol (NTP) service. The time must be synchronized between the BlueXP classification system and the BlueXP Connector system.

- **Firewalld considerations:** If you are planning to use `firewalld`, we recommend that you enable it before installing BlueXP classification. Run the following commands to configure `firewalld` so that it is compatible with BlueXP classification:

```
firewall-cmd --permanent --add-service=http
firewall-cmd --permanent --add-service=https
firewall-cmd --permanent --add-port=80/tcp
firewall-cmd --permanent --add-port=8080/tcp
firewall-cmd --permanent --add-port=443/tcp
firewall-cmd --reload
```

If you're planning to use additional BlueXP classification hosts as scanner nodes (in a distributed model),

add these rules to your primary system at this time:

```
firewall-cmd --permanent --add-port=2377/tcp
firewall-cmd --permanent --add-port=7946/udp
firewall-cmd --permanent --add-port=7946/tcp
firewall-cmd --permanent --add-port=4789/udp
```

Note that you must restart Docker or Podman whenever you enable or update `firewalld` settings.

## Enable outbound internet access from BlueXP classification

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



This section is not required for host systems installed in sites without internet connectivity.

Endpoints	Purpose
<a href="https://api.bluelxp.netapp.com">https://api.bluelxp.netapp.com</a>	Communication with the BlueXP service, which includes NetApp accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://auth0.com">https://auth0.com</a>	Communication with the BlueXP website for centralized user authentication.
<a href="https://support.compliance.api.bluelxp.netapp.com/">https://support.compliance.api.bluelxp.netapp.com/</a> <a href="https://hub.docker.com">https://hub.docker.com</a> <a href="https://auth.docker.io">https://auth.docker.io</a> <a href="https://registry-1.docker.io">https://registry-1.docker.io</a> <a href="https://index.docker.io/">https://index.docker.io/</a> <a href="https://dseasb33srrn.cloudfront.net/">https://dseasb33srrn.cloudfront.net/</a> <a href="https://production.cloudflare.docker.com/">https://production.cloudflare.docker.com/</a>	Provides access to software images, manifests, templates, and to send logs and metrics.
<a href="https://support.compliance.api.bluelxp.netapp.com/">https://support.compliance.api.bluelxp.netapp.com/</a>	Enables NetApp to stream data from audit records.
<a href="https://github.com/docker">https://github.com/docker</a> <a href="https://download.docker.com">https://download.docker.com</a>	Provides prerequisite packages for docker installation.
<a href="http://packages.ubuntu.com/">http://packages.ubuntu.com/</a> <a href="http://archive.ubuntu.com">http://archive.ubuntu.com</a>	Provides prerequisite packages for Ubuntu installation.

## Verify that all required ports are enabled

You must ensure that all required ports are open for communication between the Connector, BlueXP classification, Active Directory, and your data sources.



Connection Type	Ports	Description
Connector <> BlueXP classification	8080 (TCP), 443 (TCP), and 80. 9000	<p>The firewall or routing rules for the Connector must allow inbound and outbound traffic over port 443 to and from the BlueXP classification instance.</p> <p>Make sure port 8080 is open so you can see the installation progress in BlueXP.</p> <p>If a firewall is used on the Linux host, port 9000 is required for internal processes within an Ubuntu server.</p>
Connector <> ONTAP cluster (NAS)	443 (TCP)	BlueXP discovers ONTAP clusters using HTTPS. If you use custom firewall policies, the Connector host must allow outbound HTTPS access through port 443. If the Connector is in the cloud, all outbound communication is allowed by the predefined firewall or routing rules.

### Run the BlueXP classification Prerequisites script

Follow these steps to run the BlueXP classification Prerequisites script.

[Watch this video](#) to see how to run the Prerequisites script and interpret the results.

#### Before you begin

- Verify that your Linux system meets the [host requirements](#).
- Verify that the system has the two prerequisite software packages installed (Docker Engine or Podman, and Python 3).
- Make sure you have root privileges on the Linux system.

#### Steps

1. Download the BlueXP classification Prerequisites script from the [NetApp Support Site](#). The file you should select is named **standalone-pre-requisite-tester-<version>**.
2. Copy the file to the Linux host you plan to use (using `scp` or some other method).
3. Assign permissions to run the script.

```
chmod +x standalone-pre-requisite-tester-v1.25.0
```

4. Run the script using the following command.

```
./standalone-pre-requisite-tester-v1.25.0 <--darksite>
```

Add the option "--darksite" only if you are running the script on a host that doesn't have internet access. Certain prerequisite tests are skipped when the host is not connected to the internet.

5. The script prompts you for the IP address of the BlueXP classification host machine.



- Enter the IP address or host name.
6. The script prompts whether you have an installed BlueXP Connector.
    - Enter **N** if you do not have an installed Connector.
    - Enter **Y** if you do have an installed Connector. And then enter the IP address or host name of the BlueXP Connector so the test script can test this connectivity.
  7. The script runs a variety of tests on the system and it displays results as it progresses. When it finishes it writes a log of the session to a file named `prerequisites-test-<timestamp>.log` in the directory `/opt/netapp/install_logs`.

## Result

If all the prerequisites tests ran successfully, you can install BlueXP classification on the host when you are ready.

If any issues were discovered, they are categorized as "Recommended" or "Required" to be fixed. Recommended issues are typically items that would make the BlueXP classification scanning and categorizing tasks run slower. These items do not need to be corrected - but you may want to address them.

If you have any "Required" issues, you should fix the issues and run the Prerequisites test script again.

# Activate scanning on your data sources

## Scan data sources overview with BlueXP classification

BlueXP classification scans the data in the repositories (the volumes, database schemas, or other user data) that you select to identify personal and sensitive data. BlueXP classification then maps your organizational data, categorizes each file, and identifies predefined patterns in the data. The result of the scan is an index of personal information, sensitive personal information, data categories, and file types.

After the initial scan, BlueXP classification continuously scans your data in a round-robin fashion to detect incremental changes. This is why it's important to keep the instance running.

You can enable and disable scans at the volume level or at the database schema level.

## What's the difference between Mapping and Classification scans

You can conduct two types of scans in BlueXP classification:

- **Mapping-only scans** provide only a high-level overview of your data and are performed on selected data sources. Mapping-only scans take less time than map and classify scans because they do not access files to see the data inside. You might want to do this initially to identify areas of research and then perform a Map & Classify scan on those areas.
- **Map & Classify scans** provide deep-level scanning of your data.

The table below shows some of the differences:

Feature	Map & classify scans	Mapping-only scans
Scan speed	Slow	Fast

Feature	Map & classify scans	Mapping-only scans
Pricing	Free	Free
Capacity	Limited to 500 TiB*	Limited to 500 TiB*
List of file types and used capacity	Yes	Yes
Number of files and used capacity	Yes	Yes
Age and size of files	Yes	Yes
Ability to run a <a href="#">Data Mapping Report</a>	Yes	Yes
Data Investigation page to view file details	Yes	No
Search for names within files	Yes	No
Create <a href="#">saved searches</a> that provide custom search results	Yes	No
Ability to run other reports	Yes	No
Ability to see metadata from files*	No	Yes

\* include::\_include/connector-limit.adoc[]

\*The following metadata is extracted from files during mapping scans:

- Working environment
- Working environment type
- Storage repository
- File type
- Used capacity
- Number of files
- File size
- File creation
- File last access
- File last modified
- File discovered time
- Permissions extraction

**Governance dashboard differences:**

<b>Feature</b>	<b>Map &amp; Classify</b>	<b>Map</b>
Stale data	Yes	Yes
Non-business data	Yes	Yes
Duplicated files	Yes	Yes
Predefined saved searches	Yes	No
Default saved searches	Yes	Yes
DDA report	Yes	Yes
Mapping report	Yes	Yes
Sensitivity level detection	Yes	No
Sensitive data with wide permissions	Yes	No
Open permissions	Yes	Yes
Age of data	Yes	Yes
Size of data	Yes	Yes
Categories	Yes	No
File types	Yes	Yes

**Compliance dashboard differences:**

<b>Feature</b>	<b>Map &amp; Classify</b>	<b>Map</b>
Personal information	Yes	No
Sensitive personal information	Yes	No
Privacy risk assessment report	Yes	No
HIPAA report	Yes	No
PCI DSS report	Yes	No

### Investigation filters differences:

Feature	Map & Classify	Map
Saved searches	Yes	Yes
Working environment type	Yes	Yes
Working environment	Yes	Yes
Storage repository	Yes	Yes
File type	Yes	Yes
File size	Yes	Yes
Created time	Yes	Yes
Discovered time	Yes	Yes
Last modified	Yes	Yes
Last access	Yes	Yes
Open permissions	Yes	Yes
File directory path	Yes	Yes
Category	Yes	No
Sensitivity level	Yes	No
Number of identifiers	Yes	No
Personal data	Yes	No
Sensitive personal data	Yes	No
Data subject	Yes	No
Duplicates	Yes	Yes
Classification status	Yes	Status is always "Limited insights"
Scan analysis event	Yes	Yes
File hash	Yes	Yes
Number of users with access	Yes	Yes
User/group permissions	Yes	Yes
File owner	Yes	Yes
Directory type	Yes	Yes

### How quickly does BlueXP classification scan data

The scan speed is affected by network latency, disk latency, network bandwidth, environment size, and file distribution sizes.

- When performing Mapping-only scans, BlueXP classification can scan between 100-150 TiBs of data per

day.

- When performing Map & classify scans, BlueXP classification can scan between 15-40 TiBs of data per day.

## Scan Azure NetApp Files volumes with BlueXP classification

Complete a few steps to get started with BlueXP classification for Azure NetApp Files.

### Discover the Azure NetApp Files system that you want to scan

If the Azure NetApp Files system you want to scan is not already in BlueXP as a working environment, you can add it to the canvas at this time.

[See how to discover the Azure NetApp Files system in BlueXP.](#)

### Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

BlueXP classification must be deployed in the cloud when scanning Azure NetApp Files volumes, and it must be deployed in the same region as the volumes you wish to scan.

**Note:** Deploying BlueXP classification in an on-premises location is not currently supported when scanning Azure NetApp Files volumes.

### Enable BlueXP classification in your working environments

You can enable BlueXP classification on your Azure NetApp Files volumes.

1. From the BlueXP left navigation menu, click **Governance > Classification**.
2. From the BlueXP classification menu, select **Configuration**.



3. Select how you want to scan the volumes in each working environment. [Learn about mapping and classification scans:](#)
  - To map all volumes, select **Map all Volumes**.
  - To map and classify all volumes, select **Map & Classify all Volumes**.
  - To customize scanning for each volume, select **Or select scanning type for each volume**, and then choose the volumes you want to map and/or classify.

See [Enable and disable compliance scans on volumes](#) for details.

4. In the confirmation dialog box, select **Approve** to have BlueXP classification start scanning your volumes.

## Result

BlueXP classification starts scanning the volumes you selected in the working environment. Results will be available in the Compliance dashboard as soon as BlueXP classification finishes the initial scans. The time that it takes depends on the amount of data—it could be a few minutes or hours. You can track the progress of the initial scan by navigating to the **Configuration** menu then selecting the **Working Environment configuration**. The progress of each scan is shown as a progress bar. You can also hover over the progress bar to see the number of files scanned relative to the total files in the volume.



- By default, if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, the system won't scan the files in your volumes because BlueXP classification can't revert the "last access time" to the original timestamp. If you don't care if the last access time is reset, click **Or select scanning type for each volume**. The resulting page has a setting you can enable so that BlueXP classification will scan the volumes regardless of permissions.
- BlueXP classification scans only one file share under a volume. If you have multiple shares in your volumes, you'll need to scan those other shares separately as a shares group. [See more details about this BlueXP classification limitation.](#)

## Verify that BlueXP classification has access to volumes


Make sure that BlueXP classification can access volumes by checking your networking, security groups, and export policies. You'll need to provide BlueXP classification with CIFS credentials so it can access CIFS volumes.



For Azure NetApp Files, BlueXP classification can only scan volumes that are in the same region as BlueXP.

## Steps

1. Make sure that there's a network connection between the BlueXP classification instance and each network that includes volumes for Azure NetApp Files.
2. Ensure the following ports are open to the BlueXP classification instance:
  - For NFS – ports 111 and 2049.
  - For CIFS – ports 139 and 445.
3. Ensure that NFS volume export policies include the IP address of the BlueXP classification instance so it can access the data on each volume.
4. If you use CIFS, provide BlueXP classification with Active Directory credentials so it can scan CIFS volumes.
  - a. From the BlueXP left navigation menu, select **Governance > Classification**.
5. From the BlueXP classification menu, select **Configuration**.

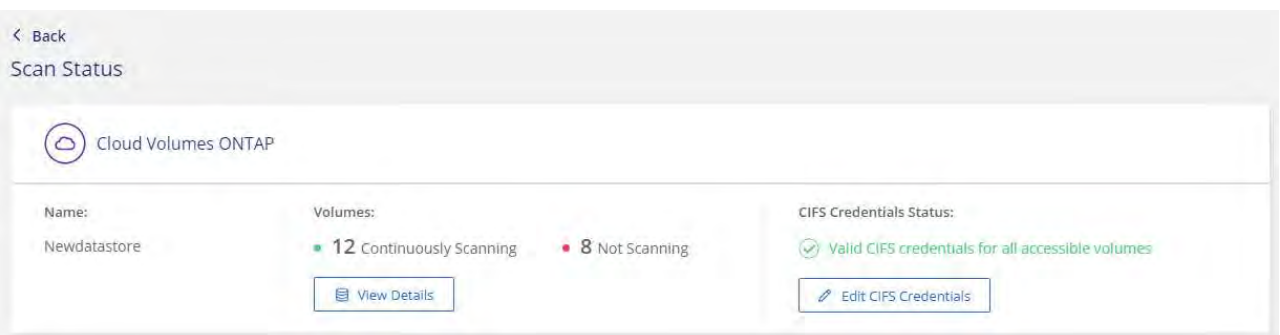


- a. For each working environment, select **Edit CIFS Credentials** and enter the user name and password that BlueXP classification needs to access CIFS volumes on the system.

The credentials can be read-only, but providing admin credentials ensures that BlueXP classification can read any data that requires elevated permissions. The credentials are stored on the BlueXP classification instance.

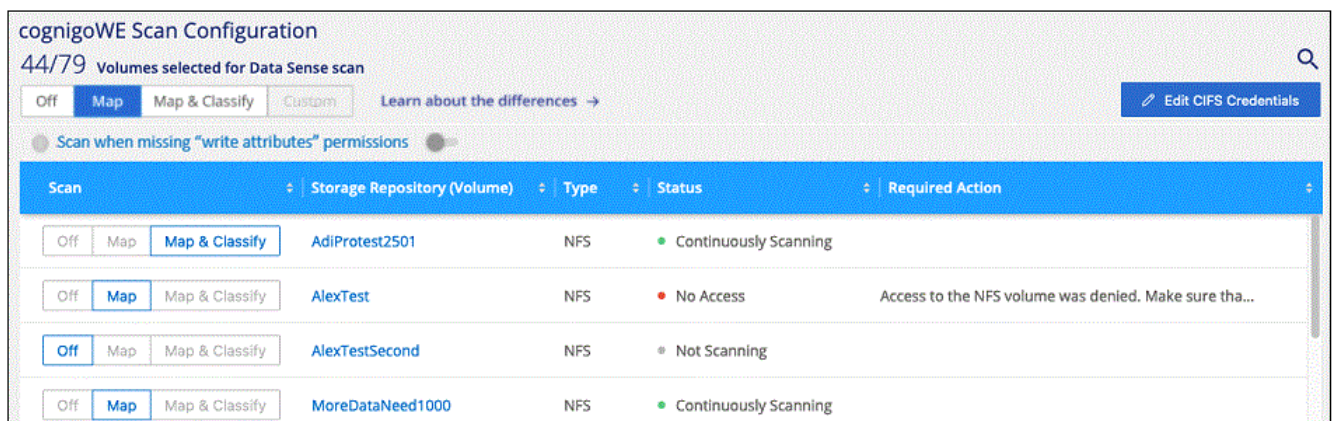
If you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, it's recommended the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, configure the Active Directory user as part of a parent group in the organization which has permissions to all files.

After you enter the credentials, you should see a message that all CIFS volumes were authenticated successfully.



6. On the Configuration page, select **View Details** to review the status for each CIFS and NFS volume and correct any errors.

For example, the following image shows four volumes; one of which BlueXP classification can't scan due to network connectivity issues between the BlueXP classification instance and the volume.



## Enable and disable compliance scans on volumes

You can start or stop mapping-only scans, or mapping and classification scans, in a working environment at any time from the Configuration page. You can also change from mapping-only scans to mapping and classification scans, and vice-versa. We recommend that you scan all volumes.



New volumes added to the working environment are automatically scanned only when you have set the **Map** or **Map & Classify** setting in the heading area. When set to **Custom** or **Off** in the heading area, you'll need to activate mapping and/or full scanning on each new volume you add in the working environment.

The switch at the top of the page for **Scan when missing "write attributes" permissions** is disabled by default. This means that if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, that the system won't scan the files because BlueXP classification can't revert the "last access time" to the original timestamp. If you don't care if the last access time is reset, turn the switch ON and all files are scanned regardless of the permissions. [Learn more](#).

Scan	Storage Repository (Volume)	Type	Status	Required Action
<input type="radio"/> Off <input checked="" type="radio"/> Map <input type="radio"/> Map & Classify	AdiNFSVoL_copy	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
<input type="radio"/> Off <input type="radio"/> Map <input checked="" type="radio"/> Map & Classify	AdiProtest2501	NFS	Continuously Scanning	
<input type="radio"/> Off <input checked="" type="radio"/> Map <input type="radio"/> Map & Classify	AlexTest	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
<input type="radio"/> Off <input type="radio"/> Map <input type="radio"/> Map & Classify	AlexTestSecond	NFS	Not Scanning	

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. Do one of the following:
  - To enable mapping-only scans on a volume, in the volume area, select **Map**. To enable on all volumes, in the heading area, select **Map**.
  - To enable full scanning on a volume, in the volume area, select **Map & Classify**. To enable on all volumes, in the heading area, select **Map & Classify**.
  - To disable scanning on a volume, in the volume area, select **Off**. To disable scanning on all volumes, in the heading area, select **Off**.

## Scan Amazon FSx for ONTAP volumes with BlueXP classification

Complete a few steps to get started scanning Amazon FSx for ONTAP volume with BlueXP classification.

### Before you begin

- You need an active Connector in AWS to deploy and manage BlueXP classification.
- The security group you selected when creating the working environment must allow traffic from the BlueXP



classification instance. You can find the associated security group using the ENI connected to the FSx for ONTAP file system and edit it using the AWS Management Console.

[AWS security groups for Linux instances](#)

[AWS security groups for Windows instances](#)

[AWS elastic network interfaces \(ENI\)](#)

- Ensure the following ports are open to the BlueXP classification instance:
  - For NFS – ports 111 and 2049.
  - For CIFS – ports 139 and 445.

## Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

You should deploy BlueXP classification in the same AWS network as the Connector for AWS and the FSx volumes you wish to scan.

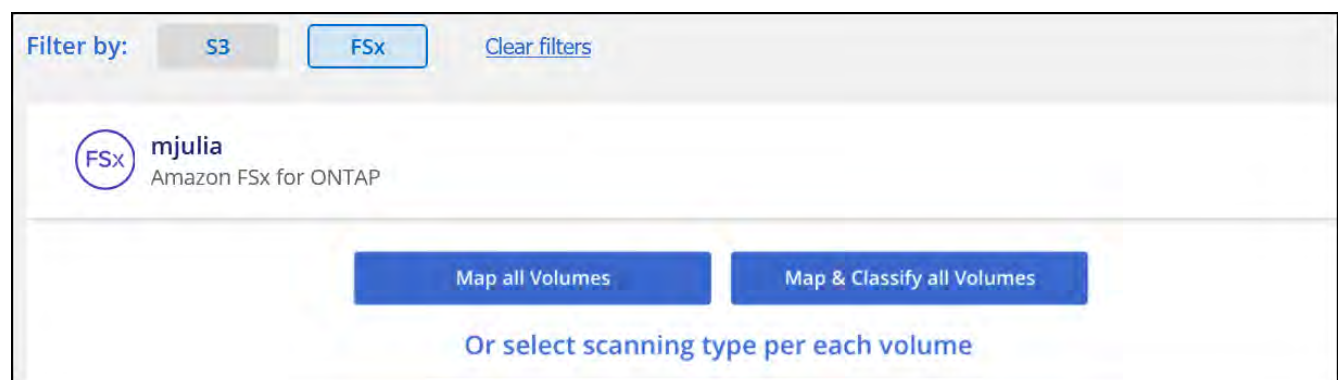
**Note:** Deploying BlueXP classification in an on-premises location is not currently supported when scanning FSx volumes.

Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

## Enable BlueXP classification in your working environments

You can enable BlueXP classification for FSx for ONTAP volumes.

1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. From the BlueXP classification menu, select **Configuration**.



3. Select how you want to scan the volumes in each working environment. [Learn about mapping and classification scans](#):
  - To map all volumes, click **Map all Volumes**.
  - To map and classify all volumes, click **Map & Classify all Volumes**.
  - To customize scanning for each volume, click **Or select scanning type for each volume**, and then choose the volumes you want to map and/or classify.
4. In the confirmation dialog box, click **Approve** to have BlueXP classification start scanning your volumes.

## Result

BlueXP classification starts scanning the volumes you selected in the working environment. Results will be available in the Compliance dashboard as soon as BlueXP classification finishes the initial scans. The time that it takes depends on the amount of data—it could be a few minutes or hours. You can track the progress of the initial scan by navigating to the **Configuration** menu then selecting the **Working Environment configuration**. The progress of each scan is shown as a progress bar. You can also hover over the progress bar to see the number of files scanned relative to the total files in the volume.



- By default, if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, the system won't scan the files in your volumes because BlueXP classification can't revert the "last access time" to the original timestamp. If you don't care if the last access time is reset, click **Or select scanning type for each volume**. The resulting page has a setting you can enable so that BlueXP classification will scan the volumes regardless of permissions.
- BlueXP classification scans only one file share under a volume. If you have multiple shares in your volumes, you'll need to scan those other shares separately as a shares group. [See more details about this BlueXP classification limitation](#).

## Verify that BlueXP classification has access to volumes

Make sure BlueXP classification can access volumes by checking your networking, security groups, and export policies.

You'll need to provide BlueXP classification with CIFS credentials so it can access CIFS volumes.

## Steps

1. From the BlueXP classification menu, select **Configuration**.
2. On the Configuration page, select **View Details** to review the status and correct any errors.

For example, the following image shows a volume BlueXP classification can't scan due to network connectivity issues between the BlueXP classification instance and the volume.

Scan	Storage Repository (Volume)	Type	Status	Required Action
Off	jrncione	NFS	No Access	Check network connectivity between the Data Sense ...

3. Make sure there's a network connection between the BlueXP classification instance and each network that includes volumes for FSx for ONTAP.



For FSx for ONTAP, BlueXP classification can scan volumes only in the same region as BlueXP.

4. Ensure NFS volume export policies include the IP address of the BlueXP classification instance so it can access the data on each volume.
5. If you use CIFS, provide BlueXP classification with Active Directory credentials so it can scan CIFS volumes.
  - a. From the BlueXP classification menu, select **Configuration**.
  - b. For each working environment, select **Edit CIFS Credentials** and enter the user name and password that BlueXP classification needs to access CIFS volumes on the system.

The credentials can be read-only, but providing admin credentials ensures that BlueXP classification

can read any data that requires elevated permissions. The credentials are stored on the BlueXP classification instance.

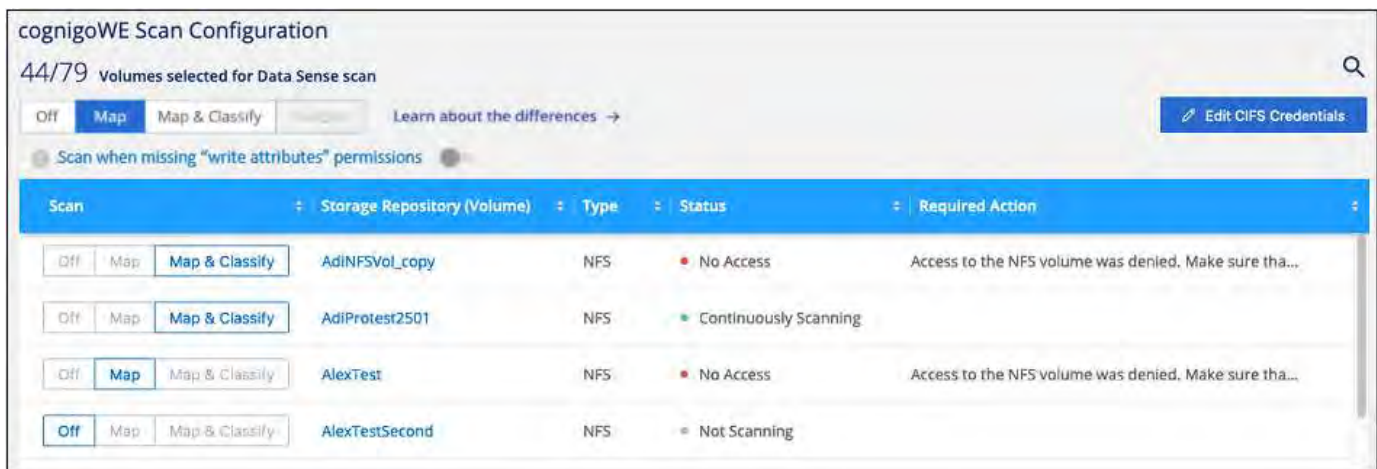
If you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, it's recommended the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, configure the Active Directory user as part of a parent group in the organization which has permissions to all files.

After you enter the credentials, you should see a message that all CIFS volumes were authenticated successfully.

## Enable and disable compliance scans on volumes

You can start or stop mapping-only scans, or mapping and classification scans, in a working environment at any time from the Configuration page. You can also change from mapping-only scans to mapping and classification scans, and vice-versa. We recommend that you scan all volumes.

The switch at the top of the page for **Scan when missing "write attributes" permissions** is disabled by default. This means that if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, that the system won't scan the files because BlueXP classification can't revert the "last access time" to the original timestamp. If you don't care if the last access time is reset, turn the switch ON and all files are scanned regardless of the permissions. [Learn more.](#)



Scan	Storage Repository (Volume)	Type	Status	Required Action
<input type="radio"/> Off <input type="radio"/> Map <input type="radio"/> Map & Classify	AdiNFSVol_copy	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
<input type="radio"/> Off <input type="radio"/> Map <input type="radio"/> Map & Classify	AdiProtest2501	NFS	Continuously Scanning	
<input type="radio"/> Off <input type="radio"/> Map <input type="radio"/> Map & Classify	AlexTest	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
<input type="radio"/> Off <input type="radio"/> Map <input type="radio"/> Map & Classify	AlexTestSecond	NFS	Not Scanning	

1. From the BlueXP classification menu, select **Configuration**.
2. In the Configuration page, locate the working environment with the volumes you want to scan.
3. Do one of the following:
  - To enable mapping-only scans on a volume, in the volume area, select **Map**. Or, to enable on all volumes, in the heading area, select **Map**.  
To enable full scanning on a volume, in the volume area, select **Map & Classify**. Or, to enable on all volumes, in the heading area, select **Map & Classify**.
  - To disable scanning on a volume, in the volume area, select **Off**. To disable scanning on all volumes, in the heading area, select **Off**.

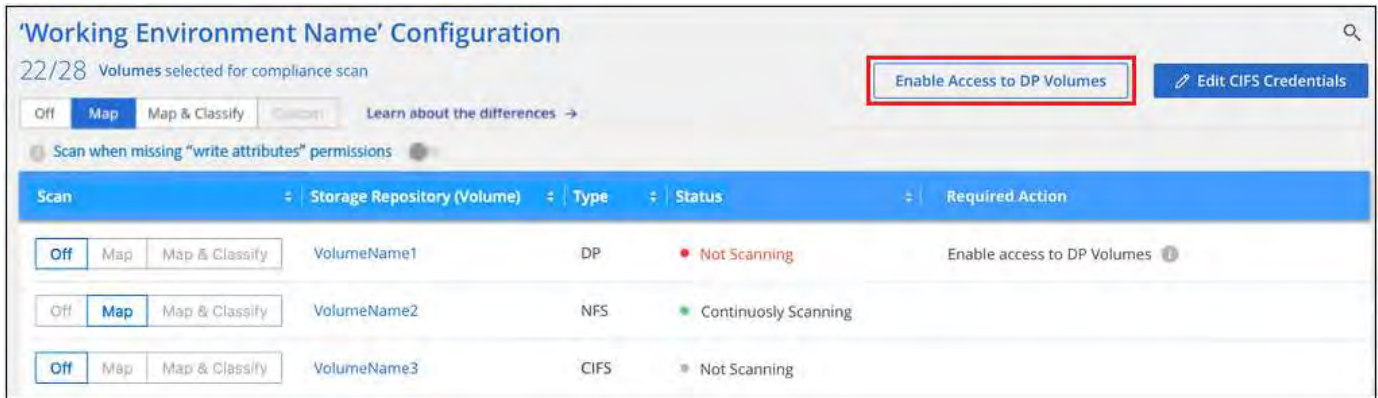


New volumes added to the working environment are automatically scanned only when you have set the **Map** or **Map & Classify** setting in the heading area. When set to **Custom** or **Off** in the heading area, you'll need to activate mapping and/or full scanning on each new volume you add in the working environment.

## Scan data protection volumes

By default, data protection (DP) volumes are not scanned because they are not exposed externally and BlueXP classification cannot access them. These are the destination volumes for SnapMirror operations from an FSx for ONTAP file system.

Initially, the volume list identifies these volumes as *Type DP* with the *Status Not Scanning* and the *Required Action Enable Access to DP volumes*.



## Steps

If you want to scan these data protection volumes:

1. From the BlueXP classification menu, select **Configuration**.
2. Select **Enable Access to DP volumes** at the top of the page.
3. Review the confirmation message and select **Enable Access to DP volumes** again.
  - Volumes that were initially created as NFS volumes in the source FSx for ONTAP file system are enabled.
  - Volumes that were initially created as CIFS volumes in the source FSx for ONTAP file system require that you enter CIFS credentials to scan those DP volumes. If you already entered Active Directory credentials so that BlueXP classification can scan CIFS volumes you can use those credentials, or you can specify a different set of Admin credentials.

The screenshot shows the 'Provide Active Directory Credentials' dialog box. It has two radio buttons: 'Use existing CIFS Scanning Credentials (user1@domain2)' (selected) and 'Use Custom Credentials'. Below are fields for 'Active Directory Domain' and 'DNS IP Address'. A blue button 'Enable Access to DP Volumes' and a 'Cancel' button are at the bottom.

The screenshot shows the 'Provide Active Directory Credentials' dialog box. It has two radio buttons: 'Use existing CIFS Scanning Credentials (user1@domain2)' and 'Use Custom Credentials' (selected). Below are fields for 'Username' and 'Password', and 'Active Directory Domain' and 'DNS IP Address'. A blue button 'Enable Access to DP Volumes' and a 'Cancel' button are at the bottom.

4. Activate each DP volume that you want to scan.

## Result

Once enabled, BlueXP classification creates an NFS share from each DP volume that was activated for

scanning. The share export policies only allow access from the BlueXP classification instance.

If you had no CIFS data protection volumes when you initially enabled access to DP volumes, and later add some, the button **Enable Access to CIFS DP** appears at the top of the Configuration page. Select this button and add CIFS credentials to enable access to these CIFS DP volumes.



Active Directory credentials are registered only in the storage VM of the first CIFS DP volume, so all DP volumes on that SVM will be scanned. Any volumes that reside on other SVMs will not have the Active Directory credentials registered, so those DP volumes won't be scanned.

## Scan Cloud Volumes ONTAP and on-premises ONTAP volumes with BlueXP classification

Complete a few steps to start scanning your Cloud Volumes ONTAP and on-premises ONTAP volumes using BlueXP classification.

### Prerequisites

Before you enable BlueXP classification, make sure you have a supported configuration.

- If you are scanning Cloud Volumes ONTAP and on-premises ONTAP systems that are accessible over the internet, you can [deploy BlueXP classification in the cloud](#) or [in an on-premises location that has internet access](#).
- If you are scanning on-premises ONTAP systems that have been installed in a dark site that has no internet access, you need to [deploy BlueXP classification in the same on-premises location that has no internet access](#). This also requires that the BlueXP Connector is deployed in that same on-premises location.

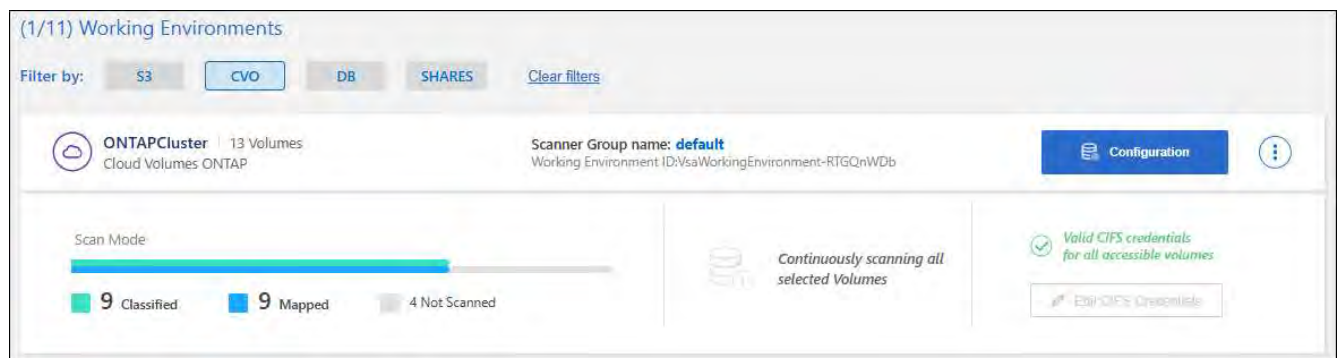
### Enable BlueXP classification scanning in your working environments

You can enable BlueXP classification scanning on Cloud Volumes ONTAP systems in any supported cloud provider, and on on-premises ONTAP clusters.

### Steps

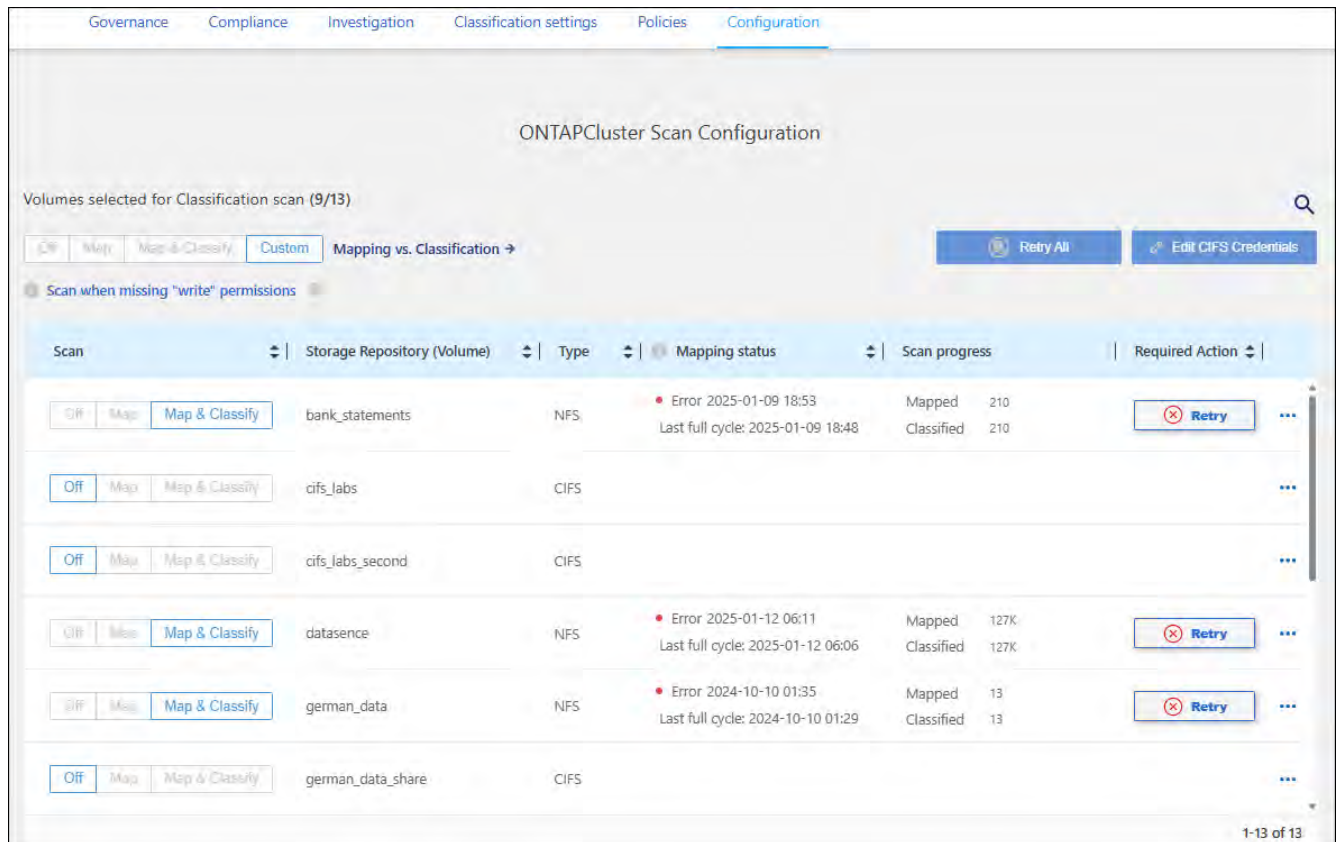
1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. From the BlueXP classification menu, select **Configuration**.

The Configuration page shows multiple working environments.



3. Choose a working environment and select **Configuration**.





- If you don't care if the last access time is reset, turn the **Scan when missing "write attributes" permissions** switch ON and all files are scanned regardless of the permissions.

The switch at the top of the page for **Scan when missing "write attributes" permissions** is disabled by default. This means that if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, that the system won't classify the files because BlueXP classification can't revert the "last access time" to the original timestamp. [Learn more.](#)

- Select how you want to scan the volumes in each working environment. [Learn about mapping and classification scans:](#)
  - To map all volumes, select **Map**.
  - To map and classify all volumes, select **Map & Classify**.
  - To customize scanning for each volume, select **Custom**, and then choose the volumes you want to map and/or classify.
- In the confirmation dialog box, select **Approve** to have BlueXP classification start scanning your volumes.

## Result

BlueXP classification starts scanning the volumes you selected in the working environment. Results start to appear in the Compliance dashboard as soon as BlueXP classification starts the scan. The time that it takes to complete depends on the amount of data—it could be a few minutes or hours.



BlueXP classification scans only one file share under a volume. If you have multiple shares in your volumes, you'll need to scan those other shares separately as a shares group. [See more details about this BlueXP classification limitation.](#)

## Verify that BlueXP classification has access to volumes

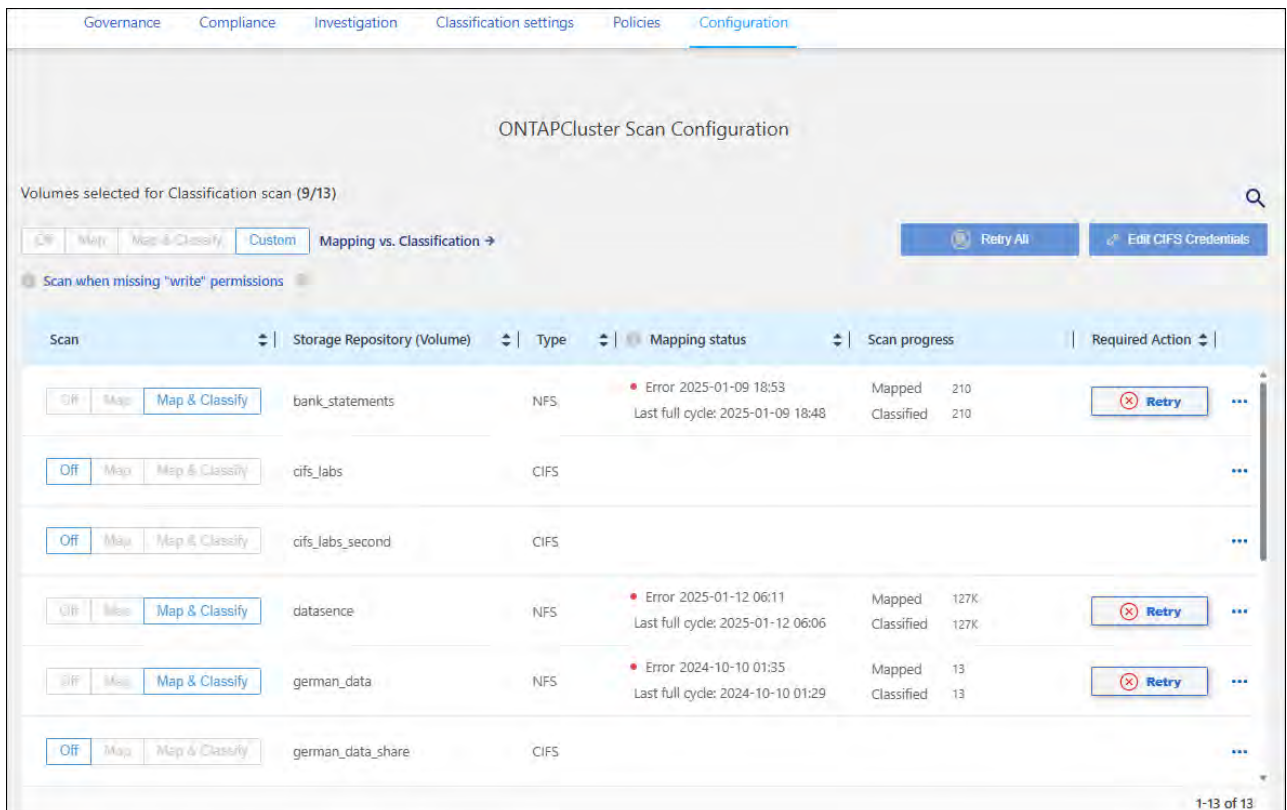
Make sure that BlueXP classification can access volumes by checking your networking, security groups, and export policies. You'll need to provide BlueXP classification with CIFS credentials so it can access CIFS volumes.

### Steps

1. Make sure that there's a network connection between the BlueXP classification instance and each network that includes volumes for Cloud Volumes ONTAP or on-prem ONTAP clusters.
2. Ensure that the security group for Cloud Volumes ONTAP allows inbound traffic from the BlueXP classification instance.

You can either open the security group for traffic from the IP address of the BlueXP classification instance, or you can open the security group for all traffic from inside the virtual network.

3. Ensure that NFS volume export policies include the IP address of the BlueXP classification instance so it can access the data on each volume.
4. If you use CIFS, provide BlueXP classification with Active Directory credentials so it can scan CIFS volumes.
  - a. From the BlueXP left navigation menu, select **Governance > Classification**.
  - b. From the BlueXP classification menu, select **Configuration**.



The screenshot displays the 'ONTAPCluster Scan Configuration' page in the BlueXP interface. It shows a table of volumes selected for classification scan (9/13). The table includes columns for Scan status, Storage Repository (Volume), Type, Mapping status, Scan progress, and Required Action. The 'Scan' column has buttons for 'Off', 'Map', and 'Map & Classify'. The 'Mapping status' column shows error messages for some volumes, such as 'Error 2025-01-09 18:53' for 'bank\_statements' and 'Error 2025-01-12 06:11' for 'datasence'. The 'Scan progress' column shows 'Mapped' and 'Classified' counts. The 'Required Action' column has a 'Retry' button for volumes with errors. There are also buttons for 'Retry All' and 'Edit CIFS Credentials' at the top right of the table area.

Scan	Storage Repository (Volume)	Type	Mapping status	Scan progress	Required Action
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	bank_statements	NFS	<span style="color: red;">•</span> Error 2025-01-09 18:53 Last full cycle: 2025-01-09 18:48	Mapped 210 Classified 210	<input type="button" value="Retry"/> <input type="button" value="..."/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	cifs_labs	CIFS			<input type="button" value="..."/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	cifs_labs_second	CIFS			<input type="button" value="..."/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	datasence	NFS	<span style="color: red;">•</span> Error 2025-01-12 06:11 Last full cycle: 2025-01-12 06:06	Mapped 127K Classified 127K	<input type="button" value="Retry"/> <input type="button" value="..."/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	german_data	NFS	<span style="color: red;">•</span> Error 2024-10-10 01:35 Last full cycle: 2024-10-10 01:29	Mapped 13 Classified 13	<input type="button" value="Retry"/> <input type="button" value="..."/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	german_data_share	CIFS			<input type="button" value="..."/>

- c. For each working environment, select **Edit CIFS Credentials** and enter the user name and password that BlueXP classification needs to access CIFS volumes on the system.

The credentials can be read-only, but providing admin credentials ensures that BlueXP classification can read any data that requires elevated permissions. The credentials are stored on the BlueXP classification instance.

If you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, it's recommended the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, configure the Active Directory user as part of a parent group in the organization which has permissions to all files.

After you enter the credentials, you should see a message that all CIFS volumes were authenticated successfully.

5. On the Configuration page, select **Configuration** to review the status for each CIFS and NFS volume and correct any errors.

### Disable compliance scans on volumes

You can start or stop mapping-only scans, or mapping and classification scans, in a working environment at any time from the Configuration page. You can also change from mapping-only scans to mapping and classification scans, and vice-versa. We recommend that you scan all volumes.



New volumes added to the working environment are automatically scanned only when you have set the **Map** or **Map & Classify** setting in the heading area. When the option is set to **Custom** or **Off** in the heading area, you'll need to activate mapping and/or full scanning on each new volume you add in the working environment.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. Select the **Configuration** button for the working environment that you want to change.

The screenshot shows the 'ONTAPCluster Scan Configuration' page. At the top, there are navigation tabs: Governance, Compliance, Investigation, Classification settings, Policies, and Configuration. Below the tabs, it says 'Volumes selected for Classification scan (9/13)'. There are buttons for 'Off', 'Map', 'Map & Classify', and 'Custom'. A 'Mapping vs. Classification' link is also present. On the right, there are 'Retry All' and 'Edit CIFS Credentials' buttons. Below this is a table with columns: Scan, Storage Repository (Volume), Type, Mapping status, Scan progress, and Required Action. The table lists several volumes, some with error messages in the Mapping status column.

Scan	Storage Repository (Volume)	Type	Mapping status	Scan progress	Required Action
Off   Map   Map & Classify	bank_statements	NFS	Error 2025-01-09 18:53 Last full cycle: 2025-01-09 18:48	Mapped 210 Classified 210	Retry
Off   Map   Map & Classify	cifs_labs	CIFS			
Off   Map   Map & Classify	cifs_labs_second	CIFS			
Off   Map   Map & Classify	datasence	NFS	Error 2025-01-12 06:11 Last full cycle: 2025-01-12 06:06	Mapped 127K Classified 127K	Retry
Off   Map   Map & Classify	german_data	NFS	Error 2024-10-10 01:35 Last full cycle: 2024-10-10 01:29	Mapped 13 Classified 13	Retry
Off   Map   Map & Classify	german_data_share	CIFS			

3. Do one of the following:



- To disable scanning on a volume, in the volume area, select **Off**.
- To disable scanning on all volumes, in the heading area, select **Off**.

## Scan database schemas with BlueXP classification

Complete a few steps to start scanning your database schemas with BlueXP classification.

### Review prerequisites

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.

### Supported databases

BlueXP classification can scan schemas from the following databases:

- Amazon Relational Database Service (Amazon RDS)
- MongoDB
- MySQL
- Oracle
- PostgreSQL
- SAP HANA
- SQL Server (MSSQL)



The statistics gathering feature **must be enabled** in the database.

### Database requirements

Any database with connectivity to the BlueXP classification instance can be scanned, regardless of where it is hosted. You just need the following information to connect to the database:

- IP Address or host name
- Port
- Service name (only for accessing Oracle databases)
- Credentials that allow read access to the schemas

When choosing a user name and password, it's important to choose one that has full read permissions to all the schemas and tables you want to scan. We recommend that you create a dedicated user for the BlueXP classification system with all the required permissions.

**Note:** For MongoDB, a read-only Admin role is required.

### Deploy the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

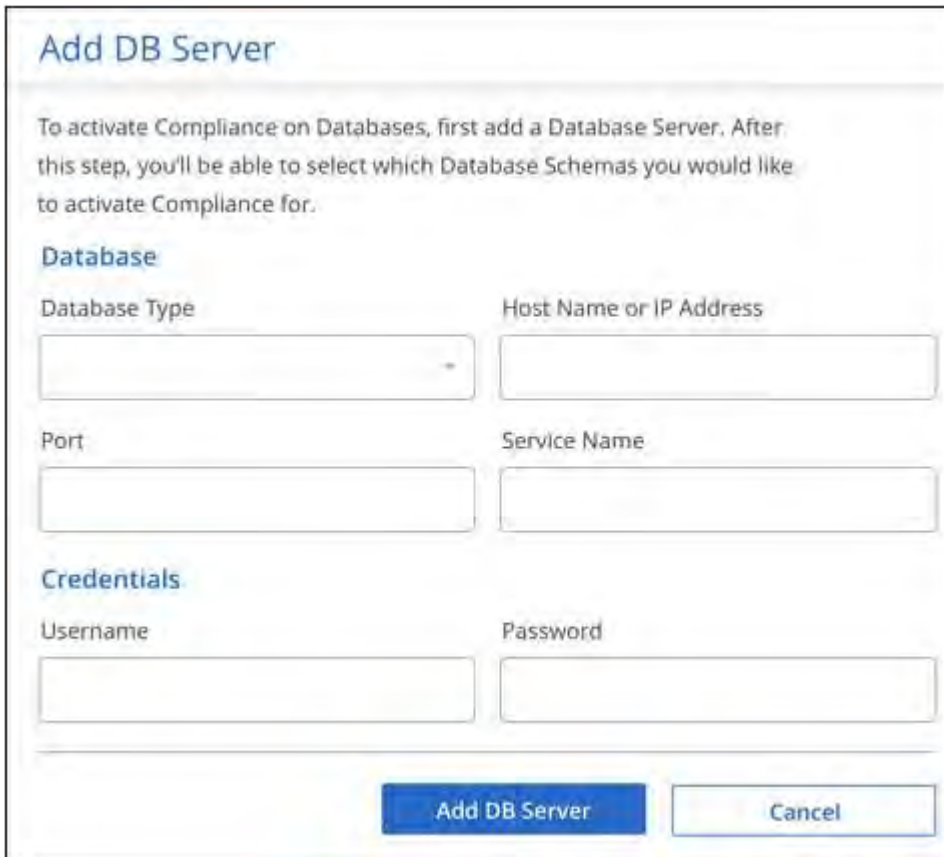
If you are scanning database schemas that are accessible over the internet, you can [deploy BlueXP classification in the cloud](#) or [deploy BlueXP classification in an on-premises location that has internet access](#).

If you are scanning database schemas that have been installed in a dark site that has no internet access, you need to [deploy BlueXP classification in the same on-premises location that has no internet access](#). This also requires that the BlueXP Connector is deployed in that same on-premises location.

### Add the database server

Add the database server where the schemas reside.

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration page, select **Add Working Environment > Add Database Server**.
3. Enter the required information to identify the database server.
  - a. Select the database type.
  - b. Enter the port and the host name or IP address to connect to the database.
  - c. For Oracle databases, enter the Service name.
  - d. Enter the credentials so that BlueXP classification can access the server.
  - e. Click **Add DB Server**.



The screenshot shows a web form titled "Add DB Server". At the top, there is a blue header with the title. Below the header is a paragraph of instructional text: "To activate Compliance on Databases, first add a Database Server. After this step, you'll be able to select which Database Schemas you would like to activate Compliance for." The form is divided into three sections: "Database", "Credentials", and "Add DB Server" buttons. The "Database" section contains four input fields: "Database Type" (a dropdown menu), "Host Name or IP Address", "Port", and "Service Name". The "Credentials" section contains two input fields: "Username" and "Password". At the bottom right, there are two buttons: "Add DB Server" (a blue button) and "Cancel" (a white button with a blue border).

The database is added to the list of working environments.

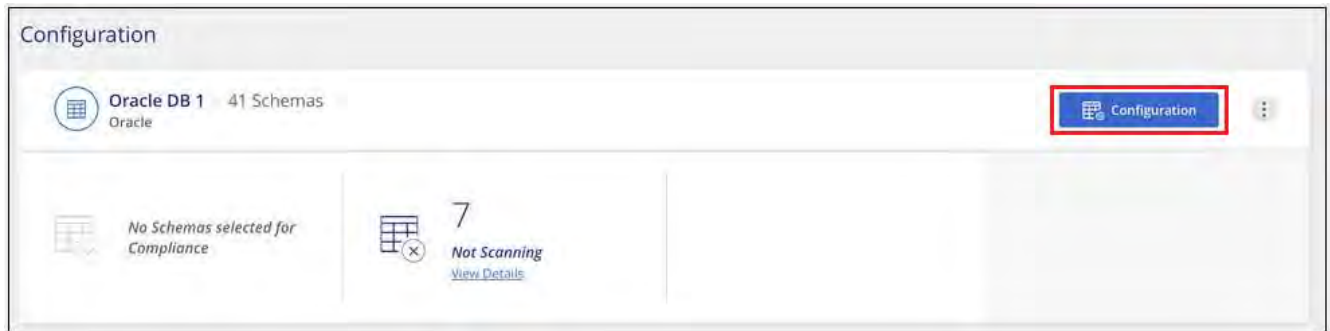
### Enable and disable compliance scans on database schemas

You can stop or start full scanning of your schemas at any time.




There is no option to select mapping-only scans for database schemas.

1. From the Configuration page, select the **Configuration** button for the database you want to configure.



2. Select the schemas that you want to scan by moving the slider to the right.



## Result

BlueXP classification starts scanning the database schemas that you enabled. You can track the progress of the initial scan by navigating to the **Configuration** menu then selecting the **Working Environment configuration**. The progress of each scan is shown as a progress bar. You can also hover over the progress bar to see the number of files scanned relative to the total files in the volume. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

BlueXP classification scans your databases once per day; databases are not continuously scanned like other data sources.

## Scan file shares with BlueXP classification

To scan file shares, you must first create a file shares group in BlueXP classification. File shares groups are for NFS or CIFS (SMB) shares hosted on-premises or in the cloud.



Scanning data from non-NetApp file shares is not supported in the BlueXP classification core version.

## Prerequisites

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.

- The shares can be hosted anywhere, including in the cloud or on-premises. CIFS shares from older NetApp 7-Mode storage systems can be scanned as file shares.
  - BlueXP classification can't extract permissions or the "last access time" from 7-Mode systems.
  - Because of a known issue between some Linux versions and CIFS shares on 7-Mode systems, you must configure the share to use only SMBv1 with NTLM authentication enabled.
- There needs to be network connectivity between the BlueXP classification instance and the shares.
- You can add a DFS (Distributed File System) share as a regular CIFS share. Because BlueXP classification is unaware that the share is built upon multiple servers/volumes combined as a single CIFS share, you might receive permission or connectivity errors about the share when the message really only applies to one of the folders/shares that is located on a different server/volume.
- For CIFS (SMB) shares, ensure that you have Active Directory credentials that provide read access to the shares. Admin credentials are preferred in case BlueXP classification needs to scan any data that requires elevated permissions.

If you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, it's recommended the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, configure the Active Directory user as part of a parent group in the organization which has permissions to all files.

- All CIFS file shares in a group must use the same Active Directory credentials.
- You can mix NFS and CIFS (using either Kerberos or NTLM) shares. You must add the shares to the group separately. That is, you must complete the process twice—once per protocol.
  - You cannot create a file shares group that mixes CIFS authentication types (Kerberos and NTLM).
- If you're using CIFS with Kerberos authentication, ensure the IP address provided is accessible to the BlueXP classification service. The file shares can't be added if the IP address is unreachable.

## Create a file shares group

When you add file shares to the group, you must use the format `<host_name>:/<share_path>`.

+

You can add file shares individually or you can enter a line-separated list of the file shares you want to scan. You can add up to 100 shares at a time.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration page, select **Add Working Environment > Add File Shares Group**.
3. In the Add File Shares Group dialog, enter the name for the group of shares then select **Continue**.
4. Select the protocol for the file shares you are adding.

**.If you're adding CIFS shares with NTLM authentication, enter the Active Directory credentials to access the CIFS volumes. Although read-only credentials are supported, it's recommended you provide full access with administrator credentials. Select Save.**

1. Add the file shares that you want to scan (one file share per line). Then select **Continue**.
2. A confirmation dialog displays the number of shares that were added.

If the dialog lists any shares that could not be added, capture this information so that you can resolve the issue. If the issue pertains to a naming convention, you can re-add the share with a corrected name.

### 3. Configure scanning on the volume:

- To enable mapping-only scans on file shares, select **Map**.
- To enable full scans on file shares, select **Map & Classify**.
- To disable scanning on file shares, select **Off**.



The switch at the top of the page for **Scan when missing "write attributes" permissions** is disabled by default. This means that if BlueXP classification doesn't have write attributes permissions in CIFS or write permissions in NFS, the system won't scan the files because BlueXP classification can't revert the "last access time" to the original timestamp.

If you switch **Scan when missing "write attributes" permissions** to **On**, the scan resets the last accessed time and scans all files regardless of permissions.

To learn more about the last accessed time stamp, see [xref:./Metadata collected from data sources in BlueXP classification](#).

## Result

BlueXP classification starts scanning the files in the file shares you added. You can [Track the scanning progress](#) and view the results of the scan in the **Dashboard**.



If the scan doesn't complete successfully for a CIFS configuration with Kerberos authentication, check the **Configuration** tab for errors.

## Edit a file shares group

After you create a file shares group, you can edit the CIFS protocol or add and remove file shares.

### Edit the CIFS protocol configuration

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration page, select the file shares group you want to modify.
3. Select **Edit CIFS Credentials**.

## Edit CIFS Authentication

Classification requires Active Directory credentials to access CIFS Volumes in Micky.

The credentials can be read-only, but providing admin credentials ensures that Classification can read any data that requires elevated permissions.

### Select Authentication Method

NTLM

Kerberos

Username

Password

domain\user or user@domain

Password

Save

Cancel

4. Choose the authentication method: **NTLM** or **Kerberos**.
5. Enter the Active Directory **Username** and **Password**.
6. Select **Save** to complete the process.

### Add file shares to compliance scans

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration page, select the file shares group you want to modify.
3. Select **+ Add shares**.
4. Select the protocol for the file shares you are adding.

If you're adding file shares to a protocol you've already configured, no changes are required.

If you're adding file shares with a second protocol, ensure you've properly configured the authentication properly as detailed in the [prerequisites](#).

5. Add the file shares you want to scan (one file share per line) using the format `<host_name>:/<share_path>`.
6. Select **Continue** to complete adding the file shares.

### Remove a file share from compliance scans

1. From the BlueXP classification menu, select **Configuration**.
2. Select the working environment from which you want to remove file shares.
3. Select **Configuration**.
4. From the Configuration page, select the Actions **...** for the file share you want to remove.
5. From the Actions menu, select **Remove Share**.

### Track the scanning progress

You can track the progress of the initial scan.

1. Select the **Configuration** menu.
2. Select the **Working Environment Configuration**.

The progress of each scan is shown as a progress bar.

3. Hover over the progress bar to see the number of files scanned relative to the total files in the volume.

### Scan StorageGRID data with BlueXP classification

Complete a few steps to start scanning data within StorageGRID directly with BlueXP classification.

#### Review StorageGRID requirements

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.

- You need to have the endpoint URL to connect with the object storage service.
- You need to have the Access Key and Secret Key from StorageGRID so that BlueXP classification can access the buckets.

#### Deploy the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

If you are scanning data from StorageGRID that is accessible over the internet, you can [deploy BlueXP classification in the cloud](#) or [deploy BlueXP classification in an on-premises location that has internet access](#).

If you are scanning data from StorageGRID that has been installed in a dark site that has no internet access, you need to [deploy BlueXP classification in the same on-premises location that has no internet access](#). This also requires that the BlueXP Connector is deployed in that same on-premises location.

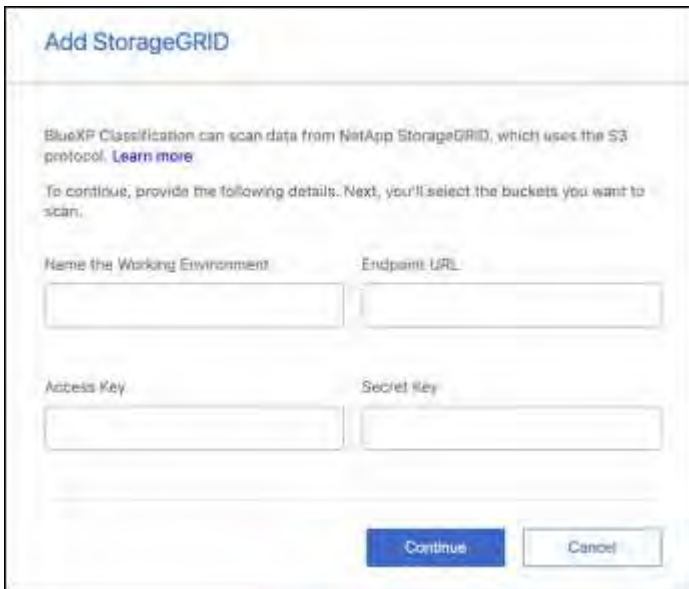
#### Add the StorageGRID service to BlueXP classification

Add the StorageGRID service.

##### Steps

1. From the BlueXP classification menu, select the **Configuration** option.
2. From the Configuration page, select **Add Working Environment > Add StorageGRID**.
3. In the Add StorageGRID Service dialog, enter the details for the StorageGRID service and click **Continue**.

- a. Enter the name you want to use for the Working Environment. This name should reflect the name of the StorageGRID service to which you are connecting.
- b. Enter the Endpoint URL to access the object storage service.
- c. Enter the Access Key and Secret Key so that BlueXP classification can access the buckets in StorageGRID.



The screenshot shows a web form titled "Add StorageGRID". Below the title, there is explanatory text: "BlueXP Classification can scan data from NetApp StorageGRID, which uses the S3 protocol. [Learn more](#)". This is followed by instructions: "To continue, provide the following details. Next, you'll select the buckets you want to scan." The form contains four input fields arranged in a 2x2 grid. The top row has "Name the Working Environment" and "Endpoint URL". The bottom row has "Access Key" and "Secret Key". At the bottom right of the form are two buttons: "Continue" (highlighted in blue) and "Cancel".

## Result

StorageGRID is added to the list of working environments.

## Enable and disable compliance scans on StorageGRID buckets

After you enable BlueXP classification on StorageGRID, the next step is to configure the buckets that you want to scan. BlueXP classification discovers those buckets and displays them in the working environment you created.

## Steps

1. In the Configuration page, locate the StorageGRID working environment.
2. On the StorageGRID working environment tile, select **Configuration**.



Buckets selected for Classification scan (5/8)

Scan	Storage Repository (Bucket)	Mapping status	Classification status	Required Action
Off   Map   <b>Map &amp; Classify</b>	bucketadipro	Finished 2024-09-05 10:33 Last full cycle: 2024-09-05 10:33	Mapped: 84 Classified: 5	...
Off   Map   <b>Map &amp; Classify</b>	datasense-0-files	Finished 2024-09-05 08:00 Last full cycle: 2024-09-05 08:00		...
Off   Map   <b>Map &amp; Classify</b>	datasense-10tb	Running 2024-09-04 07:25	Mapped: 3.7M Classified: 2.1M	...
Off   <b>Map</b>   Map & Classify	datasense-1tb	Running 2024-09-05 09:05 Last full cycle: 2024-09-05 03:04	Mapped: 1.3M	...
Off   <b>Map</b>   Map & Classify	datasense-1tb-2	Running 2024-09-05 09:06 Last full cycle: 2024-09-05 03:05	Mapped: 1.3M	...
Off   Map   Map & Classify	datasense-1tb-3	Not scanning		...

3. Complete one of the following steps to enable or disable scanning:

- To enable mapping-only scans on a bucket, select **Map**.
- To enable full scans on a bucket, select **Map & Classify**.
- To disable scanning on a bucket, select **Off**.

### Result

BlueXP classification starts scanning the buckets that you enabled. You can track the progress of the initial scan by navigating to the **Configuration** menu then selecting the **Working Environment configuration**. The progress of each scan is shown as a progress bar. You can also hover over the progress bar to see the number of files scanned relative to the total files in the volume. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

## Integrate your Active Directory with BlueXP classification

You can integrate a global Active Directory with BlueXP classification to enhance the results that BlueXP classification reports about file owners and which users and groups have access to your files.

When you set up certain data sources (listed below), you need to enter Active Directory credentials in order for BlueXP classification to scan CIFS volumes. This integration provides BlueXP classification with file owner and permissions details for the data that resides in those data sources. The Active Directory entered for those data sources might differ from the global Active Directory credentials you enter here. BlueXP classification will look in all integrated Active Directories for user and permission details.

This integration provides additional information in the following locations in BlueXP classification:

- You can use the "File Owner" [filter](#) and see results in the file's metadata in the Investigation pane. Instead of the file owner containing the SID (Security Identifier), it is populated with the actual user name.

You can also view more details about the file owner: account name, email address, and SAM account name, or view items owned by that user.

- You can see [full file permissions](#) for each file and directory when you click the "View all Permissions"

button.

- In the [Governance dashboard](#), the Open Permissions panel will show a greater level of detail about your data.



Local user SIDs, and SIDs from unknown domains, are not translated to the actual user name.

## Supported data sources

An Active Directory integration with BlueXP classification can identify data from within the following data sources:

- On-premises ONTAP systems
- Cloud Volumes ONTAP
- Azure NetApp Files
- FSx for ONTAP
- OneDrive accounts and SharePoint accounts (for legacy versions 1.30 and earlier)

There is no support for identifying user and permission information from Database schemas, Google Drive accounts, Amazon S3 accounts, or Object Storage that uses the Simple Storage Service (S3) protocol.

## Connect to your Active Directory server

After you've deployed BlueXP classification and have activated scanning on your data sources, you can integrate BlueXP classification with your Active Directory. Active Directory can be accessed using a DNS Server IP address or an LDAP Server IP address.

The Active Directory credentials can be read-only, but providing admin credentials ensures that BlueXP classification can read any data that requires elevated permissions. The credentials are stored on the BlueXP classification instance.

For CIFS volumes/file shares, if you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, we recommend that the user has Write Attributes permission. If possible, we recommend making the Active Directory configured user part of a parent group in the organization which has permissions to all files.

### Requirements

- You must have an Active Directory already set up for the users in your company.
- You must have the information for the Active Directory:
  - DNS Server IP address, or multiple IP addressesor
  - LDAP Server IP address, or multiple IP addresses
  - User Name and Password to access the server
  - Domain Name (Active Directory Name)
  - Whether you are using secure LDAP (LDAPS) or not
  - LDAP Server Port (typically 389 for LDAP, and 636 for secure LDAP)

- The following ports must be open for outbound communication by the BlueXP classification instance:

Protocol	Port	Destination	Purpose
TCP & UDP	389	Active Directory	LDAP
TCP	636	Active Directory	LDAP over SSL
TCP	3268	Active Directory	Global Catalog
TCP	3269	Active Directory	Global Catalog over SSL

## Steps

1. From the BlueXP classification Configuration page, click **Add Active Directory**.



2. In the Connect to Active Directory dialog, enter the Active Directory details and click **Connect**.

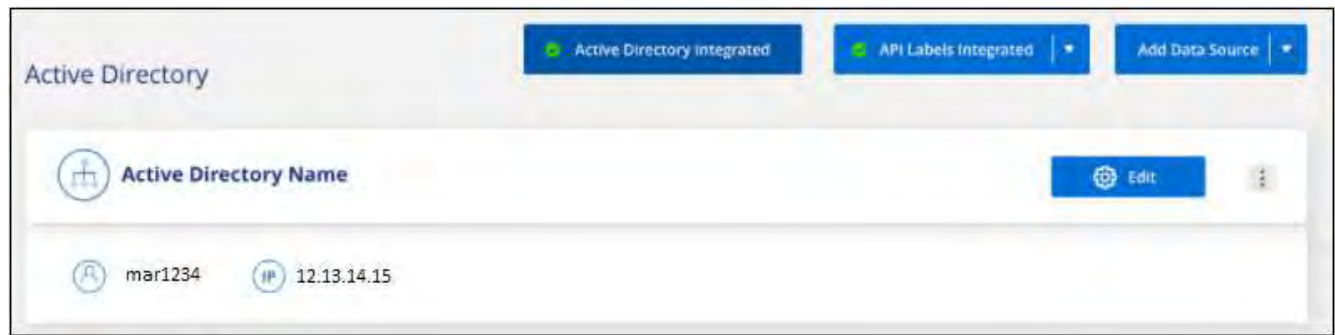
You can add multiple IP addresses, if required, by clicking **Add IP**.

 A screenshot of the "Connect to Active Directory" dialog box. The title is "Connect to Active Directory". It has several input fields:
 

- "Username" with the value "mar1234".
- "Password" with a masked value "\*\*\*\*\*".
- "DNS Server IP address:" with a radio button selected, containing the value "12.20.70.00" and a "+ Add IP" button.
- "Domain Name" with the value "mar@netapp.com".
- "LDAP Server IP Address:" with a radio button unselected and an empty field with a "+ Add IP" button.
- "LDAP Server Port" with the value "389".
- "LDAP Secure Connection" checkbox, which is unchecked.


 At the bottom right, there are two buttons: "Connect" (highlighted with a red rectangular box) and "Cancel".

BlueXP classification integrates to the Active Directory, and a new section is added to the Configuration page.



## Manage your Active Directory integration

If you need to modify any values in your Active Directory integration, click the **Edit** button and make the changes.

You can also delete the integration selecting the  button then **Remove Active Directory**.

# Use BlueXP classification

## View governance details about the data stored in your organization with BlueXP classification

Gain control of the costs related to the data on your organization's storage resources. BlueXP classification identifies the amount of stale data, non-business data, duplicate files, and very large files in your systems so you can decide whether you want to remove or tier some files to less expensive object storage.

This is where you should begin your research. From the Governance dashboard, you can select an area for further investigation.

Additionally, if you're planning to migrate data from on-premises locations to the cloud, you can view the size of the data and whether any of the data contains sensitive information before moving it.

### Review the Governance dashboard

The Governance dashboard provides information so that you can increase the efficiency and control the costs related to the data stored on your storage resources.

**NetApp BlueXP** | Demo | BlueXP Search | Organization: Demo SIM | Project: group1 | Connector: OCCMassDem... | Notifications | Settings | Help | Logout

**Classification** | Governance | Compliance | Investigation | Classification settings | Policies | Configuration

---

### Savings Opportunities

Scale Data

216.2K items | 227.8 GB

Optimize Storage

Non-Business Data

18.6K items | 24.3 GB

Optimize Storage

Duplicate Files

159.8K items | 188.6 GB

Optimize Storage

### Policies

View All

bb1	132.3K items
Data 3-5 years old	105.9K items

---

### Data Overview

[Data Discovery Assessment Report](#) | 
 [Full Data Mapping Overview Report](#) | 
 Scanned: 265.3 GB | 
 270.6K Files | 
 141 Tables

#### Top Data Repositories by Sensitivity Level

Repository	Non-Sensitive	Personal	Sensitive	Total Items
CVD	~50%	~30%	~20%	133.8K items
Amazon S3	~50%	~30%	~20%	131.5K items
File Shares	~50%	~30%	~20%	5.2K items
Database	~50%	~30%	~20%	141 items

---

#### Sensitive Data and Wide Permissions

Y-axis: Detected sensitive classes (Sensitive to Not sensitive)  
X-axis: Access level (Restrictive to Permissive)  
Z-axis: Number of files (0 to 10K)

#### Open Permissions

■ 51% - No Open Permissions | 
 ■ 48% - Open to Organization | 
 ■ 1% - Open to Public

---

#### Age of Data

Modified | Created | Last Accessed

Age	Count
30 days	~10K
31-60 days	~10K
61-90 days	~10K
91-180 days	~10K
181-365 days	~50K
1-3 years	~100K
3-5 years	~100K
5-7 years	~100K
Over 7 years	~10K

#### Size of Data

Category	Size
PDF files	~20K
XML files	~20K
CSV files	~210K
JSON files	~20K
Image files	~5K
Spreadsheet files	~5K
Log files	~5K
Other files	~5K

---

### Classification

#### 40 Categories

View All

Miscellaneous	18.2K items
Code	18K items
Bank Statements	13K items
Miscellaneous Spreadsheets	5.2K items

#### 127 File Types

View All

PDF	99.9K items
TXT	88.9K items
DOCX	31.4K items
MAP	16K items

Classification Documentation | v1.41.0

## Steps

1. From the BlueXP menu, select **Governance > Classification**.
2. Select **Governance**.

The Governance dashboard appears.

## Review savings opportunities

The *Saving Opportunities* component shows data that you can delete or tier to less expensive object storage. The data in *Saving Opportunities* update every 2 hours and can be manually updated.

## Steps

1. From the BlueXP classification menu, select **Governance**.
2. Within each Savings Opportunities tile of the Governance dashboard, select **Optimize Storage** to view the filtered results in the Investigation page. To discover any data you should delete or tier to less expensive storage, investigate the the *Saving Opportunities*.
  - **Stale Data** - Data that was last modified over 3 years ago.
  - **Non-Business Data** - Data considered not to be business related, based on their Category or File Type. This includes:
    - Application Data
    - Audio
    - Executables
    - Images
    - Logs
    - Videos
    - Miscellaneous (general "other" category)
  - **Duplicate Files** - Files that are duplicated in other locations in the data sources you are scanning. [See what types of duplicate files are displayed.](#)



If any of your data sources implement data tiering, old data that already resides in object storage can be identified in the *Stale Data* category.

## Review saved searches with the largest number of results

In the *Saved searches* tab, the searches with the greatest number of results appear at the top of the list. This data updates every two hours.

For details about saved searches, see [Create saved searches](#).

## Steps

1. From the BlueXP classification menu, select **Governance**.
2. In the Governance dashboard, locate the Saved Searches tile. Select the name of a saved search to display the results in the Investigation page.
3. Select **View All** to view the list of all available saved searches.  
In the *Saved searches* area, the searches with the greatest number of results appear at the top of the list.

## Create the Data Discovery Assessment Report

The Data Discovery Assessment Report provides a high-level analysis of the scanned environment to show areas of concern and potential remediation steps. The results are based on both mapping and classifying your data. The goal of this report is to raise awareness of three significant aspects of your dataset:


Feature	Description
Data governance concerns	A detailed picture of all the data you own and areas where you may be able to reduce the amount of data to save costs.
Data security exposures	Areas where your data is accessible to internal or external attacks because of broad access permissions.
Data compliance gaps	Where your personal or sensitive personal information is located for both security and for DSARs (data subject access requests).

Using this report, you might take the following actions:

- Reduce storage costs by changing your retention policy, or by moving or deleting certain data (stale, duplicate, or non-business data)
- Protect your data that has broad permissions by revising global group management policies
- Protect your data that has personal or sensitive personal information by moving PII to more secure data stores

### Steps

1. From the BlueXP menu, select **Governance > Classification**.
2. Select **Governance**.
3. Select **Data Discovery Assessment Report**.



### Result

BlueXP classification generates a PDF report that you can review and send to other groups as needed.

## Create the Data Mapping Overview Report

The Data Mapping Overview Report provides an overview of the data being stored in your corporate data sources to assist you with decisions of migration, back up, security, and compliance processes. The report summarizes all working environments and data sources. It also provides an analysis for each working environment.


The report includes the following information:




Category	Description
Usage Capacity	For all working environments: Lists the number of files and the used capacity for each working environment. For single working environments: Lists the files that are using the most capacity.
Age of Data	Provides three charts and graphs for when files were created, last modified, or last accessed. Lists the number of files, and their used capacity, based on certain date ranges.
Size of Data	Lists the number of files that exist within certain size ranges in your working environments.
File Types	Lists the total number of files and the used capacity for each type of file being stored in your working environments.

### Steps

1. From the BlueXP menu, select **Governance > Classification**.
2. Select **Governance**.
3. Select **Full Data Mapping Overview Report**.



4. To customize the company name that appears on the first page of the report, from the top right of the BlueXP classification page, select . Then select **Change company name**. The next time you generate the report, it will include the new name.

### Result

BlueXP classification generates a .pdf report that you can review and send to other groups as needed.

If the report is larger than 1 MB, the .pdf file is retained on the BlueXP classification instance and you'll see a pop-up message about the exact location. When BlueXP classification is installed on a Linux machine on your premises, or on a Linux machine you deployed in the cloud, you can navigate directly to the .pdf file. When BlueXP classification is deployed in the cloud, you'll need to SSH to the BlueXP classification instance to download the .pdf file.

### Review the top data repositories listed by data sensitivity

The *Top Data Repositories by Sensitivity Level* area of the Data Mapping Overview report lists the top four data repositories (working environments and data sources) that contain the most sensitive items. The bar chart for each working environment is divided into:

- Non-Sensitive data
- Personal data
- Sensitive Personal data

This data refreshes every two hours and can be manually refreshed.

### Steps

1. To see the total number of items in each category, position your cursor over each section of the bar.
2. To filter results that will appear in the Investigation page, select each area in the bar and investigate further.

### Review sensitive data and wide permissions

The *Sensitive Data and Wide Permissions* area of the Data Mapping Overview report shows the percentage of files that contain sensitive data and have wide permissions. The chart shows the following types of permissions:

- From the most restrictive permissions to the most permissive restrictions on the horizontal axis.
- From the least sensitive data to the most sensitive data on the vertical axis.

#### Steps

1. To see the total number of files in each category, position your cursor over each box.
2. To filter results that will appear in the Investigation page, select a box and investigate further.

### Review data listed by types of open permissions

The *Open Permissions* area of the Data Mapping Overview report shows the percentage for each type of permissions that exist for all files that are being scanned. The chart shows the following types of permissions:

- No Open Permissions
- Open to Organization
- Open to Public
- Unknown Access

#### Steps

1. To see the total number of files in each category, position your cursor over each box.
2. To filter results that will appear in the Investigation page, select a box and investigate further.

### Review the age and size of data

You might want to investigate the items in the *Age* and *Size* graphs of the Data Mapping Overview report to see if there is any data you should delete or tier to less expensive object storage.

#### Steps

1. In the Age of Data chart, to see details about the age of the data, position your cursor over a point in the chart.
2. To filter by an age or size range, select that age or size.
  - **Age of Data graph** - Categorizes data based on the time it was created, the last time it was accessed, or the last time it was modified.
  - **Size of Data graph** - Categorizes data based on size.



If any of your data sources implement data tiering, old data that already resides in object storage might be identified in the *Age of Data* graph.

## Review the most identified data classifications in your data

The *Classification* area of the Data Mapping Overview report provides a list of the most identified [Categories](#) and [File types](#) in your scanned data.

Categories can help you understand what's happening with your data by showing you the types of information that you have. For example, a category like "resumes" or "employee contracts" can include sensitive data. When you investigate the results, you might find that employee contracts are stored in a nonsecure location. You can then correct that issue.

See [Viewing files by categories](#) for more information.

### Steps

1. From the BlueXP menu, select **Governance > Classification**.
2. Select **Governance** then the **Data Discovery Assessment Report** button.

### Result

BlueXP classification generates a .pdf report that you can review and send to other groups as needed.


## View compliance details about the private data stored in your organization with BlueXP classification

Gain control of your private data by viewing details about the personal data (PII) and sensitive personal data (SPII) in your organization. You can also gain visibility by reviewing the categories and file types that BlueXP classification found in your data.



The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.

By default, the BlueXP classification dashboard displays compliance data for all working environments and databases. To see data for only some of the working environments, select them.



Filter the results from the Data Investigation page and download a report of the results as a CSV file. See [Filtering data in the Data Investigation page](#) for details.

## View files that contain personal data


BlueXP classification automatically identifies specific words, strings, and patterns (Regex) inside the data. For example, Personal Identification Information (PII), credit card numbers, social security numbers, bank account numbers, passwords, and more. [See the full list](#). BlueXP classification identifies this type of information in individual files, in files within directories (shares and folders), and in database tables.

You can also create custom search terms to identify personal data specific to your organization. For more information, see [Create a custom classification](#).

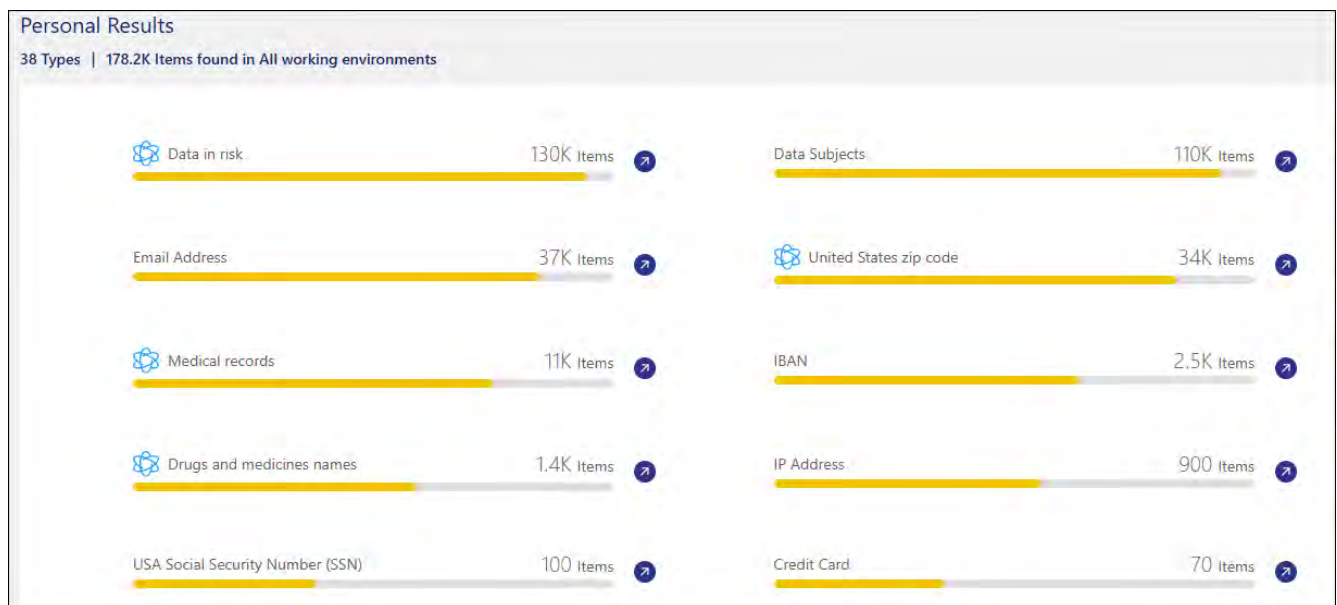
For some types of personal data, BlueXP classification uses *proximity validation* to validate its findings. The validation occurs by looking for one or more predefined keywords in proximity to the personal data that was found. For example, BlueXP classification identifies a U.S. social security number (SSN) as an SSN if it sees a proximity word next to it—for example, *SSN* or *social security*. [The table of personal data](#) shows when BlueXP classification uses proximity validation.

### Steps

1. From the BlueXP classification menu, select the **Compliance** tab.
2. To investigate the details for all personal data, select the icon next to the personal data percentage.



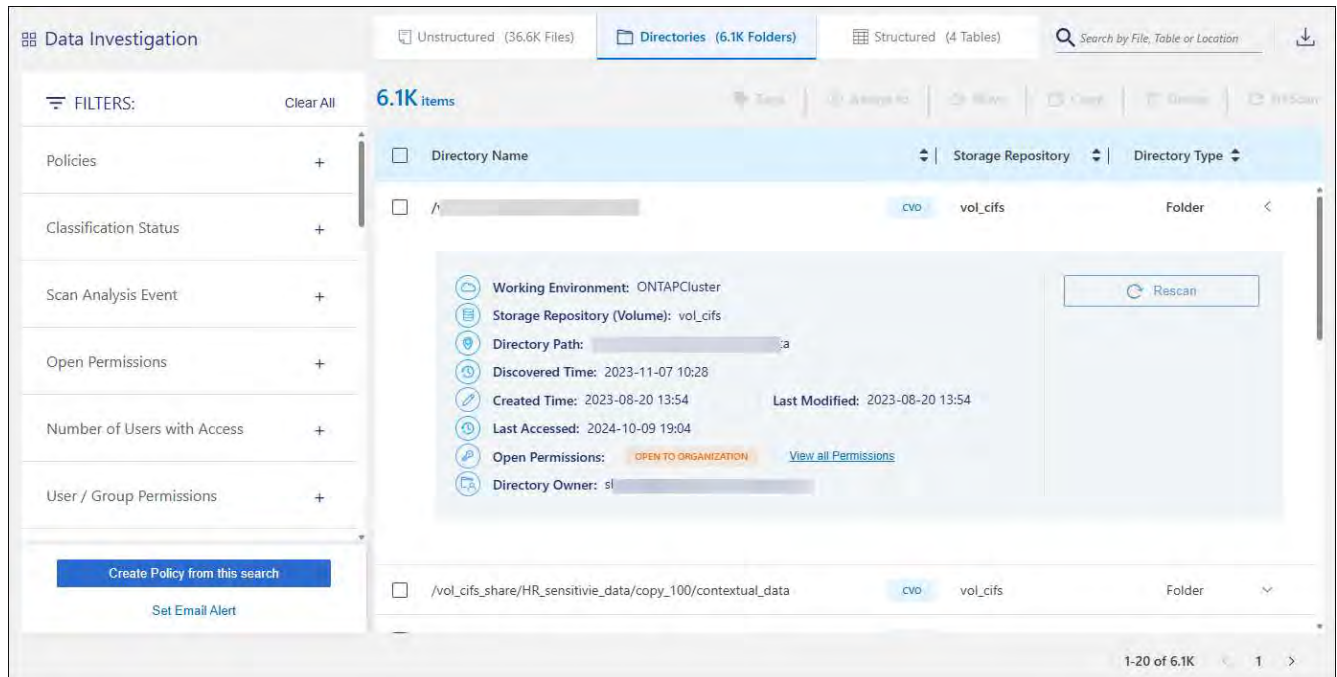
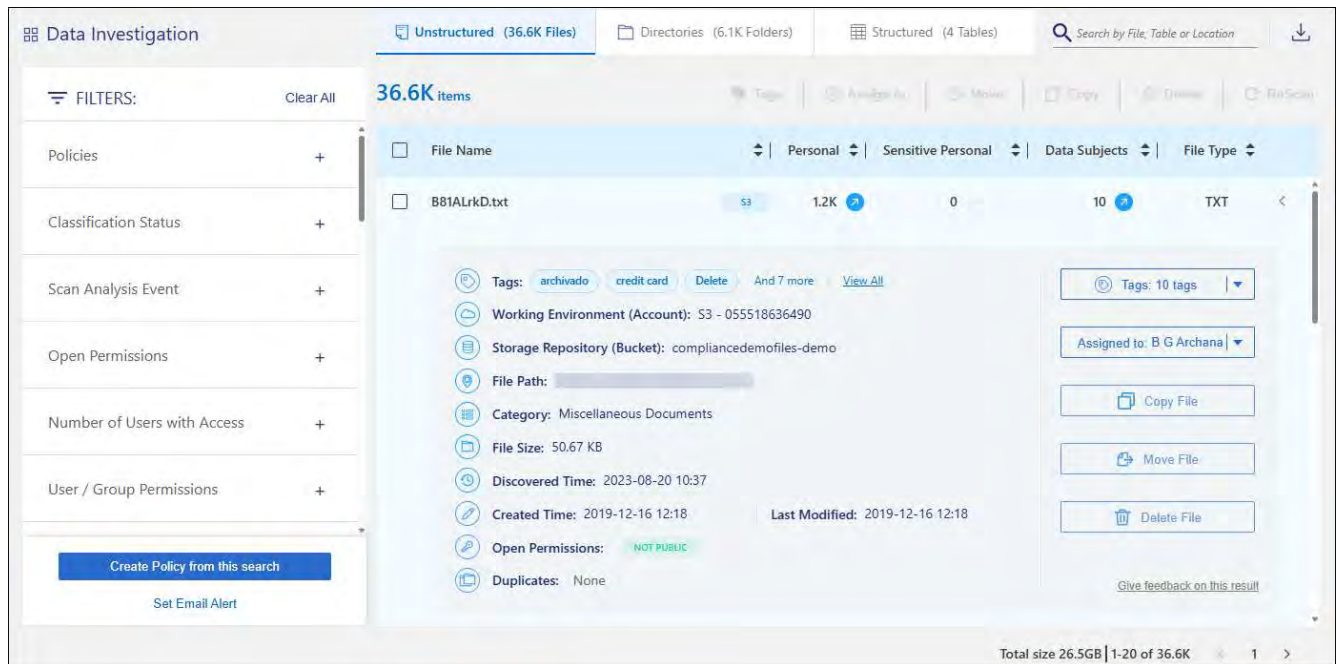
- To investigate the details for a specific type of personal data, select **View All** and then select the **Investigate Results** arrow icon for a specific type of personal data, for example, email addresses.



- Investigate the data by searching, sorting, expanding details for a specific file, selecting the **Investigate Results** arrow to see masked information, or by downloading the file list.

The two screenshots below show personal data found in individual files, and found in files within directories (shares and folders). You can also select the **Structured** tab to view personal data found in databases.





## View files that contain sensitive personal data

BlueXP classification automatically identifies special types of sensitive personal information, as defined by privacy regulations such as [articles 9 and 10 of the GDPR](#). For example, information regarding a person's health, ethnic origin, or sexual orientation. [See the full list](#). BlueXP classification identifies this type of information in individual files, in files within directories (shares and folders), and in database tables.

BlueXP classification uses artificial intelligence (AI), natural language processing (NLP), machine learning (ML), and cognitive computing (CC) to understand the meaning of the content that it scans in order to extract entities and categorize it accordingly.

For example, one sensitive GDPR data category is ethnic origin. Because of its NLP abilities, BlueXP classification can distinguish the difference between a sentence that reads "George is Mexican" (indicating

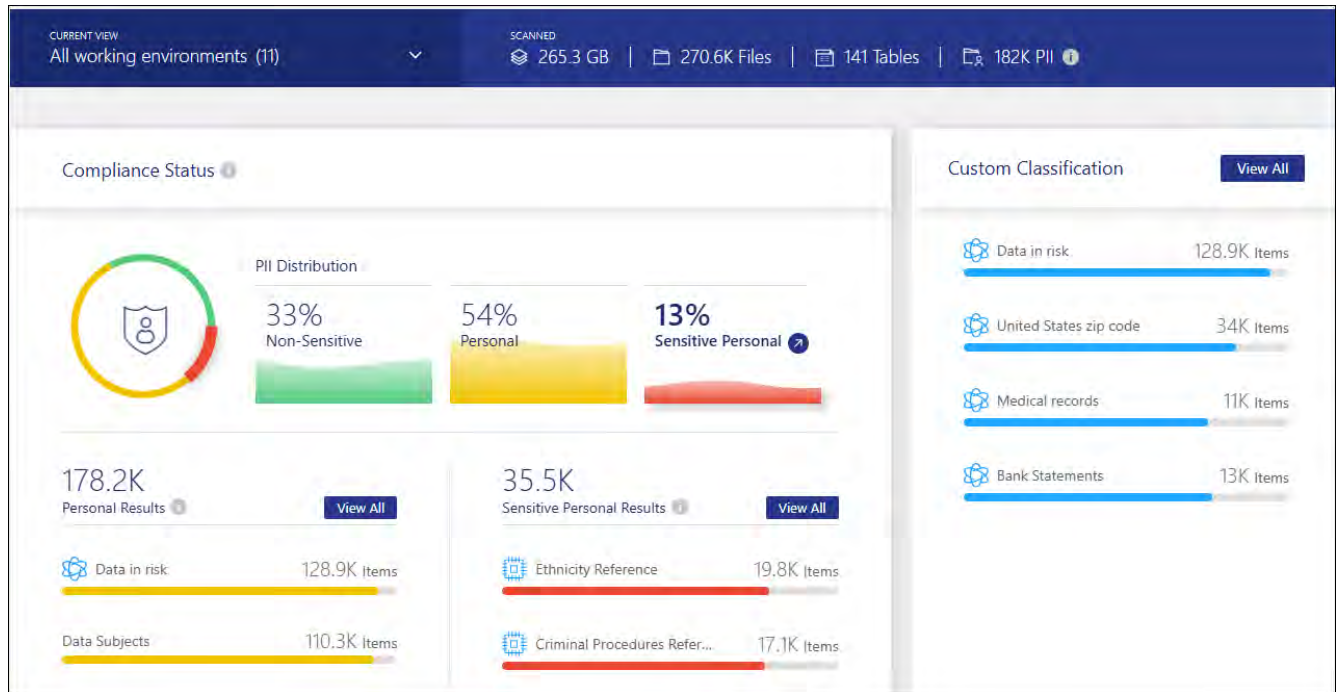
sensitive data as specified in article 9 of the GDPR), versus "George is eating Mexican food."



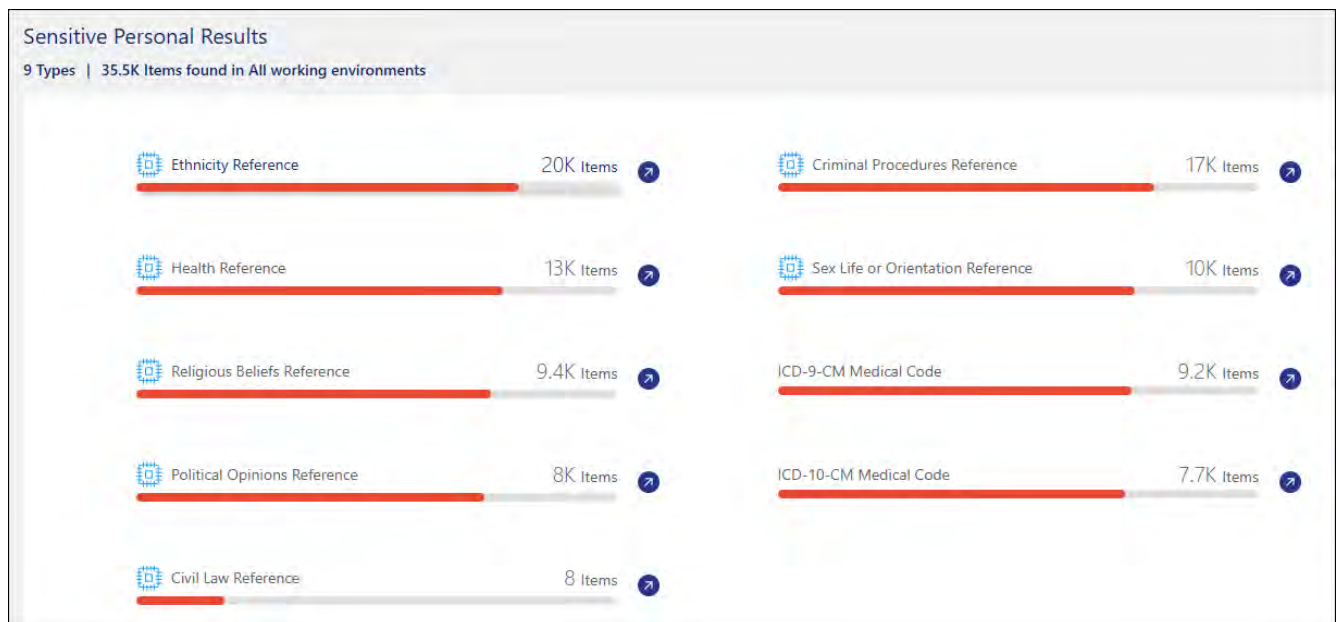
Only English is supported when scanning for sensitive personal data. Support for more languages will be added later.

### Steps

1. From the BlueXP classification menu, select **Compliance**.
2. To investigate the details for all sensitive personal data, select the icon next to the sensitive personal data percentage.



3. To investigate the details for a specific type of sensitive personal data, select **View All** and then select the **Investigate Results** arrow icon for a specific type of sensitive personal data.



- Investigate the data by searching, sorting, expanding details for a specific file, clicking **Investigate Results** to see masked information, or by downloading the file list.

## View files by categories

BlueXP classification takes the data that it scanned and divides it into different types of categories. Categories are topics based on AI analysis of the content and metadata of each file. [See the list of categories.](#)

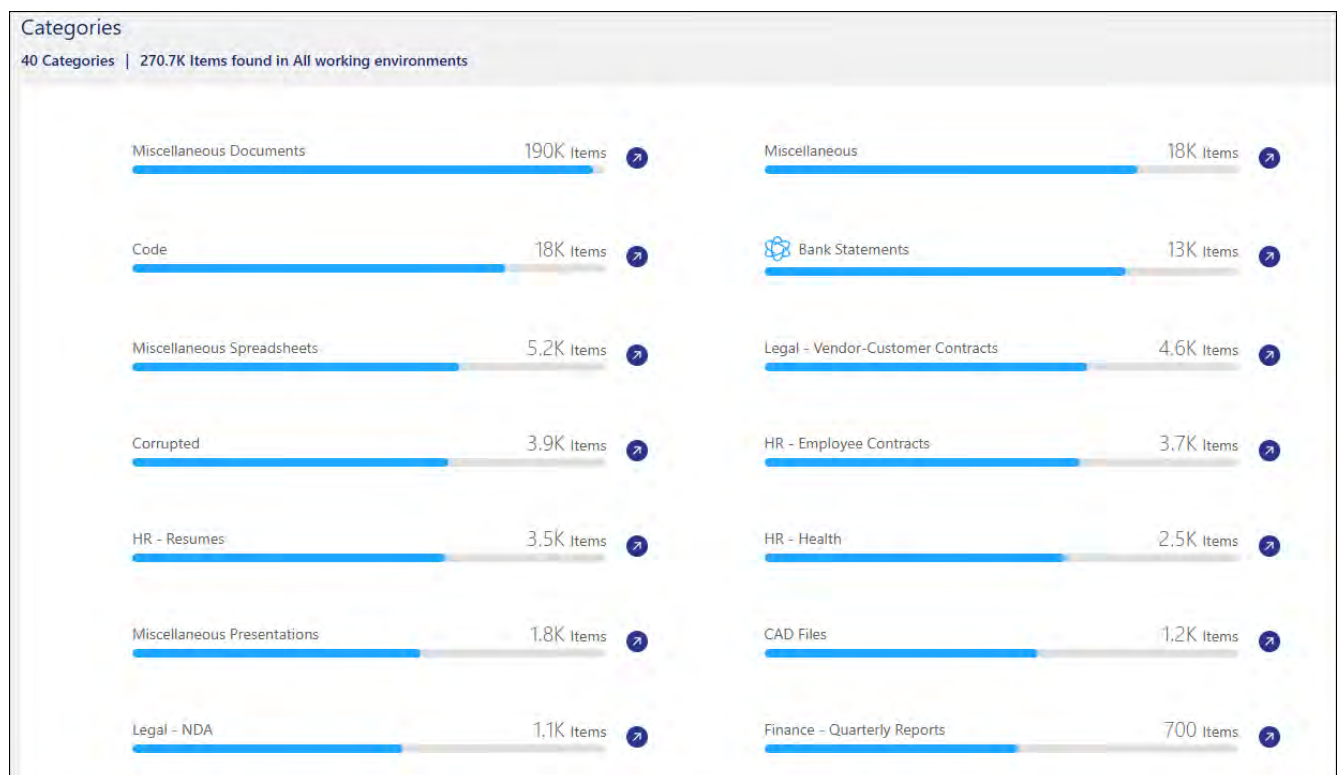
Categories can help you understand what's happening with your data by showing you the types of information that you have. For example, a category like resumes or employee contracts can include sensitive data. When you investigate the results, you might find that employee contracts are stored in a nonsecure location. You can then correct that issue.



English, German, and Spanish are supported for categories. Support for more languages will be added later.

### Steps

- From the BlueXP classification menu, select the **Compliance** tab.
- Select the **Investigate Results** arrow icon for one of the top 4 categories directly from the main screen, or select **View All** and then select the icon for any of the categories.



- Investigate the data by searching, sorting, expanding details for a specific file, selecting the **Investigate Results** arrow to see masked information, or by downloading the file list.

## View files by file types


BlueXP classification takes the data that it scanned and breaks it down by file type. Reviewing your file types can help you control your sensitive data because you might find that certain file types are not stored correctly. [See the list of file types.](#)



For example, you might be storing CAD files that include very sensitive information about your organization. If they are unsecured, you can take control of the sensitive data by restricting permissions or moving the files to another location.

## Steps

1. From the BlueXP classification menu, select the **Compliance** tab.
2. Select the **Investigate Results** arrow icon for one of the top 4 file types directly from the main screen, or select **View All** and then select the icon for any of the file types.



3. Investigate the data by searching, sorting, expanding details for a specific file, selecting the **Investigate Results** arrow to see masked information, or by downloading the file list.

## Categories of private data in BlueXP classification

There are many types of private data that BlueXP classification can identify in your volumes and databases.

BlueXP classification identifies two types of personal data:

- **Personally identifiable information (PII)**
- **Sensitive personal information (SPII)**



If you need BlueXP classification to identify other private data types, such as additional national ID numbers or healthcare identifiers, contact your account manager.

## Types of personal data

The personal data, or *personally identifiable information* (PII), found in files can be general personal data or

national identifiers. The third column in the table below identifies whether BlueXP classification uses [proximity validation](#) to validate its findings for the identifier.

The languages in which these items can be recognized are identified in the table.

Type	Identifier	Proximity validation?	English	German	Spanish	French	Japanese
General	Credit card number	Yes	✓	✓	✓		✓
	Data Subjects	No	✓	✓	✓		
	Email Address	No	✓	✓	✓		✓
	IBAN Number (International Bank Account Number)	No	✓	✓	✓		✓
	IP Address	No	✓	✓	✓		✓
	Password	Yes	✓	✓	✓		✓

Type	Identifier	Proximity validation?	English	German	Spanish	French	Japanese
------	------------	-----------------------	---------	--------	---------	--------	----------

National Identifiers

--

Type	Identifier	Proximity validation?	English	German	Spanish	French	Japanese
------	------------	-----------------------	---------	--------	---------	--------	----------

Type	Identifier	Proximity validation?	English	German	Spanish	French	Japanese
------	------------	-----------------------	---------	--------	---------	--------	----------

Type	Identifier	Proximity	English	German	Spanish	French	Japanese
	Corporate)						
	Latvian ID	Yes	✓	✓	✓		
	Lithuanian ID	Yes	✓	✓	✓		
	Luxembourg ID	Yes	✓	✓	✓		
	Maltese ID	Yes	✓	✓	✓		
	National Health Service (NHS) Number	Yes	✓	✓	✓		
	New Zealand Bank Account	Yes	✓	✓	✓		
	New Zealand Driver's License	Yes	✓	✓	✓		
	New Zealand IRD Number (Tax ID)	Yes	✓	✓	✓		
	New Zealand NHI (National Health Index) Number	Yes	✓	✓	✓		
	New Zealand Passport Number	Yes	✓	✓	✓		
	Polish ID (PESEL)	Yes	✓	✓	✓		
	Portuguese Tax Identification Number (NIF)	Yes	✓	✓	✓		
	Romanian ID (CNP)	Yes	✓	✓	✓		
	Singapore National Registration Identity Card (NRIC)	Yes	✓	✓	✓		
	Slovenian ID (EMSO)	Yes	✓	✓	✓		
	South African ID	Yes	✓	✓	✓		
	Spanish Tax Identification Number	Yes	✓	✓	✓		
	Swedish ID	Yes	✓	✓	✓		
	UK ID (NINO)	Yes	✓	✓	✓		
	USA California Driver's License	Yes	✓	✓	✓		
	USA Indiana Driver's License	Yes	✓	✓	✓		
	USA New York Driver's License	Yes	✓	✓	✓		
	USA Texas Driver's License	Yes	✓	✓	✓		
	USA Social Security Number (SSN)	Yes	✓	✓	✓		

## Types of sensitive personal data

BlueXP classification can find the following sensitive personal information (SPII) in files.

The items in this category can be recognized only in English at this time.

- **Criminal Procedures Reference:** Data concerning a natural person's criminal convictions and offenses.
- **Ethnicity Reference:** Data concerning a natural person's racial or ethnic origin.
- **Health Reference:** Data concerning a natural person's health.
- **ICD-9-CM Medical Codes:** Codes used in the medical and health industry.
- **ICD-10-CM Medical Codes:** Codes used in the medical and health industry.

- **Philosophical Beliefs Reference:** Data concerning a natural person’s philosophical beliefs.
- **Political Opinions Reference:** Data concerning a natural person’s political opinions.
- **Religious Beliefs Reference:** Data concerning a natural person’s religious beliefs.
- **Sex Life or Orientation Reference:** Data concerning a natural person’s sex life or sexual orientation.

## Types of categories

BlueXP classification categorizes your data as follows.

Most of these categories can be recognized in English, German, and Spanish.

Category	Type	English	German	Spanish
Finance	Balance Sheets	✓	✓	✓
	Purchase Orders	✓	✓	✓
	Invoices	✓	✓	✓
	Quarterly Reports	✓	✓	✓
HR	Background Checks	✓		✓
	Compensation Plans	✓	✓	✓
	Employee Contracts	✓		✓
	Employee Reviews	✓		✓
	Health	✓		✓
	Resumes	✓	✓	✓
Legal	NDA's	✓	✓	✓
	Vendor-Customer contracts	✓	✓	✓
Marketing	Campaigns	✓	✓	✓
	Conferences	✓	✓	✓
Operations	Audit Reports	✓	✓	✓
Sales	Sales Orders	✓	✓	
Services	RFI	✓		✓
	RFP	✓		✓
	SOW	✓	✓	✓
	Training	✓	✓	✓
Support	Complaints and Tickets	✓	✓	✓

The following metadata is also categorized and identified in the same supported languages:

- Application Data
- Archive Files

- Audio
- Breadcrumbs from BlueXP classification  
Business Application Data
- CAD Files
- Code
- Corrupted
- Database and index files
- Design Files
- Email Application Data
- Encrypted (files with a high entropy score)
- Executables
- Financial Application Data
- Health Application Data
- Images
- Logs
- Miscellaneous Documents
- Miscellaneous Presentations
- Miscellaneous Spreadsheets
- Miscellaneous "Unknown"
- Password Protected files
- Structured Data
- Videos
- Zero-Byte Files

## Types of files

BlueXP classification scans all files for category and metadata insights and displays all file types in the file types section of the dashboard. When BlueXP classification detects Personal Identifiable Information (PII) or when it performs a DSAR search, only the following file formats are supported:

.CSV, .DCM, .DOC, .DOCX, .JSON, .PDF, .PPTX, .RTF, .TXT, .XLS, .XLSX, Docs, Sheets, and Slides

## Accuracy of information found

NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that BlueXP classification identifies. You should always validate the information by reviewing the data.

Based on our testing, the table below shows the accuracy of the information that BlueXP classification finds. We break it down by *precision* and *recall*:

### Precision

The probability that what BlueXP classification finds has been identified correctly. For example, a precision rate of 90% for personal data means that 9 out of 10 files identified as containing personal information,



actually contain personal information. 1 out of 10 files would be a false positive.

## Recall

The probability for BlueXP classification to find what it should. For example, a recall rate of 70% for personal data means that BlueXP classification can identify 7 out of 10 files that actually contain personal information in your organization. BlueXP classification would miss 30% of the data and it won't appear in the dashboard.

We are constantly improving the accuracy of our results. Those improvements will be automatically available in future BlueXP classification releases.

Type	Precision	Recall
Personal data - General	90%-95%	60%-80%
Personal data - Country identifiers	30%-60%	40%-60%
Sensitive personal data	80%-95%	20%-30%
Categories	90%-97%	60%-80%

## Create a custom classification in BlueXP classification


With BlueXP classification, you can create a custom search for sensitive information. The search can be scoped to a regular expression (regex).

### Create a custom classification


Custom classification is only available for Map & Classify scans, not mapping-only scans. This feature is currently in preview.

#### Steps

1. Select the **Custom classification** tab.



2. Select the **Add New Classifier** button.
3. Add a Name and Description for the new classifier.
4. To add the customization as a regular expression, select **Custom regular expression** then **Next**.
5. Add a pattern to detect the specific information of your data. Select **Validate** to confirm the syntax of your entry.



6. Select **Done** to create the custom classification.

The new customization is captured in the next scheduled scan. To view results, see [Generate compliance reports](#).

# Investigate the data stored in your organization with BlueXP classification

Investigate the data from your organization by viewing details in the Data Investigation page. Here is where you can continue your research after looking at the Governance dashboard. On the Investigation page, you can filter the data using one of the many filters to show only the results you want to see. You can also view file metadata, permissions for files and directories, and check for duplicate files in your storage systems.

You can navigate to this page from many areas of the BlueXP classification UI, including the Governance and Compliance dashboards with the filters selected already on those pages. You can export the data into a CSV or JSON file for further analysis or to share with others.




The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.

## Filter data in the Data Investigation page

You can filter the contents of the investigation page to display only the results you want to see.

### Steps

1. From the BlueXP classification menu, select **Investigation**.
2. On the Data Investigation page, do any of the following:
3. To download the contents of the page as a report after you've refined it, select the button.



4. To view the data from files (unstructured data), directories (folders and file shares), or from databases (structured data), select one of the tabs at the top.

5. To sort the results in numerical or alphabetical order, select the control at the top of each column.
6. To refine the results even more, select one of the filters in the Filter pane.



You can only view the first 10,000 results—or 500 pages—for a scan on the Data Investigation page.

### Filter data by sensitivity and content

Use the following filters to view how much sensitive information is contained in your data.

Filter	Details
Category	Select the <a href="#">types of categories</a> .
Sensitivity Level	Select the sensitivity level: Personal, Sensitive personal, or Non sensitive.
Number of identifiers	Select the range of detected sensitive identifiers per file. Includes personal data and sensitive personal data. When filtering in Directories, BlueXP classification totals the matches from all files in each folder (and sub-folders).  NOTE: The December 2023 (version 1.26.6) release removed the option to calculate the number of personal identifiable information (PII) data by Directories.
Personal Data	Select the <a href="#">types of personal data</a> .
Sensitive Personal Data	Select the <a href="#">types of sensitive personal data</a> .
Data Subject	Enter a data subject's full name or known identifier. <a href="#">Learn more about data subjects here</a> .

### Filter data by user owner and user permissions

Use the following filters to view file owners and permissions to access your data.

Filter	Details
Open Permissions	Select the type of permissions within the data and within folders/shares.
User / Group Permissions	Select one or multiple user names and/or group names, or enter a partial name.
File Owner	Enter the file owner name.
Number of users with access	Select one or multiple category ranges to show which files and folders are open to a certain number of users.

### Filter data by time

Use the following filters to view data based on time criteria.

Filter	Details
Created Time	Select a time range when the file was created. You can also specify a custom time range to further refine the search results.

Filter	Details
Discovered Time	Select a time range when BlueXP classification discovered the file. You can also specify a custom time range to further refine the search results.
Last Modified	Select a time range when the file was last modified. You can also specify a custom time range to further refine the search results.
Last Accessed	Select a time range when the file, or directory (CIFS or NFS only), was last accessed. You can also specify a custom time range to further refine the search results. For the types of files that BlueXP classification scans, this is the last time BlueXP classification scanned the file.  BlueXP classification does not extract the "last accessed time" from the following data sources: SharePoint Online, SharePoint On-premises (SharePoint Server), OneDrive, Google Drive, and Amazon S3.

### Filter data by metadata

Use the following filters to view data based on location, size, and directory or file type.

Filter	Details
File Path	Enter up to 20 partial or full paths that you want to include or exclude from the query. If you enter both include paths and exclude paths, BlueXP classification finds all files in the included paths first, then it removes files from excluded paths, and then it displays the results. Note that using "*" in this filter has no effect, and that you can't exclude specific folders from the scan - all the directories and files under a configured share will be scanned.
Directory Type	Select the directory type; either "Share" or "Folder".
File Type	Select the <a href="#">types of files</a> .
File Size	Select the file size range.
File Hash	Enter the file's hash to find a specific file, even if the name is different.

### Filter data by storage type

Use the following filters to view data by storage type.

Filter	Details
Working Environment Type	Select the type of working environment. OneDrive, SharePoint, and Google Drive are categorized under "Apps".
Working Environment name	Select specific working environments.
Storage Repository	Select the storage repository, for example, a volume or a schema.

### Filter data by saved searches

Use the following filter to view data by saved searches.

Filter	Details
Saved search	Select one saved search or multiples. Go to the <a href="#">saved searches tab</a> to view the list of existing saved searches and create new ones.

### Filter data by analysis status

Use the following filter to view data by the BlueXP classification scan status.

Filter	Details
Analysis Status	Select an option to show the list of files that are Pending First Scan, Completed being scanned, Pending Rescan, or that have Failed to be scanned.
Scan Analysis Event	Select whether you want to view files that were not classified because BlueXP classification couldn't revert last accessed time, or files that were classified even though BlueXP classification couldn't revert last accessed time.

See [details about the "last accessed time" timestamp](#) for more information about the items that appear in the Investigation page when filtering using the Scan Analysis Event.

### Filter data by duplicates

Use the following filter to view files that are duplicated in your storage.


Filter	Details
Duplicates	Select whether the file is duplicated in the repositories.

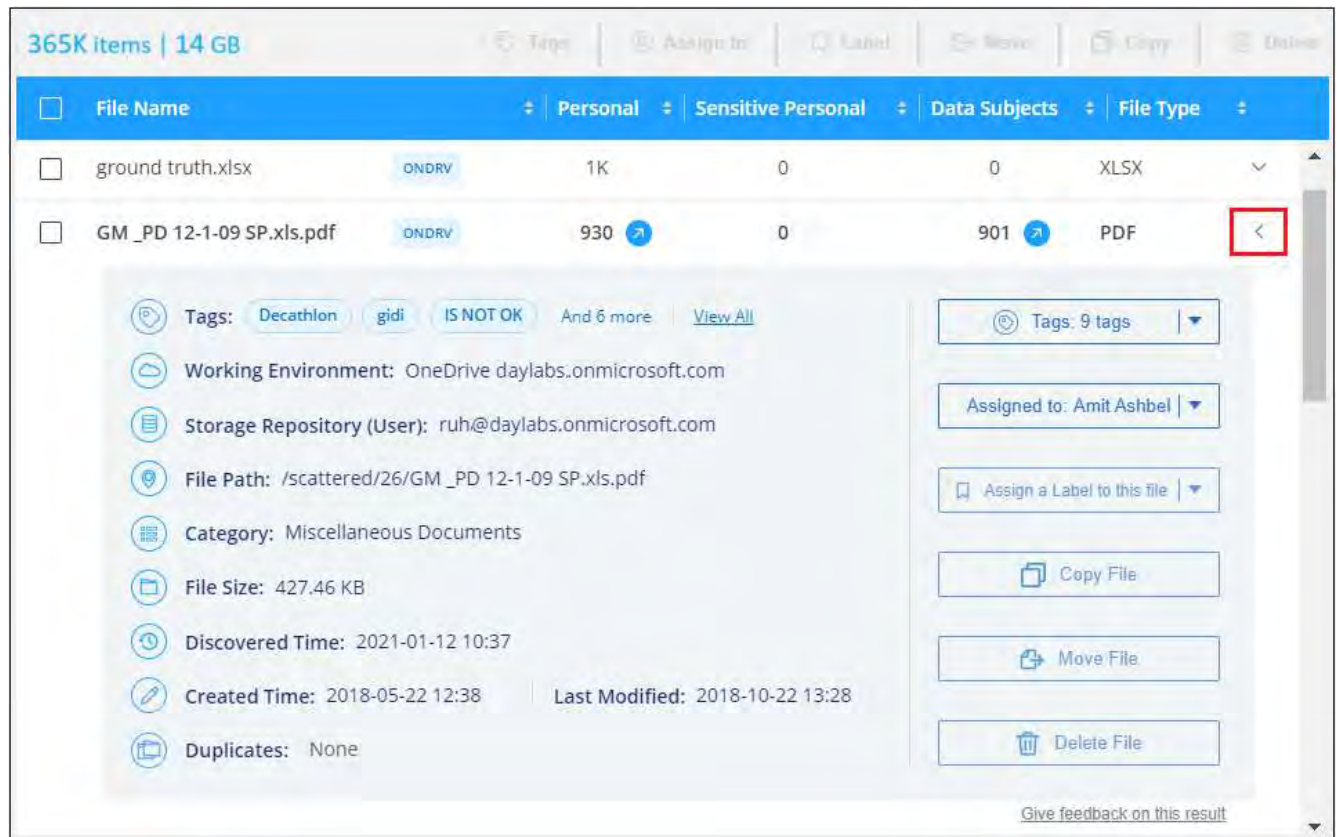
### View file metadata

In addition to showing you the working environment and volume where the file resides, the metadata shows much more information, including the file permissions, file owner, and whether there are duplicates of this file. This information is useful if you're planning to [create saved searches](#) because you can see all the information that you can use to filter your data.

The availability of information depends on the data source. For example, volume name and permissions are not shared for database files.

#### Steps

1. From the BlueXP classification menu, select **Investigation**.
2. In the Data Investigation list on the right, select the down-caret  on the right for any single file to view the file metadata.




## View users' permissions for files and directories

To view a list of all users or groups who have access to a file or to a directory and the types of permissions they have, select **View all Permissions**. This button is available only for data in CIFS shares.

Note that if you see SIDs (Security IDentifiers) instead of user and group names, you should integrate your Active Directory into BlueXP classification. [See how to do this](#).

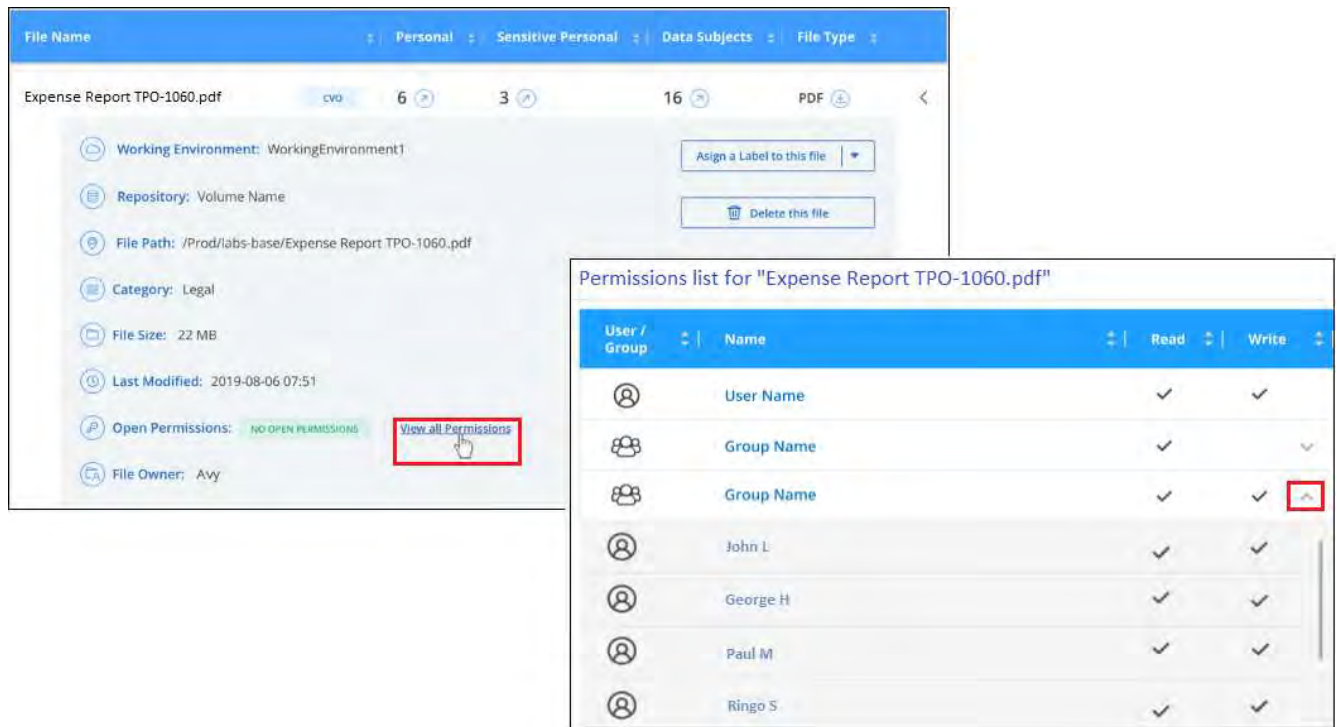
### Steps

1. From the BlueXP classification menu, select **Investigation**.
2. In the Data Investigation list on the right, select the down-caret  on the right for any single file to view the file metadata.
3. To view a list of all users or groups who have access to a file or to a directory and the types of permissions they have, in the Open Permissions field, select **View all Permissions**.



BlueXP classification shows up to 100 users in the list.





4. Select the down-caret ▼ button for any group to see the list of users who are part of the group.



You can expand one level of the group to see the users who are part of the group.

5. Select the name of a user or group to refresh the Investigation page so you can see all the files and directories that the user or group has access to.

## Check for duplicate files in your storage systems

You can check whether duplicate files are being stored in your storage systems. This is useful if you want to identify areas where you can save storage space. It's also good to ensure certain files that have specific permissions or sensitive information are not unnecessarily duplicated in your storage systems.

All of your files (not including databases) that are 1 MB or larger, or that contain personal or sensitive personal information, are compared to see if there are duplicates.

BlueXP classification uses hashing technology to determine duplicate files. If any file has the same hash code as another file, you can be 100% sure that the files are exact duplicates—even if the file names are different.

### Steps


1. From the BlueXP classification menu, select **Investigation**.
2. In the Investigation page Filters pane on the left, select "File Size" along with "Duplicates" ("Has duplicates") to see which files of a certain size range are duplicated in your environment.
3. Optionally, download the list of duplicate files and send it to your storage admin so they can decide which files, if any, can be deleted.
4. Optionally, [delete the file](#) yourself if you are confident that a specific version of the file is not needed.

### View if a specific file is duplicated

You can see if a single file has duplicates.

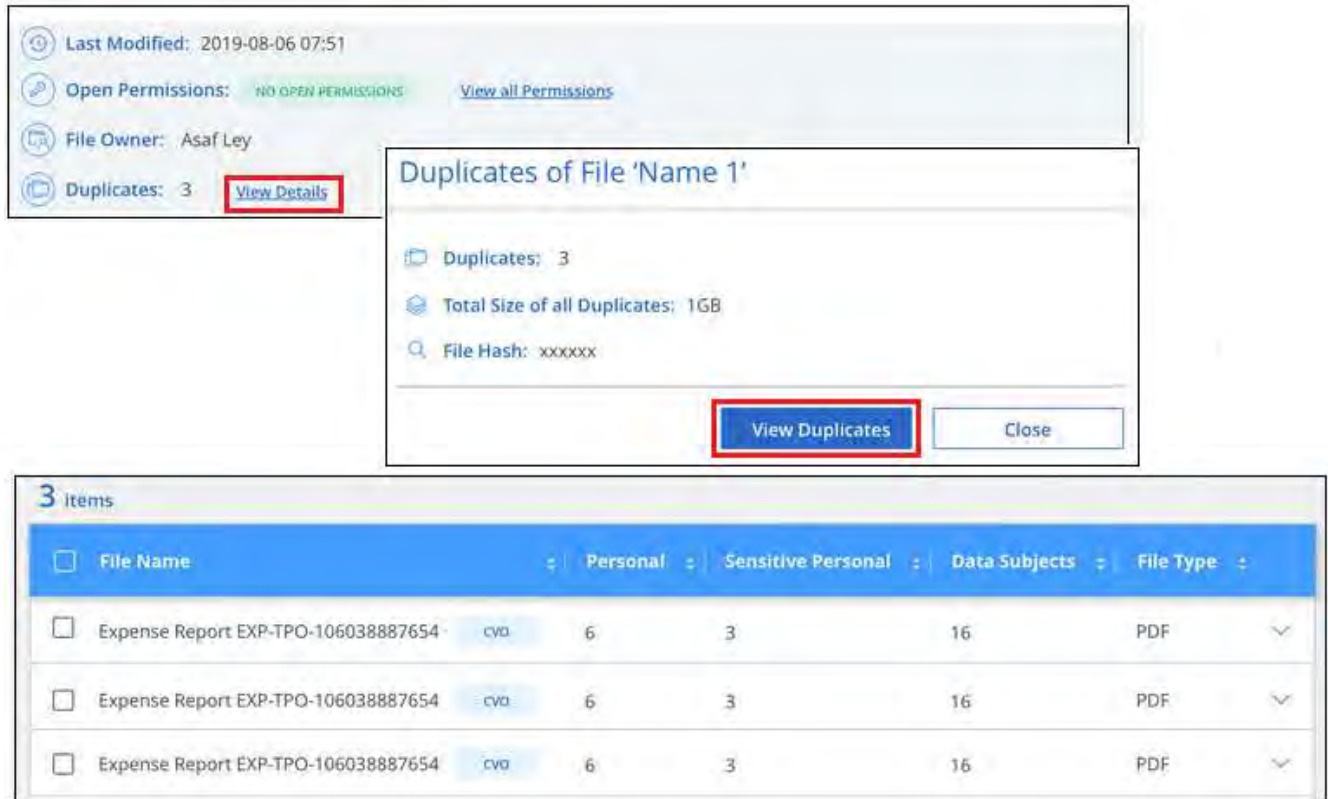


## Steps

1. From the BlueXP classification menu, select **Investigation**.
2. In the Data Investigation list, select  on the right for any single file to view the file metadata.

If duplicates exist for a file, this information appears next to the *Duplicates* field.

3. To view the list of duplicate files and where they are located, select **View Details**.
4. In the next page select **View Duplicates** to view the files in the Investigation page.



The screenshot shows the BlueXP interface. At the top, file metadata is displayed: Last Modified: 2019-08-06 07:51, Open Permissions: NO OPEN PERMISSIONS (with a link to View all Permissions), File Owner: Asaf Ley, and Duplicates: 3 (with a red box around the 'View Details' button). Below this, a modal window titled 'Duplicates of File 'Name 1'' is open, showing Duplicates: 3, Total Size of all Duplicates: 1GB, and File Hash: xxxxxx (with a red box around the 'View Duplicates' button). At the bottom, a table titled '3 Items' lists three duplicate files:

<input type="checkbox"/>	File Name		Personal	Sensitive Personal	Data Subjects	File Type
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cva	6	3	16	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cva	6	3	16	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cva	6	3	16	PDF



You can use the "file hash" value provided in this page and enter it directly in the Investigation page to search for a specific duplicate file at any time - or you can use it in a saved search.

## Create the Data Investigation Report

The Data Investigation Report is a download of the filtered contents of the Data Investigation page.

The report is available as a CSV or JSON file you can save to your local machine.

There can be up to three report files downloaded if BlueXP classification is scanning files (unstructured data), directories (folders and file shares), and databases (structured data).

The files are split into files with a fixed number of rows or records:

- JSON - 100,000 records per report that takes about 5 minutes to generate
- CSV - 200,000 records per report that takes about 4 minutes to generate



You can download a version of the CSV file to view in this browser. This version is limited to 10,000 records.

## What's included in the Data Investigation Report

The **Unstructured Files Data Report** includes the following information about your files:

- File name
- Location type
- Working environment name
- Storage repository (for example, a volume, bucket, shares)
- Repository type
- File path
- File type
- File size (in MB)
- Created time
- Last modified
- Last accessed
- File owner
  - File owner data encompasses account name, SAM account name, and e-mail address when Active Directory is configured.
- Category
- Personal information
- Sensitive personal information
- Open permissions
- Scan Analysis Error
- Deletion detection date

The deletion detection date identifies the date that the file was deleted or moved. This enables you to identify when sensitive files have been moved. Deleted files don't contribute to the file number count that appears in the dashboard or on the Investigation page. The files only appear in the CSV reports.

The **Unstructured Directories Data Report** includes the following information about your folders and file shares:


- Working environment type
- Working environment name
- Directory name
- Storage repository (for example, a folder or file shares)
- Directory owner
- Created time
- Discovered time

- Last modified
- Last accessed
- Open permissions
- Directory type

The **Structured Data Report** includes the following information about your database tables:

- DB Table name
- Location type
- Working environment name
- Storage repository (for example, a schema)
- Column count
- Row count
- Personal information
- Sensitive personal information

#### **Steps to generate the report**

1. From the Data Investigation page, select the  button on the top, right of the page.
2. Choose the report type: CSV or JSON.
3. Enter a **Report name**.
4. To download the complete report, select **Working environment** then choose the **Working Environment** and **Volume** from the respective dropdown menus. Provide a **Destination folder path**.

To download the report in the browser, select **Local** . Note this option limits the report to the first 10,000 rows and is limited to the **CSV** format. You don't need to complete any other fields if you select **Local**.

5. Select **Download Report**.

## Download Investigation Report

### Report type

CSV file       JSON file

### Report name

investigation\_report

### Export destination

Working environment       Local (limited to 10K rows)

Working environment 

Volume: 

Destination folder path

./folder/subfolder

Download Report

Cancel

### Result

A dialog displays a message that the reports are being downloaded.

## Create a saved search based on selected filters

You can create a saved search for frequently used search filters in the Data Investigation page to easily replicate those search queries.

### Steps

1. From the BlueXP classification menu, select **Investigation**.
2. On the Data Investigation page, select the filters you want to use to create a saved search.
3. At the bottom of the Filter pane, select **Create saved search from this search**.
4. Enter a name and a description for the saved search.
5. Choose any of the following:
6. Select **Create Saved Search**.



It might take up to 15 minutes for the results to appear on the Saved Searches page.

# Manage saved searches with BlueXP classification


BlueXP classification supports saving your search queries. With a saved search, you can create custom filters to sort through frequent queries of your data Investigation page. BlueXP classification also includes predefined saved searches based on common requests.



In versions of BlueXP classification earlier than 1.43, saved searches were called [policies](#).


The **Saved searches** tab in the Compliance Dashboard lists all the predefined and custom saved searches available on this instance of BlueXP classification.

Saved searches also appear in the list of filters in the Investigation page.



## View saved searches results in the Investigation page

To display the results for a saved search in the Investigation page, select the  button for a specific search then select **Investigate Results**.



## Create custom saved searches

You can create your own custom saved searches that provide results for queries specific to your organization. Results are returned for all files and directories (shares and folders) that match the search criteria.

### Steps

1. In the Investigation tab, define a search by selecting the filters you want to use. See [Filtering data in the Investigation page](#) for details.
2. Once you have all the filter characteristics set to your liking, select **Create saved search**.

## ☰ Data Investigation

☰ FILTERS: Clear All

---

Storage Repository 3 +

---

File / Directory Path +

---

Category +

---

Sensitivity Level +

---

Save this search

3. Name the saved search and add a description. The name must be unique.
4. Select **Create Saved Search**.

## Create search

---

This will save the current selected filters and search term as a saved search. You can view or delete this later from the “Saved searches” tab.

Note it may take up to 15 minutes for results to be displayed for a new saved search.

Name this search

Give it a detailed description that explains what it searches for

---

Create search

Cancel

Once you’ve created the search, you can view it in the **Saved searches** tab.

### Edit saved searches

You can modify the query criteria for a saved search (that is, the defined filters) to add or remove certain parameters.

You cannot modify default saved searches.

#### Steps

1. From the Saved searches page, select **Edit Search** for the search that you want to change.

Sensitive data

Last modified: 2024-03-04

Edit search

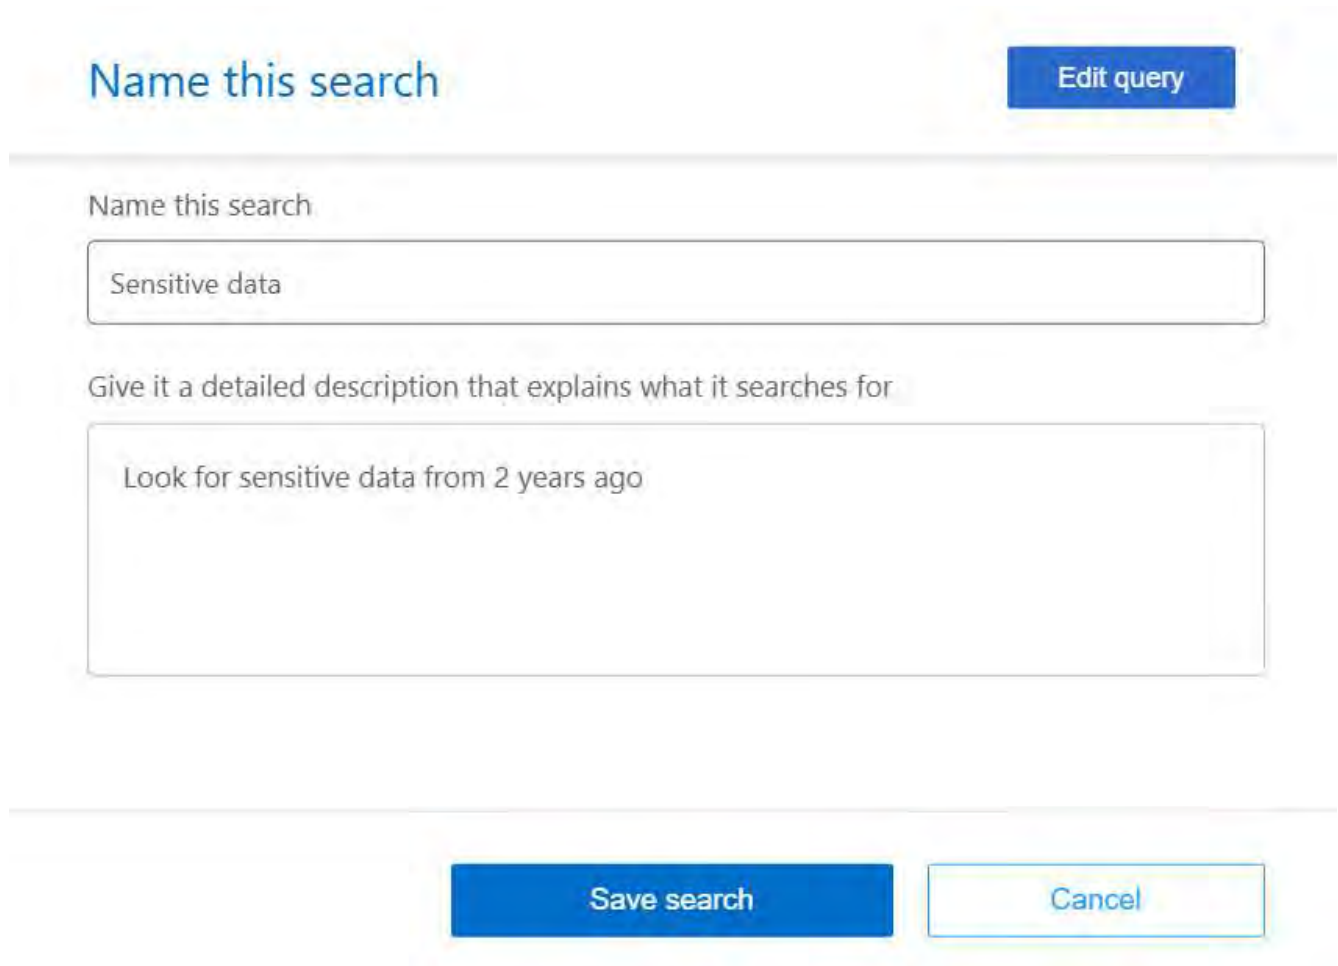


---

Look for sensitive data from 2 years ago

2. Make the changes to the name and description fields. To only change the name and description fields, select **Save search**.

To change the filters for the saved search, select **Edit query**.



Name this search

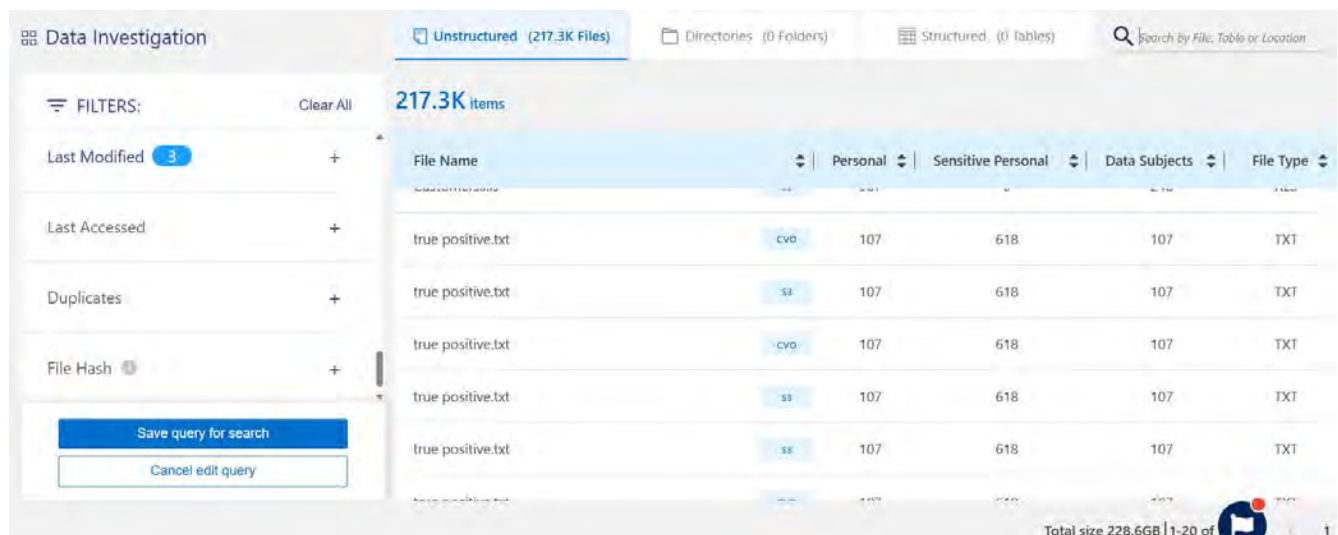
Sensitive data

Give it a detailed description that explains what it searches for

Look for sensitive data from 2 years ago

Save search Cancel

3. In the Investigation page, edit the query. You can add, remove, or modify filters. To complete your changes, select **Save query for this search**.



Data Investigation

Unstructured (217.3K Files) Directories (0 Folders) Structured (0 Tables) Search by File, Table or Location

FILTERS: Clear All 217.3K items

File Name	Personal	Sensitive Personal	Data Subjects	File Type	
true positive.txt	cvo	107	618	107	TXT
true positive.txt	ss	107	618	107	TXT
true positive.txt	cvo	107	618	107	TXT
true positive.txt	ss	107	618	107	TXT
true positive.txt	ss	107	618	107	TXT
true positive.txt	ss	107	618	107	TXT


Save query for search Cancel edit query

Total size 228.6GB | 1-20 of



## Delete saved searches

You can delete any custom saved search if you no longer need it. You can't delete default saved searches.

To delete a saved search, select the  button for a specific search, select **Delete search**, then select **Delete search** again in the confirmation dialog.

## Default searches

BlueXP classification provides the following system-defined search queries:

- **Data Subject names - High risk**

Files with more than 50 data subject names

- **Email Addresses - High risk**

Files with more than 50 email addresses or database columns with more than 50% of their rows containing email addresses

- **Personal data - High risk**

Files with more than 20 personal data identifiers or database columns with more than 50% of their rows containing personal data identifiers

- **Private data - Stale over 7 years**

Files containing personal or sensitive personal information, last modified more than 7 years ago

- **Protect - High**

Files or database columns that contain a password, credit card information, IBAN number, or social security number

- **Protect - Low**

Files that have not been accessed for more than 3 years

- **Protect - Medium**

Files that contain files or database columns with personal data identifiers including ID numbers, tax identification numbers, drivers license numbers, medicinal IDs, or passport numbers

- **Sensitive Personal data - High risk**

Files with more than 20 sensitive personal data identifiers or database columns with greater than 50% of their rows containing sensitive personal data

## Change the BlueXP classification scan settings for your repositories

You can manage how your data is being scanned in each of your working environments and data sources. You can make the changes on a "repository" basis; meaning you can

make changes for each volume, schema, user, etc. depending on the type of data source you are scanning.

Some of the things you can change are whether a repository is scanned or not, and whether BlueXP classification is performing a [mapping scan](#) or a [mapping & classification scan](#). You can also pause and resume scanning, for example, if you need to stop scanning a volume for a period of time.

## View the scan status for your repositories

You can view the individual repositories that BlueXP classification is scanning (volumes, buckets, etc.) for each working environment and data source. Additionally, you can see how many have been "Mapped", and how many have been "Classified". Classification takes a longer time as the full AI identification is being performed on all data.

You can view the scanning status of each work environment on the Configuration page:

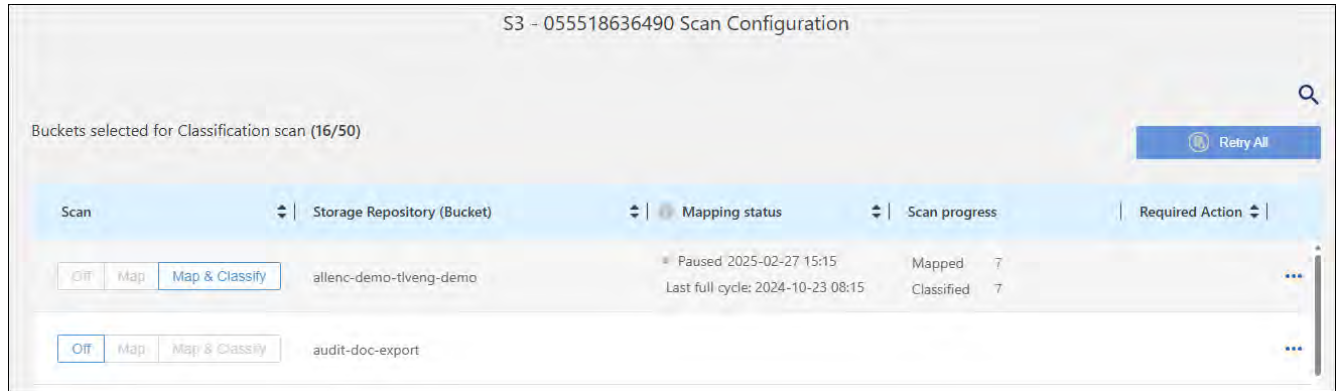
- **Initializing** (light blue dot): The map or classify configuration is activated. This appears for few seconds before starting the "pending queue" status.
- **Pending queue** (orange dot): The scan task is waiting to be listed in the scanning queue.
- **Queued** (orange dot): The task was successfully added to the scanning queue. The system will start mapping or classifying the volume when its turn in the queue arrives.
- **Running** (green dot): The scan task, which was in the queue, is actively in progress on the selected storage repository.
- **Finished** (green dot): The scan of the storage repository is complete.
- **Paused** (gray dot): You selected the "Pause" option to pause scanning. While the changes in the volume are not displayed in the system, the scanned insights are still shown.
- **Error** (red dot): The scan cannot complete because it has encountered issues. If you need to complete an action, the error appears in the tooltip under the "Required action" column. Otherwise, the system shows an "error" status and tries to recover. When it finishes, the status changes.
- **Not scanning**: The volume configuration of "Off" was selected and the system is not scanning the volume.

## Steps

1. From the BlueXP classification menu, select **Configuration**.

The screenshot shows the 'Identity Services' configuration page in the BlueXP interface. The page is titled 'Identity Services' and includes a 'Quick Navigation' sidebar with options for 'Identity Services', 'Working Environments', and 'Scanner Groups'. The main content area displays the configuration for a working environment named 'share2scan.netapp.com'. Below this, there are 11 'Working Environments' listed, with filters for 'S3', 'CVO', 'DB', and 'SHARES'. The selected environment is 'S3 - 055518636490 | 50 Buckets', with a 'Scanner Group name: default' and 'Working Environment ID: S3'. A 'Scan Mode' progress bar shows '16 Classified' (green), '16 Mapped' (blue), and '34 Not Scanned' (gray). A status message indicates 'Continuously scanning all selected Buckets'. The page also features 'Active Directory Integrated' and 'Add Working Environment' buttons at the top right, and an 'Edit' button for the selected environment.

2. From the Configuration tab, select the **Configuration** button for the working environment.
3. In the Scan Configuration page, view the scan settings for all repositories.



4. Hover your cursor over the chart in the *Mapping Status* column to see the number of files that remain to be mapped or classified in each repository (bucket in this example).

## Change the type of scanning for a repository

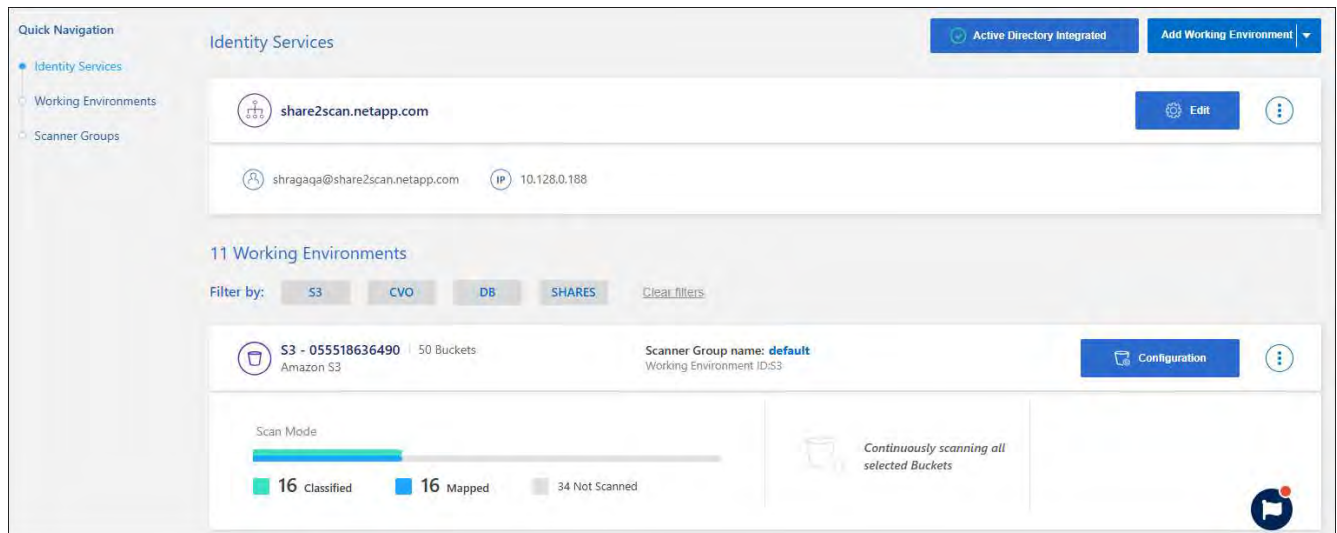
You can start or stop mapping-only scans, or mapping and classification scans, in a working environment at any time from the Configuration page. You can also change from mapping-only scans to mapping and classification scans, and vice-versa.



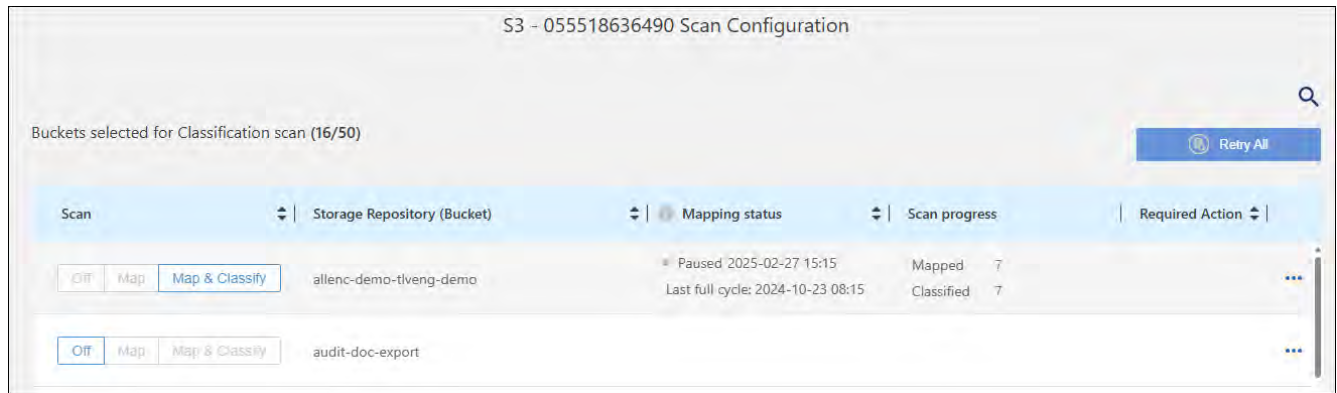
Databases can't be set to mapping-only scans. Database scanning can be Off or On; where On is equivalent to Map & Classify.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration tab, select the **Configuration** button for the working environment.




3. In the Scan Configuration page, change any of the repositories (buckets in this example) to perform **Map** or **Map & Classify** scans.



Certain types of working environments enable you to change the type of scanning globally for all repositories using a button bar at the top of the page. This is valid for Cloud Volumes ONTAP, on-premises ONTAP, Azure NetApp Files, and Amazon FSx for ONTAP systems.

The example below shows this button bar for an Azure NetApp Files system.



## Prioritize scans

You can prioritize the most important mapping-only scans or map & classify scans to ensure high priority scans are completed first.

By default, scans are queued based on the order in which they are initiated. With the ability to prioritize scans, you can move scans to the front of the queue. Multiple scans can be prioritized. Priority is designated in a first-in, first-out order, meaning the first scan you prioritize moves to the front of the queue; the second scan you prioritize becomes second in the queue, and so forth.

Priority is granted on a one-time basis. Automatic rescans of mapping data occur in the default order.

### Steps

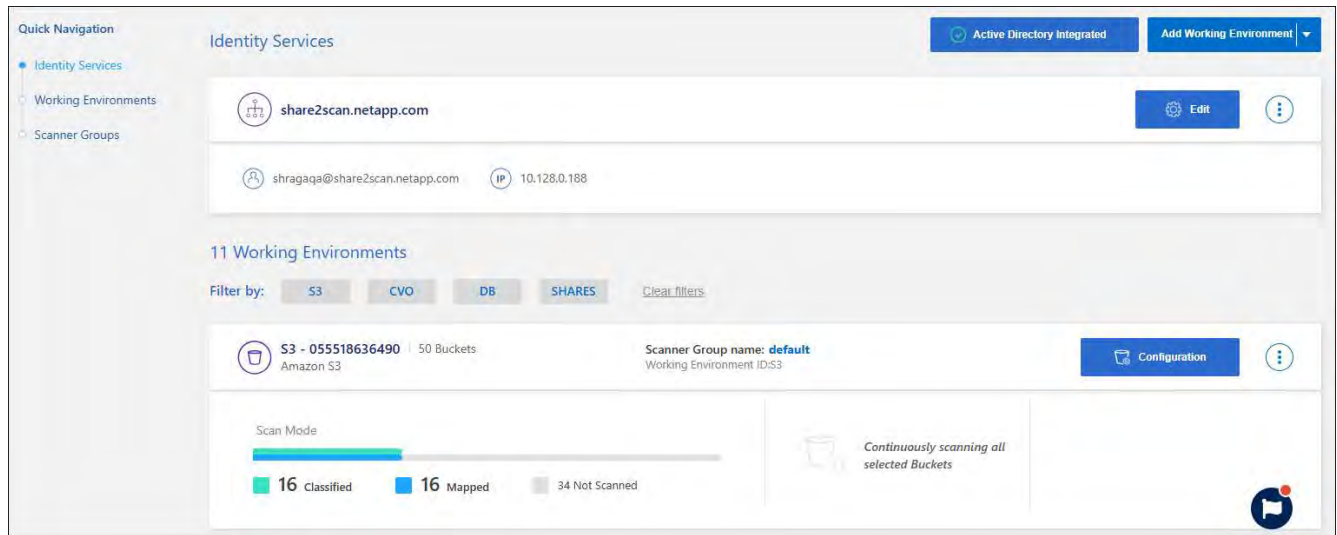
1. From the BlueXP classification menu, select **Configuration**.
2. Select the resources you want to prioritize.
3. From the Actions ... option, select **Prioritize scan**.

## Stop scanning for a repository

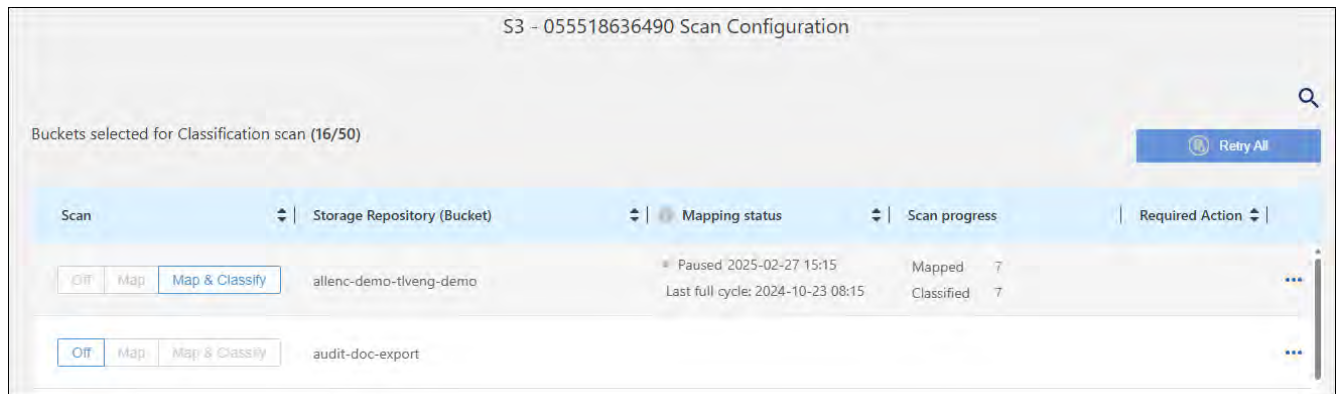
You can stop scanning a repository (for example, a volume) if you no longer need to monitor it for compliance. You do this by turning scanning "off". When scanning is turned off, all the indexing and information about that volume is removed from the system, and charging for scanning the data is stopped.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration tab, select the **Configuration** button for the working environment.



3. In the Scan Configuration page select **Off** to stop scanning for a particular bucket.




## Pause and resume scanning for a repository

You can "pause" scanning on a repository if you want to temporarily stop scanning certain content. Pausing scanning means that BlueXP classification won't perform any future scans for changes or additions to the repository, but that all the current results will still be displayed in the system. Pausing scanning does not stop charging for the scanned the data because the data still exists.

You can "resume" scanning at any time.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration tab, select the **Configuration** button for the working environment.



3. In the Scan Configuration page, select the Actions **...** icon.
4. Select **Pause** to pause scanning for a volume, or select **Resume** to resume scanning for a volume that had been previously paused.

## View BlueXP classification compliance reports

BlueXP classification provides reports that you can use to better understand the status of your organization's data privacy program.

By default, the BlueXP classification dashboards display compliance and governance data for all working environments, databases, and data sources. If you want to view reports that contain data for only some of the working environments, you can filter to see just them.



- The reports described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan can only generate the Data Mapping Report.
- NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that BlueXP classification identifies. You should always validate the information by reviewing the data.

The following reports are available for BlueXP classification:

- **Data Discovery Assessment report:** Provides a high-level analysis of the scanned environment to highlight the system's findings and to show areas of concern and potential remediation steps.
- **Data Mapping report:** Provides information about the size and number of files in your working environments. This includes usage capacity, age of data, size of data, and file types.
- **Data Subject Access Request report:** Enables you to extract a report of all files that contain information regarding a data subject's specific name or personal identifier.
- **HIPAA report:** Helps you identify the distribution of health information across your files.
- **PCI DSS report:** Helps you identify the distribution of credit card information across your files.
- **Privacy Risk Assessment report:** Provides privacy insights from your data and a privacy risk score.
- **Reports on a specific information type:** Reports are available that include details about the identified files that contain personal data and sensitive personal data. You can also see files broken down by



category and file type.

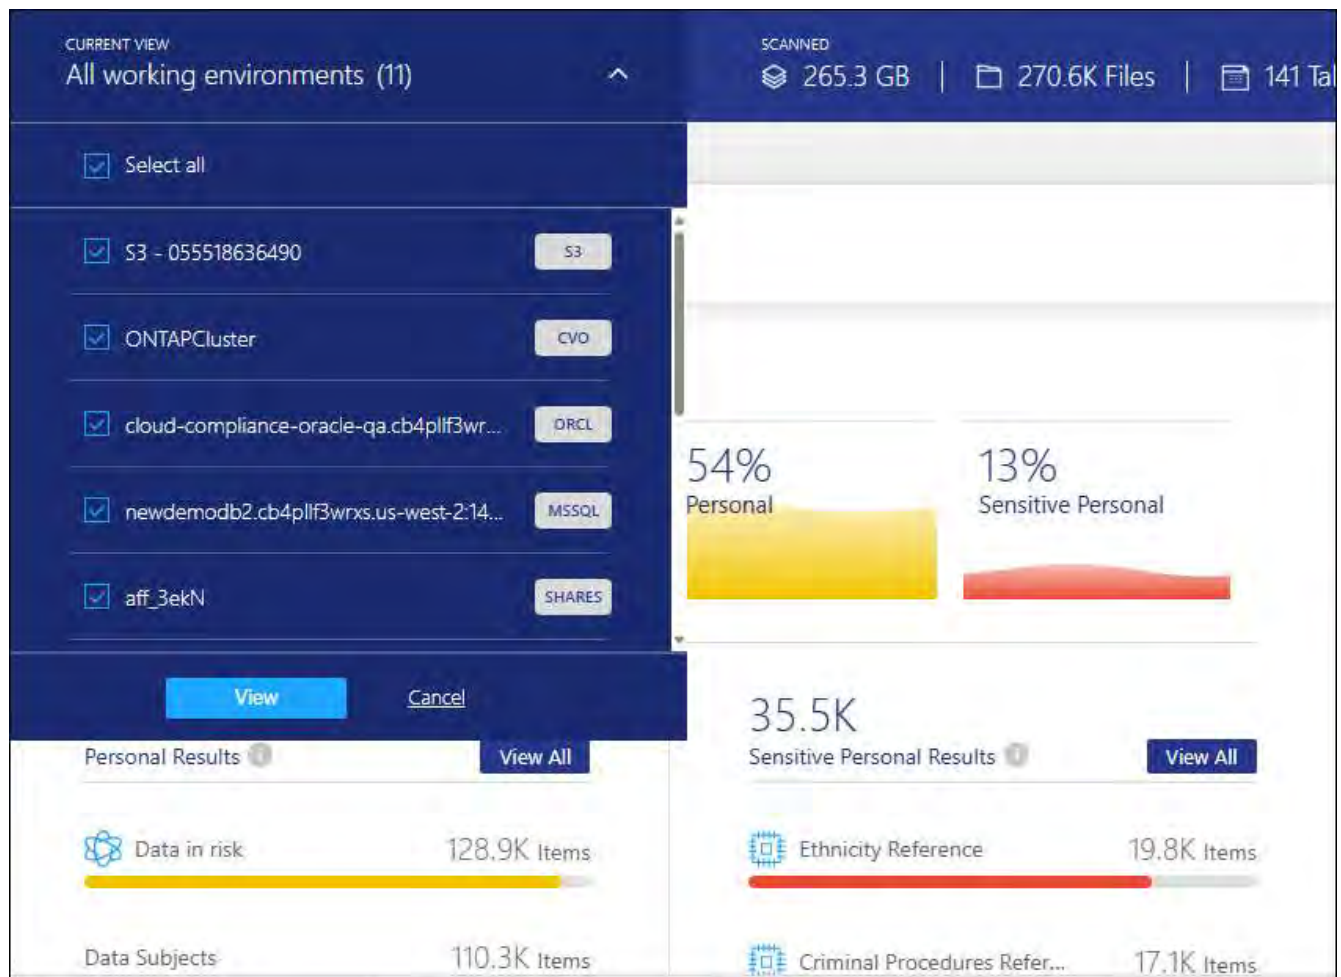
## Select the working environments for reports

You can filter the contents of the BlueXP classification Compliance dashboard to see compliance data for all working environments and databases, or for just specific working environments.

When you filter the dashboard, BlueXP classification scopes the compliance data and reports to just those working environments that you selected.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. Select the Working environments filter drop-down and select the working environments.
3. Select **View**.



## Data Subject Access Request Report

Privacy regulations such as the European GDPR grant data subjects (such as customers or employees) the right to access their personal data. When a data subject requests this information, this is known as a DSAR (data subject access request). Organizations are required to respond to these requests "without undue delay", and at the latest within one month of receipt.

You can respond to a DSAR by searching for a subject's full name or known identifier (such as an email

address) and then downloading a report. The report is designed to aid in your organization's requirement to comply with GDPR or similar data privacy laws.

### How can BlueXP classification help you respond to a DSAR?

When you perform a data subject search, BlueXP classification finds all of the files, buckets, OneDrive, and SharePoint accounts that have that person's name or identifier in it. BlueXP classification checks the latest pre-indexed data for the name or identifier. It doesn't initiate a new scan.

After the search is complete, you can then download the list of files for a Data Subject Access Request report. The report aggregates insights from the data and puts it into legal terms that you can send back to the person.



Data subject search is not supported within databases at this time.

### Search for data subjects and download reports

Search for the data subject's full name or known identifier and then download a file list report or DSAR report. You can search by [any personal information type](#).



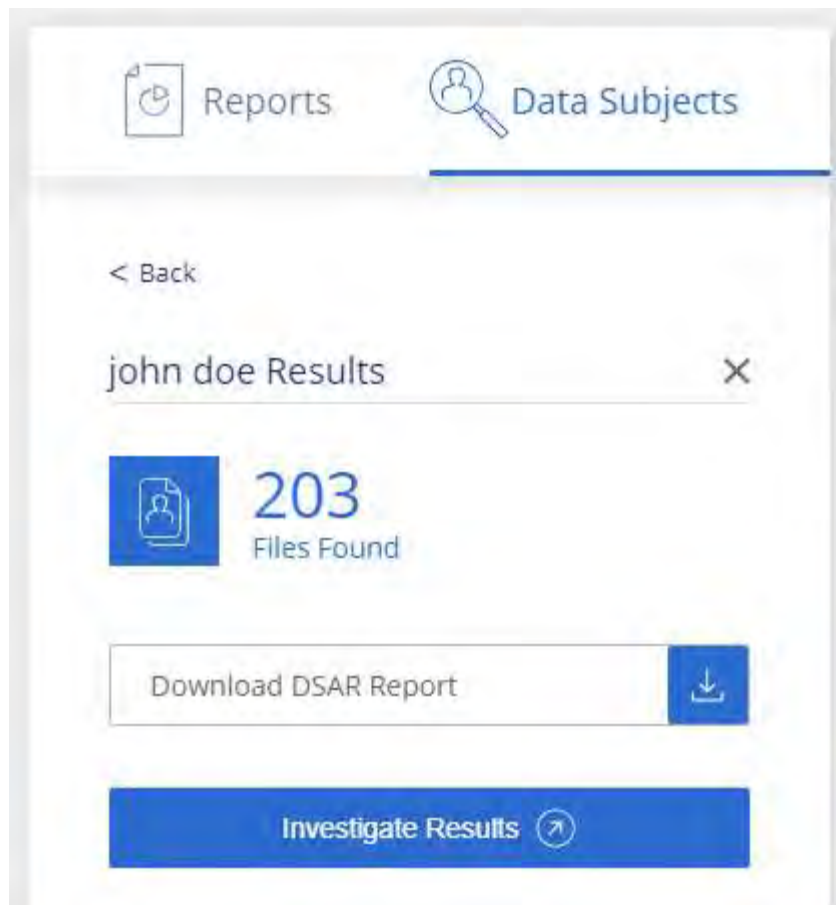
English, German, Japanese, and Spanish are supported when searching for the names of data subjects. Support for more languages will be added later.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. From the Compliance page, scroll down and select **Data Subjects**.
3. Search for the data subject's full name or known identifier.

Here's an example that shows a search for the name *john doe*:





4. Choose one of the available options:

- **Download DSAR Report:** A formal response to the access request that you can send to the data subject. This report contains automatically-generated information based on data that BlueXP classification found on the data subject and is designed to be used as a template. You should complete the form and review it internally before sending it to the data subject.
- **Investigate Results:** A page that enables you to investigate the data by searching, sorting, expanding details for a specific file, and by downloading the file list.



If there are more than 10,000 results, only the top 10,000 appear in the file list.

## Health Insurance Portability and Accountability Act (HIPAA) Report

The Health Insurance Portability and Accountability Act (HIPAA) Report can help you identify files containing health information. It is designed to aid in your organization's requirement to comply with HIPAA data privacy laws. The information BlueXP classification looks for includes:

- Health reference pattern
- ICD-10-CM Medical code
- ICD-9-CM Medical code
- HR - Health category
- Health Application Data category

The report includes the following information:

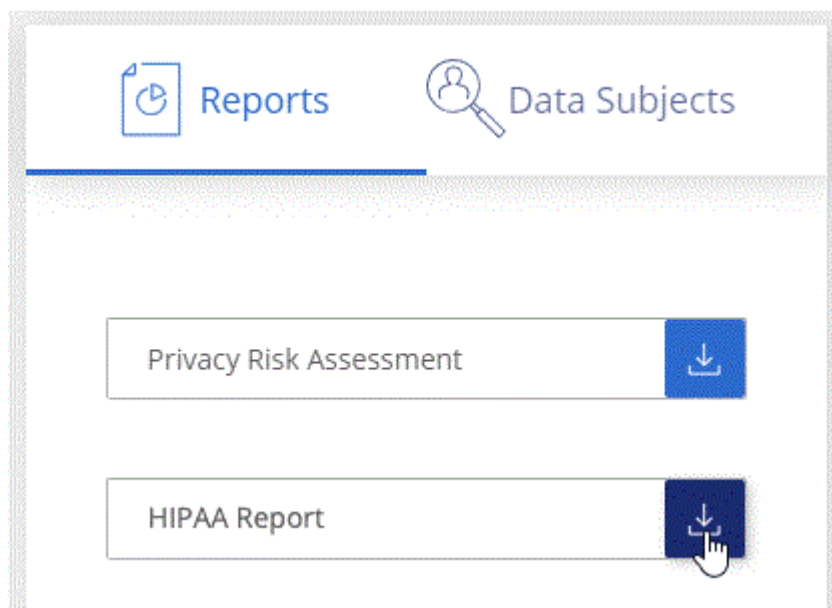
- Overview: How many files contain health information and in which working environments.
- Encryption: The percentage of files containing health information that are on encrypted or unencrypted working environments. This information is specific to Cloud Volumes ONTAP.
- Ransomware Protection: The percentage of files containing health information that are on working environments that do or don't have ransomware protection enabled. This information is specific to Cloud Volumes ONTAP.
- Retention: The timeframe in which the files were last modified. This is helpful because you shouldn't keep health information for longer than you need to process it.
- Distribution of Health Information: The working environments where the health information was found and whether encryption and ransomware protection are enabled.

## Generate the HIPAA Report

Go to the Compliance tab to generate the report.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. Scroll down and locate the **Reports** pane.
3. Select the download icon next to **HIPAA Report.**



### Result

BlueXP classification generates a PDF report that you can review and send to other groups as needed.

## Payment Card Industry Data Security Standard (PCI DSS) Report

The Payment Card Industry Data Security Standard (PCI DSS) Report can help you identify the distribution of credit card information across your files.

The report includes the following information:

- Overview: How many files contain credit card information and in which working environments.

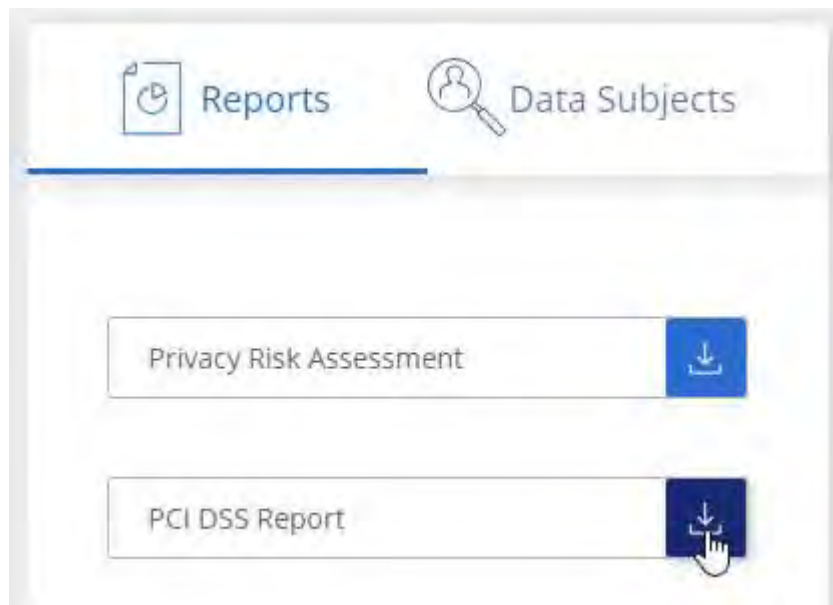
- Encryption: The percentage of files containing credit card information that are on encrypted or unencrypted working environments. This information is specific to Cloud Volumes ONTAP.
- Ransomware Protection: The percentage of files containing credit card information that are on working environments that do or don't have ransomware protection enabled. This information is specific to Cloud Volumes ONTAP.
- Retention: The timeframe in which the files were last modified. This is helpful because you shouldn't keep credit card information for longer than you need to process it.
- Distribution of Credit Card Information: The working environments where the credit card information was found and whether encryption and ransomware protection are enabled.

## Generate the PCI DSS Report

Go to the Compliance tab to generate the report.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. Scroll down and locate the **Reports** pane.
3. Select the download icon next to **PCI DSS Report**.



### Result

BlueXP classification generates a PDF report that you can review and send to other groups as needed.

## Privacy Risk Assessment Report

The Privacy Risk Assessment Report provides an overview of your organization's privacy risk status, as required by privacy regulations such as GDPR and CCPA.

The report includes the following information:

- Compliance status: A severity score and the distribution of data, whether it's non-sensitive, personal, or sensitive personal.
- Assessment overview: A breakdown of the types of personal data found, as well as the categories of data.

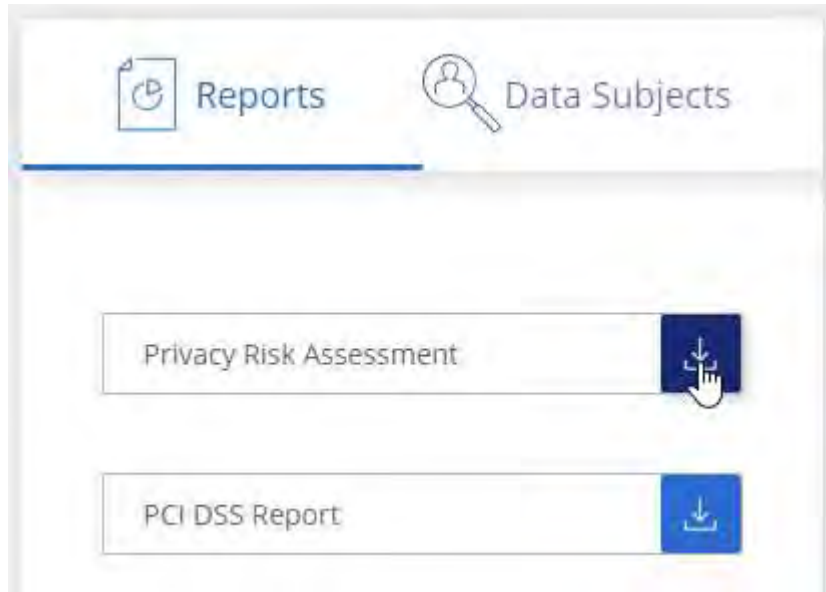
- Data subjects in this assessment: The number of people, by location, for which national identifiers were found.

## Generate the Privacy Risk Assessment Report

Go to the Compliance tab to generate the report.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. Scroll down and locate the **Reports** pane.
3. Select the download icon next to **Privacy Risk Assessment.**



### Result

BlueXP classification generates a PDF report that you can review and send to other groups as needed.

### Severity score

BlueXP classification calculates the severity score for the Privacy Risk Assessment Report on the basis of three variables:

- The percentage of personal data out of all data.
- The percentage of sensitive personal data out of all data.
- The percentage of files that include data subjects, determined by national identifiers such as national IDs, Social Security numbers, and tax ID numbers.

The logic used to determine the score is as follows:

Severity score	Logic
0	All three variables are exactly 0%
1	One of the variables are larger than 0%
2	One of the variables are larger than 3%

<b>Severity score</b>	<b>Logic</b>
3	Two of the variables are larger than 3%
4	Three of the variables are larger than 3%
5	One of the variables are larger than 6%
6	Two of the variables are larger than 6%
7	Three of the variables are larger than 6%
8	One of the variables are larger than 15%
9	Two of the variables are larger than 15%
10	Three of the variables are larger than 15%

# Manage BlueXP classification

## Exclude specific directories from BlueXP classification scans

If you want BlueXP classification to exclude scanning data that resides in certain data source directories, you can add these directory names to a configuration file. After you apply this change, the BlueXP classification engine will exclude scanning data in those directories.

Note that BlueXP classification is configured by default to exclude scanning volume snapshot data because that content is identical to the content in the volume.

This functionality is available in BlueXP classification version 1.29 and greater (starting in March 2024).

### Supported data sources

Excluding specific directories from BlueXP classification scans is supported for NFS and CIFS shares in the following data sources:

- On-premises ONTAP
- Cloud Volumes ONTAP
- Amazon FSx for NetApp ONTAP
- Azure NetApp Files
- General file shares

### Define the directories to exclude from scanning

Before you can exclude directories from classification scanning, you need to log into the BlueXP classification system so you can edit a configuration file and run a script. See how to [log in to the BlueXP classification system](#) depending on whether you manually installed the software on a Linux machine or if you deployed the instance in the cloud.



- You can exclude a maximum of 50 directory paths per BlueXP classification system.
- Excluding directory paths may affect scanning times.

### Steps

1. On the BlueXP classification system, go to `/opt/netapp/config/custom_configuration` and open the file `data_provider.yaml`.
2. In the `"data_providers"` section, under the line `"exclude:"`, enter the directory paths to exclude. For example:

```
exclude:  
- "folder1"  
- "folder2"
```

Do not change anything else in this file.

3. Save the changes to the file.

4. Go to `/opt/netapp/Datasense/tools/customer_configuration/data_providers` and run the following script:

```
update_data_providers_from_config_file.sh
```

This command commits the directories to be excluded from scanning to the classification engine.

## Result

All subsequent scans of your data will exclude scanning of those specified directories.

You can add, edit, or delete items from the exclude list using these same steps. The revised exclude list will be updated after you run the script to commit your changes.

## Examples

### Configuration 1:

Every folder that contains "folder1" anywhere in the name will be excluded from all data sources.

```
data_providers:  
  exclude:  
    - "folder1"
```

### Expected results for paths that will be excluded:

- /CVO1/folder1
- /CVO1/folder1name
- /CVO1/folder10
- /CVO1/\*folder1
- /CVO1/+folder1name
- /CVO1/notfolder10
- /CVO22/folder1
- /CVO22/folder1name
- /CVO22/folder10

### Examples for paths that will not be excluded:

- /CVO1/\*folder
- /CVO1/foldername
- /CVO22/\*folder20

### Configuration 2:

Every folder that contains "folder1" only at the start of the name will be excluded.

```
data_providers:
  exclude:
    - "\\*folder1"
```

#### Expected results for paths that will be excluded:

- /CVO/\*folder1
- /CVO/\*folder1name
- /CVO/\*folder10

#### Examples for paths that will not be excluded:

- /CVO/folder1
- /CVO/folder1name
- /CVO/not\*folder10

#### Configuration 3:

Every folder in data source "CVO22" that contains "folder1" anywhere in the name will be excluded.

```
data_providers:
  exclude:
    - "CVO22/folder1"
```

#### Expected results for paths that will be excluded:

- /CVO22/folder1
- /CVO22/folder1name
- /CVO22/folder10

#### Examples for paths that will not be excluded:

- /CVO1/folder1
- /CVO1/folder1name
- /CVO1/folder10

## Escaping special characters in folder names

If you have a folder name that contains one of the following special characters and you want to exclude data in that folder from being scanned, you'll need to use the escape sequence `\\` before the folder name.

```
., +, *, ?, ^, $, (, ), [, ], {, }, |
```

For example:

Path in source: `/project/*not_to_scan`

Syntax in exclude file: `"\\*not_to_scan"`



## View the current exclusion list

It's possible for the contents of the `data_provider.yaml` configuration file to be different than what has actually been committed after running the `update_data_providers_from_config_file.sh` script. To view the current list of directories that you've excluded from BlueXP classification scanning, run the following command from `/opt/netapp/Datasense/tools/customer_configuration/data_providers`:

```
get_data_providers_configuration.sh
```

## Define additional group IDs as open to organization in BlueXP classification

When group IDs (GIDs) are attached to files or folders in NFS file shares they define the permissions for the file or folder; such as whether they are "open to the organization". If some group IDs (GIDs) are not initially set up with the "Open to Organization" permission level, you can add that permission to the GID so that any files and folders that have that GID attached will be deemed "open to the organization".

After you make this change and BlueXP classification rescans your files and folders, any files and folders that have these group IDs attached will show this permission in the Investigation Details page, and they'll also appear in reports where you are displaying file permissions.

To activate this functionality, you need to log into the BlueXP classification system so you can edit a configuration file and run a script. See how to [log in to the BlueXP classification system](#) depending on whether you manually installed the software on a Linux machine or if you deployed the instance in the cloud.

### Add the "open to organization" permission to group IDs

You need to have the group ID numbers (GIDs) before starting this task.

#### Steps

1. On the BlueXP classification system, go to `/opt/netapp/config/custom_configuration` and open the file `data_provider.yaml`.
2. In the line `organization_group_ids: []` add the group IDs. For example:

```
organization_group_ids: [1014, 1015, 21, 2021, 1013, 2020, 1018, 1019]
```

Do not change anything else in this file.

3. Save the changes to the file.
4. Go to `/opt/netapp/Datasense/tools/customer_configuration/data_providers` and run the following script:

```
update_data_providers_from_config_file.sh
```

This command commits the revised group ID permissions to the classification engine.

## Result

All subsequent scans of your data will identify files or folders that have these group IDs attached as "open to organization".

You can edit the list of group IDs and delete any group IDs you added in the past using these same steps. The revised list of group IDs will be updated after you run the script to commit your changes.

## View the current list of group IDs

It's possible for the contents of the `data_provider.yaml` configuration file to differ from what has actually been committed after running the `update_data_providers_from_config_file.sh` script. To view the current list of group IDs that you've added to BlueXP classification, run the following command from `/opt/netapp/Datasense/tools/customer_configuration/data_providers`:


```
get_data_providers_configuration.sh
```

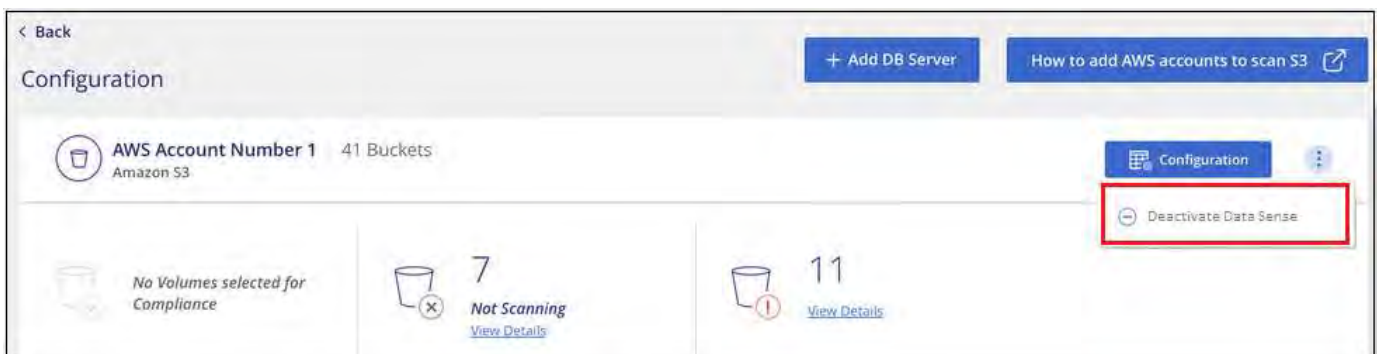
## Remove data sources from BlueXP classification

If you need to, you can stop BlueXP classification from scanning one or more working environments, databases, or file share groups.

### Deactivate compliance scans for a working environment

When you deactivate scans, BlueXP classification no longer scans the data on the working environment and it removes the indexed compliance insights from the BlueXP classification instance (the data from the working environment itself isn't deleted).


1. From the *Configuration* page, select the  button in the row for the working environment then **Deactivate Data Sense**.



You can also disable compliance scans for a working environment from the Services panel when you select the working environment.

## Remove a database from BlueXP classification

If you no longer want to scan a certain database, you can delete it from the BlueXP classification interface and stop all scans.


1. From the *Configuration* page, select the  button in the row for the database then **Remove DB Server**.



## Remove a group of file shares from BlueXP classification

If you no longer want to scan user files from a file shares group, you can delete the File Shares Group from the BlueXP classification interface and stop all scans.

### Steps

1. From the *Configuration* page, select the  button in the row for the File Shares Group then **Remove File Shares Group**.



2. Click **Delete Group of Shares** from the confirmation dialog.


## Uninstall BlueXP classification

You can uninstall BlueXP classification to troubleshoot issues or to permanently remove the software from the host. Deleting the instance also deletes the associated disks where the indexed data resides, meaning all the information BlueXP classification has scanned will be permanently deleted.

The steps that you need to use depend on whether you deployed BlueXP classification in the cloud or on an on-premises host.

### Uninstall BlueXP classification from a cloud deployment

You can uninstall and delete the BlueXP classification instance from the cloud provider environment if you no longer want to use BlueXP classification.

1. At the top of the BlueXP classification page, select  then **Uninstall Classification**.



2. In the *Uninstall Classification* dialog, type **uninstall** to confirm that you want to disconnect the BlueXP classification instance from the BlueXP Connector then select **Uninstall**.
3. Go to your cloud provider's console and delete the BlueXP classification instance. The instance is named *CloudCompliance* with a generated hash (UUID) concatenated to it. For example: *CloudCompliance-16bb6564-38ad-4080-9a92-36f5fd2f71c7*

This deletes the instance and all associated data that had been collected by BlueXP classification.

## Uninstall BlueXP classification from an on-premises deployment

You can uninstall BlueXP classification from a host if you no longer want to use BlueXP classification, or if you had an issue that requires reinstallation.

1. At the top of the BlueXP classification page, select  then **Uninstall Classification**.



2. In the *Uninstall Classification* dialog, type **uninstall** to confirm that you want to disconnect the BlueXP classification instance from the BlueXP Connector then select **Uninstall**.
3. To uninstall the software from the host, run the `cleanup.sh` script on the BlueXP classification host machine, for example:

```
cleanup.sh
```

The script is located in the `/install/light_probe/onprem_installer/cleanup.sh` directory.

See how to [log in to the BlueXP classification host machine](#).

# Deprecated features

## BlueXP classification deprecated features

BlueXP classification is available as a core capability within BlueXP at no additional charge. By including BlueXP classification as a core BlueXP capability available to all customers, NetApp is enabling you to access tailored data management with core features.

There are some features and functionality that are deprecated in the BlueXP core version starting with version 1.31 and later and are still supported in legacy versions 1.30 and earlier.

### Supported data sources

Data source	Legacy versions 1.30 and earlier	BlueXP core versions 1.31 and later
Cloud Volumes ONTAP (deployed in AWS, Azure, or GCP)	Yes	Yes
On-premises ONTAP clusters	Yes	Yes
StorageGRID	Yes	Yes
Azure NetApp Files	Yes	Yes
Amazon FSx for ONTAP	Yes	Yes
Google Cloud NetApp Volumes	Yes	Yes
Cloud Volumes Service for Google Cloud	Yes	Yes
Databases	Yes	Yes
Amazon S3	Yes	No
Google Cloud Storage	Yes	No
OneDrive	Yes	No
SharePoint Online	Yes	No
SharePoint On-premises (SharePoint Server)	Yes	No
Google Drive	Yes	No

### Compliance features

Feature	Legacy versions 1.30 and earlier	BlueXP core versions 1.31 and later
Identify Personal Identifiable Information (PII)	Yes	Yes
Identify sensitive personal information	Yes	Yes
Respond to Data Subject Access Requests (DSAR)	Yes	Yes

<b>Feature</b>	<b>Legacy versions 1.30 and earlier</b>	<b>BlueXP core versions 1.31 and later</b>
Create a custom list of "personal data" that is identified	Yes	No
Notify users through email when files contain certain PII. (You define this criteria using <a href="#">Policies</a> .)	Yes	No
Use directory-level filters	Yes	Yes
Use directory-level PII analysis	Yes	No

## Features to manage your data

<b>Feature</b>	<b>Legacy versions 1.30 and earlier</b>	<b>BlueXP core versions 1.31 and later</b>
Move, copy, and delete source files	Yes	No
Categorize data using Status tags	Yes	No
Categorize data using AIP labels	Yes	No
Assign files to users	Yes	No
Rescan data on demand	Yes	No
Create custom classifiers	Yes	No
Exclude directories from scanning	Yes	Yes
Search for names within files	Yes	Yes
Export data to NFS/CIFS from investigation	Yes	Yes
Export data to CSV from investigation	Yes	Yes
Support multiple scanners	Yes	No
Integrate Active Directory	Yes	Yes
Use permission analysis and filters	Yes	Yes
Use the file card	Yes	Yes
Use the heatmap	Yes	Yes
Use actions on Dashboard and file card	Yes	No
Use file access audit logging	Yes	No
Enable file access from the Configuration page	Yes	No
Use certain predefined policies	Yes	No

## Deploy BlueXP classification deprecations

## Install BlueXP classification on multiple hosts for large configurations with no internet access

Complete a few steps to install BlueXP classification on multiple hosts in an on-premises site that doesn't have internet access - also known as *private mode*. This type of installation is perfect for your secure sites.

For very large configurations where you'll be scanning petabytes of data in sites without internet access, you can include multiple hosts to provide additional processing power. When using multiple host systems, the primary system is called the *Manager node* and the additional systems that provide extra processing power are called *Scanner nodes*.

Follow these steps when installing BlueXP classification software on multiple on-premises hosts in an offline environment.



The following information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

### Before you begin

- Verify that all your Linux systems for the Manager and Scanner nodes meet the host requirements.
- Verify that you have installed the two prerequisite software packages (Docker Engine or Podman, and Python 3).
- Make sure you have root privileges on the Linux systems.
- Verify that your offline environment meets the required permissions and connectivity.
- You must have the IP addresses of the scanner node hosts that you plan to use.
- The following ports and protocols must be enabled on all hosts:

Port	Protocols	Description
2377	TCP	Cluster management communications
7946	TCP, UDP	Inter-node communication
4789	UDP	Overlay network traffic
50	ESP	Encrypted IPsec overlay network (ESP) traffic
111	TCP, UDP	NFS Server for sharing files between the hosts (needed from each scanner node to manager node)
2049	TCP, UDP	NFS Server for sharing files between the hosts (needed from each scanner node to manager node)

### Steps

1. Follow steps 1 through 8 from the [Single-host installation](#) on the manager node.
2. As shown in step 9, when prompted by the installer, you can enter the required values in a series of prompts, or you can provide the required parameters as command line arguments to the installer.

In addition to the variables available for a single-host installation, a new option `-n <node_ip>` is used to specify the IP addresses of the scanner nodes. Multiple node IPs are separated by a comma.

For example, this command adds 3 scanner nodes:

```
sudo ./install.sh -a <account_id> -c <client_id> -t <user_token> --host  
<ds_host> --manager-host <cm_host> -n <node_ip1>,<node_ip2>,<node_ip3> --no  
-proxy --darksite
```

3. Before the manager node installation completes, a dialog displays the installation command needed for the scanner nodes. Copy the command (for example: `sudo ./node_install.sh -m 10.11.12.13 -t ABCDEF-1-3u69m1-1s35212`) and save it in a text file.
4. On **each** scanner node host:
  - a. Copy the Data Sense installer file (**cc\_onprem\_installer.tar.gz**) to the host machine.
  - b. Unzip the installer file.
  - c. Paste and run the command that you copied in step 3.

When the installation finishes on all scanner nodes and they have been joined to the manager node, the manager node installation finishes as well.

## Result

The BlueXP classification installer finishes installing packages, and registers the installation. Installation can take 15 to 25 minutes.

## What's Next

From the Configuration page you can select the local [on-prem ONTAP clusters](#) and local [databases](#) that you want to scan.

# Scan data deprecations

## Scan Amazon S3 buckets with BlueXP classification

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

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

## Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

### Set up the S3 requirements in your cloud environment

Ensure that your cloud environment can meet the requirements for BlueXP classification, including preparing an IAM role and setting up connectivity from BlueXP classification to S3. [See the complete list.](#)

2

### Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.



**3**

### Activate BlueXP classification on your S3 working environment

Select the Amazon S3 working environment, click **Enable**, and select an IAM role that includes the required permissions.

**4**

### Select the buckets to scan

Select the buckets that you'd like to scan and BlueXP classification will start scanning them.

## Reviewing S3 prerequisites

The following requirements are specific to scanning S3 buckets.

### Set up an IAM role for the BlueXP classification instance

BlueXP classification needs permissions to connect to the S3 buckets in your account and to scan them. Set up an IAM role that includes the permissions listed below. BlueXP prompts you to select an IAM role when you enable BlueXP classification on the Amazon S3 working environment.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:Get*",
        "s3:List*",
        "s3:PutObject"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedRolePolicies"
      ],
      "Resource": [
        "arn:aws:iam::*:policy/*",
        "arn:aws:iam::*:role/*"
      ]
    }
  ]
}
```

## Provide connectivity from BlueXP classification to Amazon S3

BlueXP classification needs a connection to Amazon S3. The best way to provide that connection is through a VPC Endpoint to the S3 service. For instructions, see [AWS Documentation: Creating a Gateway Endpoint](#).

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

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

An alternative is to provide the connection by using a NAT Gateway.



You can't use a proxy to get to S3 over the internet.

## Deploying the BlueXP classification instance

[Deploy BlueXP classification in BlueXP](#) if there isn't already an instance deployed.

You need to deploy the instance using a Connector deployed in AWS so that BlueXP automatically discovers the S3 buckets in this AWS account and displays them in an Amazon S3 working environment.

**Note:** Deploying BlueXP classification in an on-premises location is not currently supported when scanning S3 buckets.

Upgrades to BlueXP classification software are automated as long as the instance has internet connectivity.

## Activating BlueXP classification on your S3 working environment


Enable BlueXP classification on Amazon S3 after you verify the prerequisites.

### Steps


1. From the BlueXP left navigation menu, click **Storage > Canvas**.
2. Select the Amazon S3 working environment.



3. In the Services pane on the right, click **Enable** next to **Classification**.




4. When prompted, assign an IAM role to the BlueXP classification instance that has [the required permissions](#).



5. Select **Enable**.



You can also enable compliance scans for a working environment from the Configuration page by selecting the  button then **Activate BlueXP classification**.

### Result

BlueXP assigns the IAM role to the instance.

### Enabling and disabling compliance scans on S3 buckets


After BlueXP enables BlueXP classification on Amazon S3, the next step is to configure the buckets that you want to scan.

When BlueXP is running in the AWS account that has the S3 buckets you want to scan, it discovers those buckets and displays them in an Amazon S3 working environment.


BlueXP classification can also [scan S3 buckets that are in different AWS accounts](#).

### Steps

1. Select the Amazon S3 working environment.
2. In the Services pane on the right, click **Configure Buckets**.



3. Enable mapping-only scans, or mapping and classification scans, on your buckets.



To:	Do this:
Enable mapping-only scans on a bucket	Click <b>Map</b>
Enable full scans on a bucket	Click <b>Map &amp; Classify</b>
Disable scanning on a bucket	Click <b>Off</b>

### Result

BlueXP classification starts scanning the S3 buckets that you enabled. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

### Scanning buckets from additional AWS accounts

You can scan S3 buckets that are under a different AWS account by assigning a role from that account to access the existing BlueXP classification instance.

### Steps

1. Go to the target AWS account where you want to scan S3 buckets and create an IAM role by selecting **Another AWS account**.

## Create role



### Select type of trusted entity

 <b>AWS service</b> EC2, Lambda and others	 <b>Another AWS account</b> Belonging to you or 3rd party	 <b>Web identity</b> Cognito or any OpenID provider	 <b>SAML 2.0 federation</b> Your corporate directory
--	---	---	--

Allows entities in other accounts to perform actions in this account. [Learn more](#)

### Specify accounts that can use this role

Account ID\*

- Options
- Require external ID (Best practice when a third party will assume this role)
  - Require MFA

Be sure to do the following:

- Enter the ID of the account where the BlueXP classification instance resides.
- Change the **Maximum CLI/API session duration** from 1 hour to 12 hours and save that change.
- Attach the BlueXP classification IAM policy. Make sure it has the required permissions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:Get*",
        "s3:List*",
        "s3:PutObject"
      ],
      "Resource": "*"
    }
  ]
}
```

2. Go to the source AWS account where the BlueXP classification instance resides and select the IAM role that is attached to the instance.
  - a. Change the **Maximum CLI/API session duration** from 1 hour to 12 hours. Save the change.
  - b. Select **Attach policies** then **Create policy**.
  - c. Create a policy that includes the "sts:AssumeRole" action and specify the ARN of the role that you created in the target account.

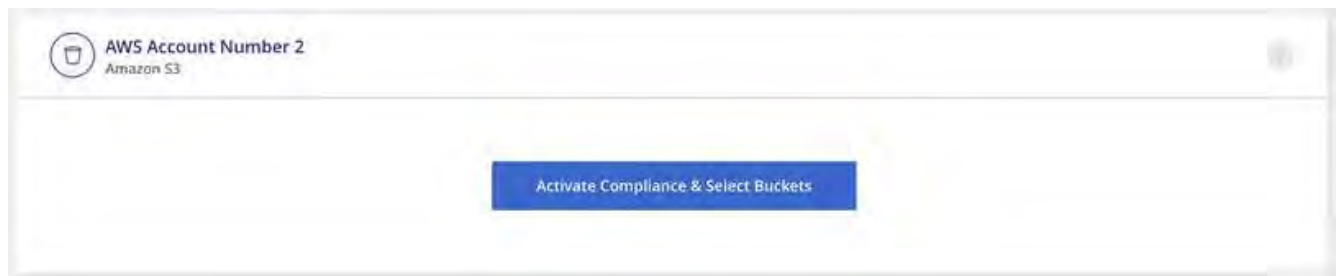
```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::<ADDITIONAL-ACCOUNT-
ID>:role/<ADDITIONAL_ROLE_NAME>"
    },
    {
      "Effect": "Allow",
      "Action": [
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedRolePolicies"
      ],
      "Resource": [
        "arn:aws:iam::*:policy/*",
        "arn:aws:iam::*:role/*"
      ]
    }
  ]
}

```

The BlueXP classification instance profile account receives access to the additional AWS account.

3. Navigate to the **Amazon S3 Configuration** page and the new AWS account is displayed. Note that it can take a few minutes for BlueXP classification to sync the new account's working environment and show this information.



4. Click **Activate BlueXP classification & Select Buckets** and select the buckets you want to scan.

### Result

BlueXP classification starts scanning the new S3 buckets that you enabled.

## Scan OneDrive accounts with BlueXP classification

Complete a few steps to start scanning files in your user's OneDrive folders with BlueXP classification.

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

## Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

### Review OneDrive prerequisites

Ensure that you have the Admin credentials to log into the OneDrive account.

2

### Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

3

### Add the OneDrive account

Using Admin user credentials, log into the OneDrive account that you want to access so that it is added as a new working environment.

4

### Add the users and select the type of scanning

Add the list of users from the OneDrive account that you want to scan and select the type of scanning. You can add up to 100 users at time.

## Reviewing OneDrive requirements

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.

- You must have the Admin login credentials for the OneDrive for Business account that provides read access to the user's files.
- You'll need a line-separated list of the email addresses for all the users whose OneDrive folders you want to scan.

## Deploying the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

BlueXP classification can be [deployed in the cloud](#) or [in an on-premises location that has internet access](#).

Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

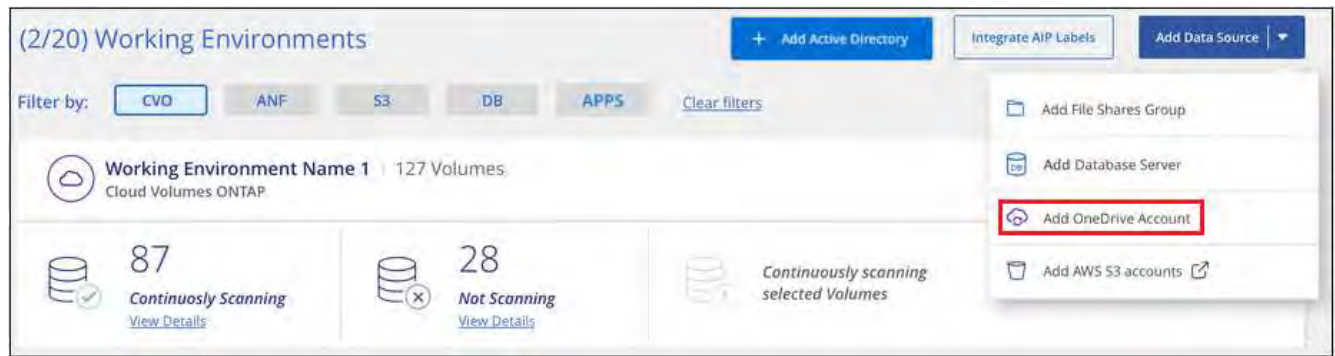
## Adding the OneDrive account

Add the OneDrive account where the user files reside.

### Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add OneDrive Account**.





2. In the Add a OneDrive account dialog, click **Sign in to OneDrive**.
3. In the Microsoft page that appears, select the OneDrive account and enter the required Admin user and password, then click **Accept** to allow BlueXP classification to read data from this account.

The OneDrive account is added to the list of working environments.

### Adding OneDrive users to compliance scans

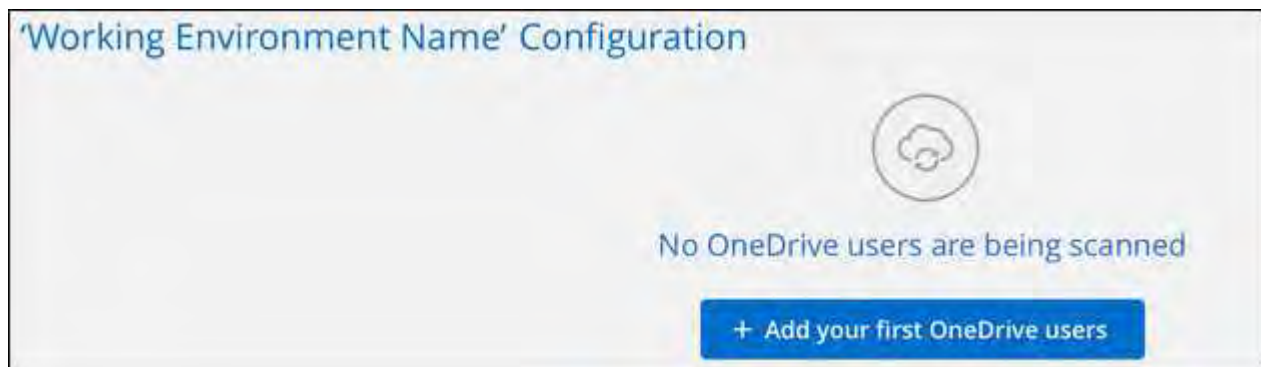
You can add individual OneDrive users, or all of your OneDrive users, so that their files will be scanned by BlueXP classification.

#### Steps


1. From the *Configuration* page, click the **Configuration** button for the OneDrive account.




2. If this is the first time adding users for this OneDrive account, click **Add your first OneDrive users**.



If you are adding additional users from a OneDrive account, click **Add OneDrive users**.



3. Add the email addresses for the users whose files you want to scan - one email address per line (up to 100 maximum per session) - and click **Add Users**.



A confirmation dialog displays the number of users who were added.

If the dialog lists any users who could not be added, capture this information so that you can resolve the issue. In some cases you can re-add the user with a corrected email address.

4. Enable mapping-only scans, or mapping and classification scans, on user files.

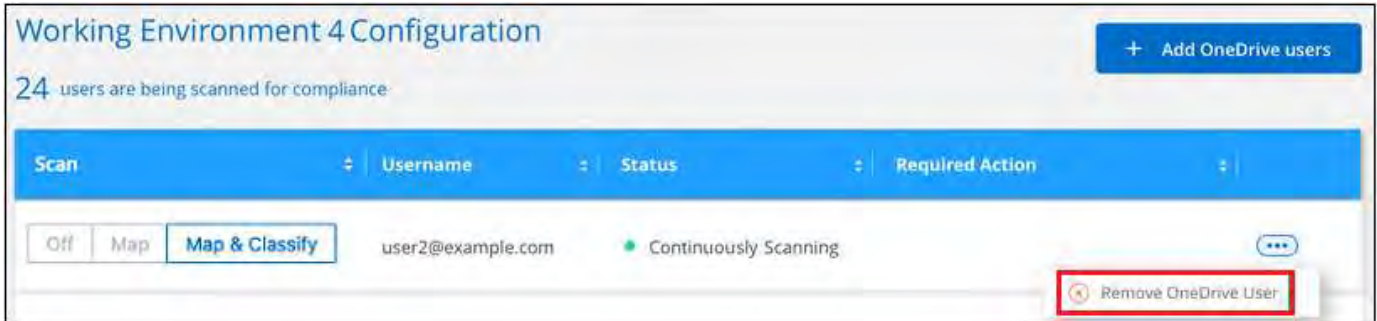
To:	Do this:
Enable mapping-only scans on user files	Click <b>Map</b>
Enable full scans on user files	Click <b>Map &amp; Classify</b>
Disable scanning on user files	Click <b>Off</b>

## Result

BlueXP classification starts scanning the files for the users you added, and the results are displayed in the Dashboard and in other locations.

### Removing a OneDrive user from compliance scans

If users leave the company or if their email address changes, you can remove individual OneDrive users from having their files scanned at any time. Just click **Remove OneDrive User** from the Configuration page.



### Scan SharePoint accounts with BlueXP classification

Complete a few steps to start scanning files in your SharePoint Online and SharePoint On-Premise accounts with BlueXP classification.

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

### Review SharePoint requirements

Review the following prerequisites to make sure you are ready to activate BlueXP classification on a SharePoint account.

- You must have the Admin user login credentials for the SharePoint account that provides read access to all SharePoint sites.
  - For SharePoint Online you can use a non-Admin account, but that user must have permission to access all the SharePoint sites that you want to scan.
- For SharePoint On-Premise, you'll also need the URL of the SharePoint Server.
- You will need a line-separated list of the SharePoint site URLs for all the data you want to scan.

### Deploy the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

- For SharePoint Online, BlueXP classification can be [deployed in the cloud](#).
- For SharePoint On-Premises, BlueXP classification can be installed [in an on-premises location that has internet access](#) or [in an on-premises location that does not have internet access](#).

When BlueXP classification is installed in a site without internet access, the BlueXP Connector also must be installed in that same site without internet access. [Learn more](#).

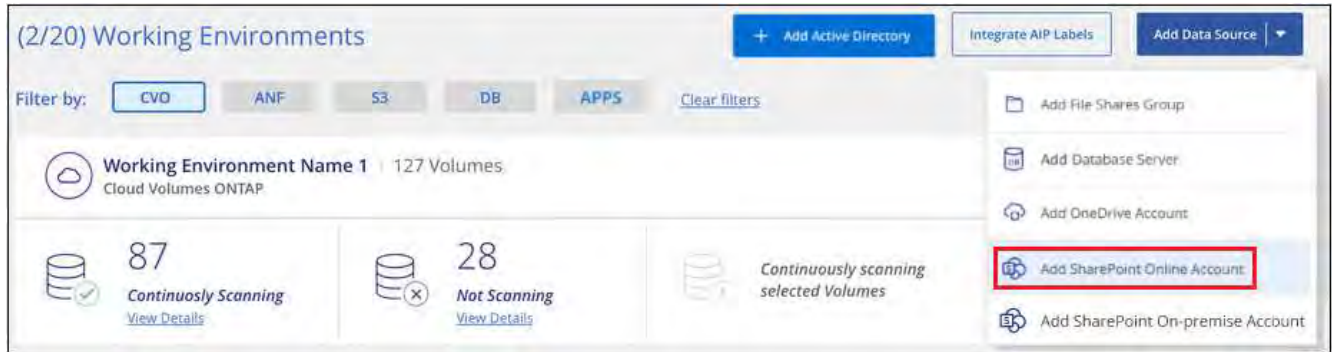
Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

## Add a SharePoint Online account

Add the SharePoint Online account where the user files reside.

### Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add SharePoint Online Account**.



2. In the Add a SharePoint Online Account dialog, click **Sign in to SharePoint**.
3. In the Microsoft page that appears, select the SharePoint account and enter the user and password (Admin user or other user with access to the SharePoint sites), then click **Accept** to allow BlueXP classification to read data from this account.

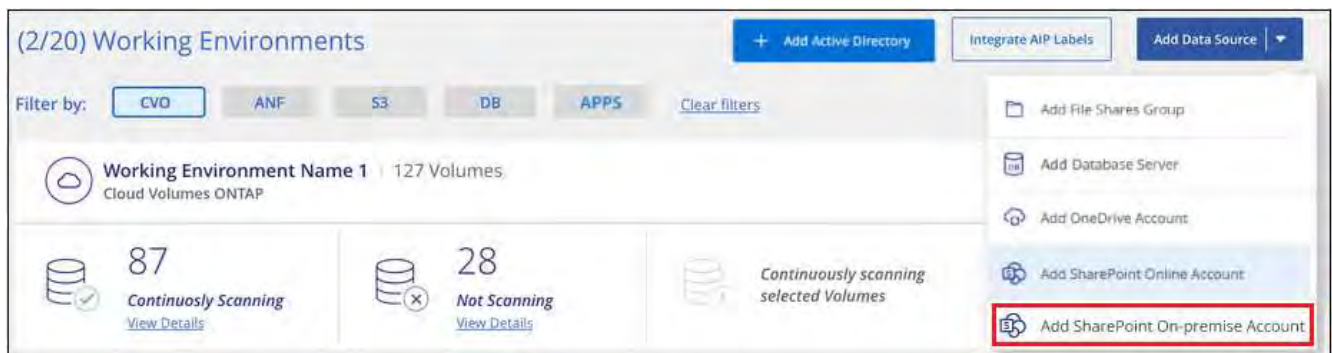
The SharePoint Online account is added to the list of working environments.

## Add a SharePoint On-premise account

Add the SharePoint On-premise account where the user files reside.

### Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add SharePoint On-premise Account**.



2. In the Log into the SharePoint On-Premise Server dialog, enter the following information:
  - Admin user in the format "domain/user" or "user@domain", and admin password
  - URL of the SharePoint Server

Log into the SharePoint On-Premises Server

To activate Data Sense on your SharePoint business account, sign in to SharePoint with an Admin user.

Username  Password

URL

3. Click **Connect**.

The SharePoint On-premise account is added to the list of working environments.

### Add SharePoint sites to compliance scans


You can add individual SharePoint sites, or up to 1,000 SharePoint sites in the account, so that the associated files will be scanned by BlueXP classification. The steps are the same whether you are adding SharePoint Online or SharePoint On-premise sites.

#### Steps

1. From the *Configuration* page, click the **Configuration** button for the SharePoint account.

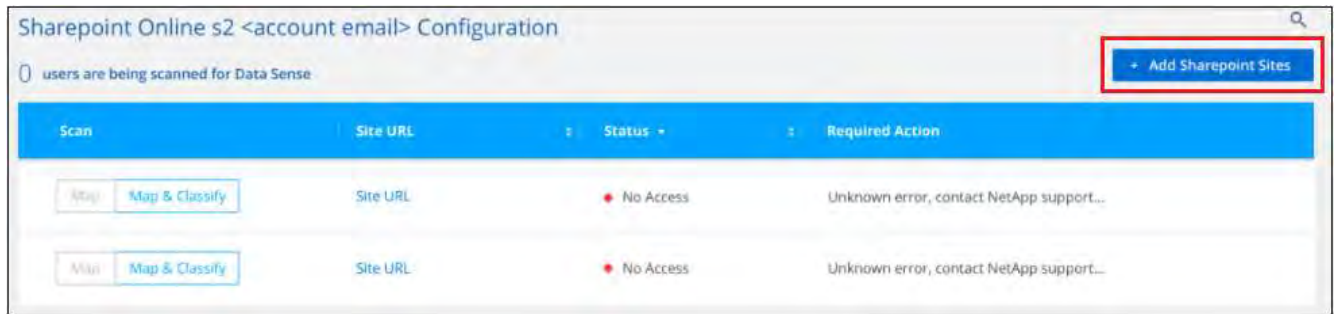


2. If this is the first time adding sites for this SharePoint account, click **Add your first SharePoint site**.

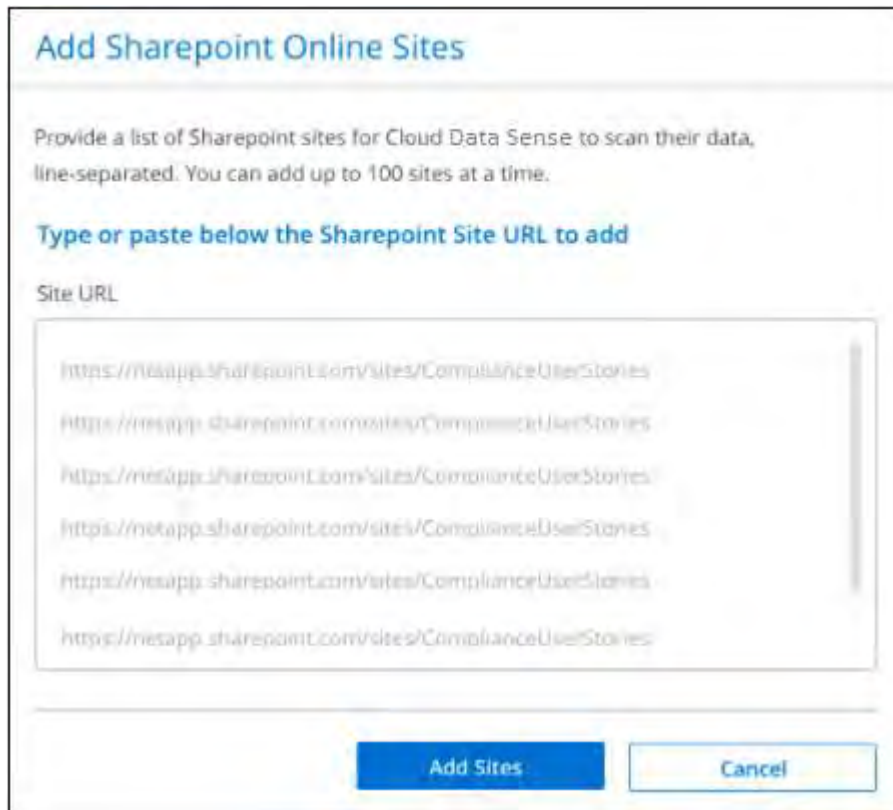


If you are adding additional users from a SharePoint account, click **Add SharePoint Sites**.





3. Add the URLs for the sites whose files you want to scan - one URL per line (up to 100 maximum per session) - and click **Add Sites**.



A confirmation dialog displays the number of sites that were added.

If the dialog lists any sites that could not be added, capture this information so that you can resolve the issue. In some cases you can re-add the site with a corrected URL.

4. If you need to add more than 100 sites for this account, just click **Add SharePoint Sites** again until you have added all your sites for this account (up to 1,000 sites total for each account).
5. Enable mapping-only scans, or mapping and classification scans, on the files in the SharePoint sites.

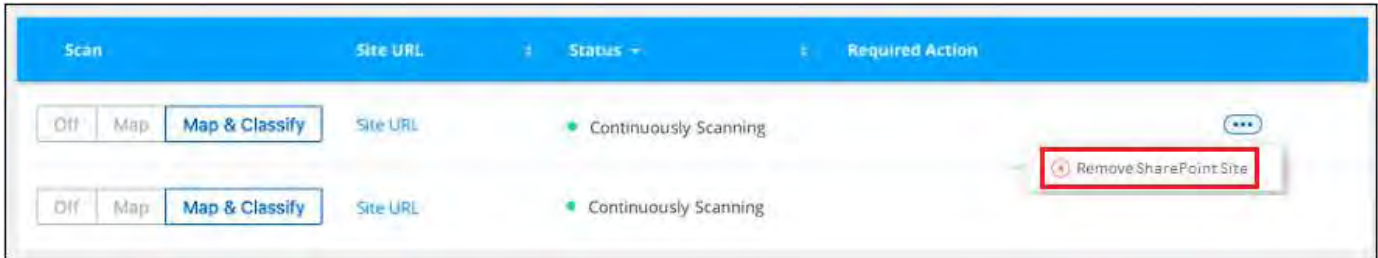
To:	Do this:
Enable mapping-only scans on files	Click <b>Map</b>
Enable full scans on files	Click <b>Map &amp; Classify</b>
Disable scanning on files	Click <b>Off</b>

## Result

BlueXP classification starts scanning the files in the SharePoint sites you added, and the results are displayed in the Dashboard and in other locations.

## Remove a SharePoint site from compliance scans

If you remove a SharePoint site in the future, or decide not to scan files in a SharePoint site, you can remove individual SharePoint sites from having their files scanned at any time. Just click **Remove SharePoint Site** from the Configuration page.



Note that you can [delete the entire SharePoint account from BlueXP classification](#) if you no longer want to scan any user data from the SharePoint account.

## Scan Google Drive accounts with BlueXP classification

Complete a few steps to start scanning user files in your Google Drive accounts with BlueXP classification.

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

### Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

#### 1 Review Google Drive prerequisites

Ensure that you have the Admin credentials to log into the Google Drive account.

#### 2 Deploy BlueXP classification

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

#### 3 Log into the Google Drive account

Using Admin user credentials, log into the Google Drive account that you want to access so that it is added as a new data source.

#### 4 Select the type of scanning for the user files

Select the type of scanning you want to perform on the user files; mapping or mapping and classifying.

## Review Google Drive requirements

Review the following prerequisites to make sure you are ready to enable BlueXP classification on a Google Drive account.

- You must have the Admin login credentials for the Google Drive account that provides read access to the user's files

## Current restrictions

The following BlueXP classification features are not currently supported with Google Drive files:

- When viewing files in the Data Investigation page, the actions in the button bar aren't active. You can't copy, move, delete, etc. any files.
- Permissions can't be identified within files in Google Drive, so no permission information is displayed in the Investigation page.

## Deploy BlueXP classification

Deploy BlueXP classification if there isn't already an instance deployed.

BlueXP classification can be [deployed in the cloud](#) or [in an on-premises location that has internet access](#).

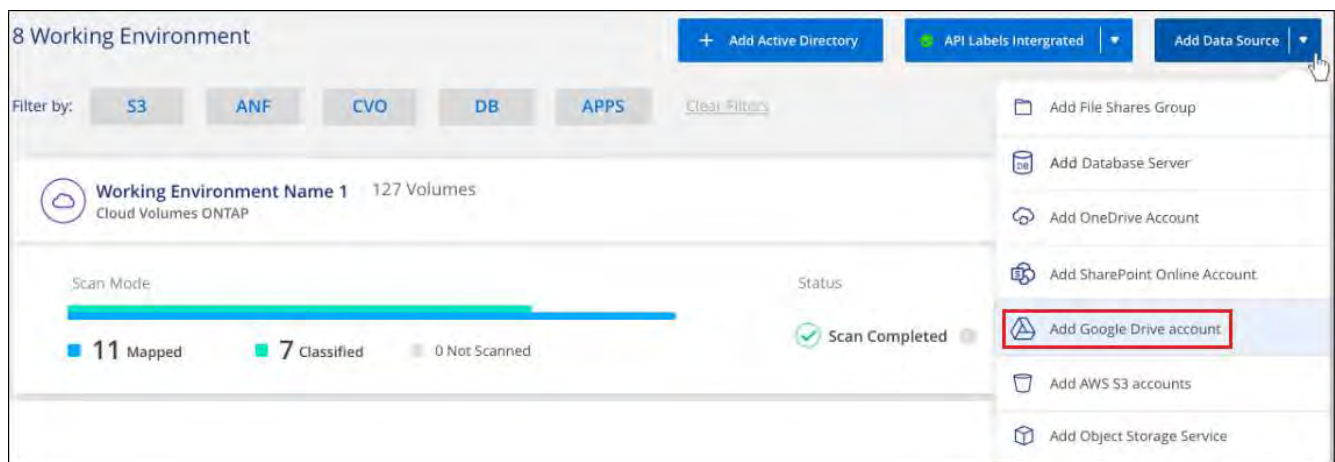
Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

## Add the Google Drive account

Add the Google Drive account where the user files reside. If you want to scan files from multiple users, you'll need to run through this step for each user.

## Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add Google Drive Account**.



2. In the Add a Google Drive Account dialog, click **Sign in to Google Drive**.
3. In the Google page that appears, select the Google Drive account and enter the required Admin user and password, then click **Accept** to allow BlueXP classification to read data from this account.



The Google Drive account is added to the list of working environments.

### Select the type of scanning for user data

Select the type of scanning that BlueXP classification will perform on the user's data.

#### Steps

1. From the *Configuration* page, click the **Configuration** button for the Google Drive account.
1. Enable mapping-only scans, or mapping and classification scans, on the files in the Google Drive account.



To:	Do this:
Enable mapping-only scans on files	Click <b>Map</b>
Enable full scans on files	Click <b>Map &amp; Classify</b>
Disable scanning on files	Click <b>Off</b>

#### Result

BlueXP classification starts scanning the files in the Google Drive account you added, and the results are displayed in the Dashboard and in other locations.

### Remove a Google Drive account from compliance scans

Since only a single user's Google Drive files are part of a single Google Drive account, if you want to stop scanning files from a user's Google Drive account, then you should [delete the Google Drive account from BlueXP classification](#).

### Scan StorageGRID data with BlueXP classification

Complete a few steps to start scanning data within object storage directly with BlueXP classification. BlueXP classification can scan data from any Object Storage service which uses the Simple Storage Service (S3) protocol. This includes NetApp StorageGRID, IBM Cloud Object Store, Linode, B2 Cloud Storage, Amazon S3, and more.



The following information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

Beginning with version 1.31, BlueXP classification is part of the core BlueXP offering. For more information, see [Scan StorageGRID data](#).

## Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

### Review object storage prerequisites

You need to have the endpoint URL to connect with the object storage service.

You need to have the Access Key and Secret Key from the object storage provider so that BlueXP classification can access the buckets.

2

### Deploy the BlueXP classification instance

Deploy [BlueXP classification](#) if there isn't already an instance deployed.

3

### Add the Object Storage Service

Add the object storage service to BlueXP classification.

4

### Select the buckets to scan

Select the buckets that you'd like to scan and BlueXP classification will start scanning them.

## Reviewing object storage requirements

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.

- You need to have the endpoint URL to connect with the object storage service.
- You need to have the Access Key and Secret Key from the object storage provider so that BlueXP classification can access the buckets.

## Deploying the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

If you are scanning data from S3 object storage that is accessible over the internet, you can [deploy BlueXP classification in the cloud](#) or [deploy BlueXP classification in an on-premises location that has internet access](#).

If you are scanning data from S3 object storage that has been installed in a dark site that has no internet access, you need to [deploy BlueXP classification in the same on-premises location that has no internet access](#). This also requires that the BlueXP Connector is deployed in that same on-premises location.

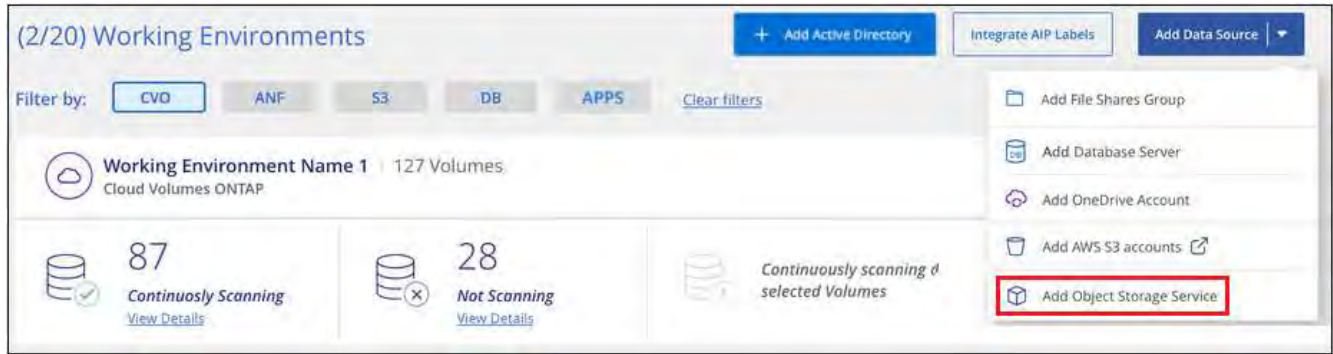
Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

## Adding the object storage service to BlueXP classification

Add the object storage service.

## Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add Object Storage Service**.



2. In the Add Object Storage Service dialog, enter the details for the object storage service and click **Continue**.
  - a. Enter the name you want to use for the Working Environment. This name should reflect the name of the object storage service to which you are connecting.
  - b. Enter the Endpoint URL to access the object storage service.
  - c. Enter the Access Key and Secret Key so that BlueXP classification can access the buckets in the object storage.

The 'Add Object Storage Service' dialog box contains the following text: 'Cloud Data Sense can scan data from any Object Storage service which uses the S3 protocol. This includes NetApp StorageGRID, IBM Object Store, and more. To continue, enter the following information. In the next steps you'll need to select the buckets you want to scan.' Below this are four input fields: 'Name the Working Environment' (value: object\_myIBM), 'Endpoint URL' (value: http://my.endpoint.com), 'Access Key' (value: AIUKDO574NDJG86795), and 'Secret Key' (value: .....). At the bottom are 'Continue' and 'Cancel' buttons.

## Result

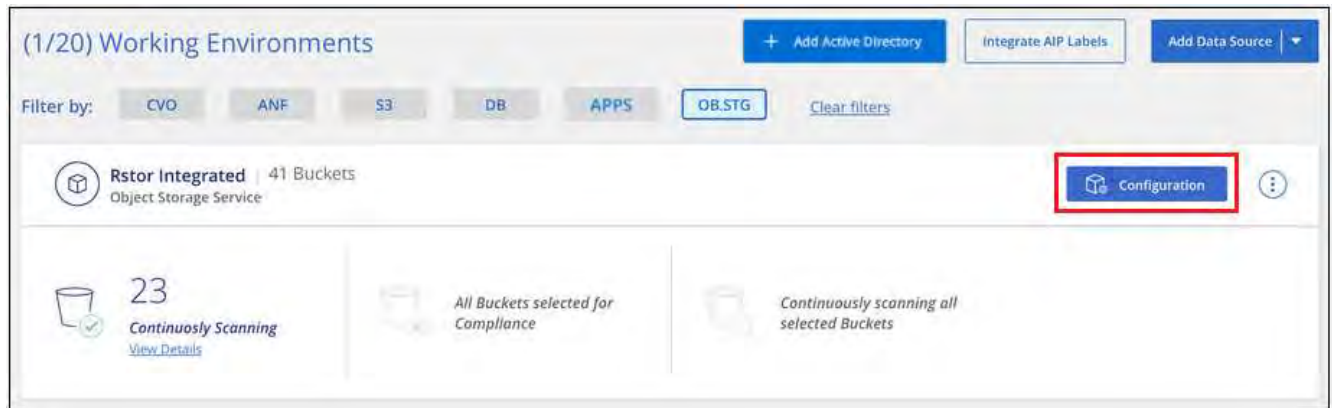
The new Object Storage Service is added to the list of working environments.

## Enabling and disabling compliance scans on object storage buckets

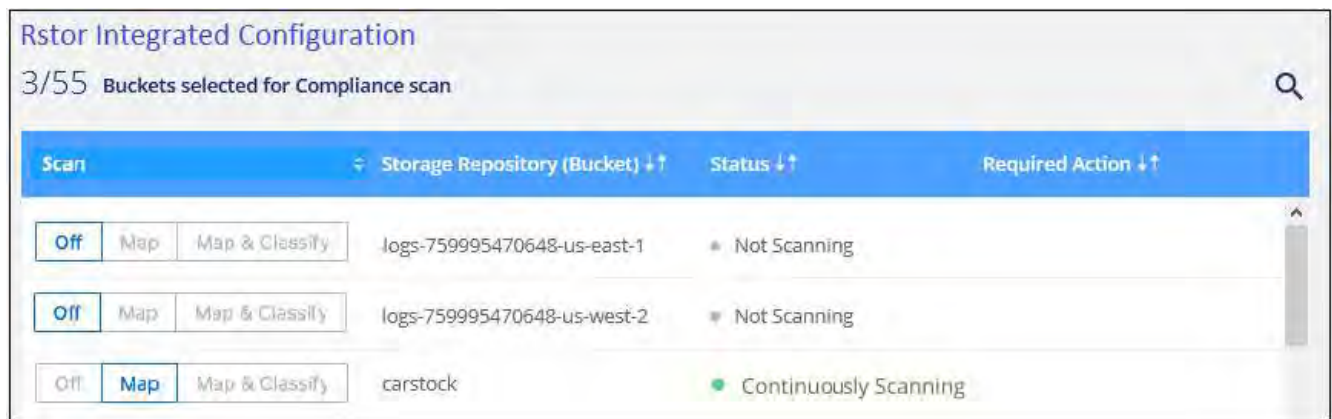
After you enable BlueXP classification on your Object Storage Service, the next step is to configure the buckets that you want to scan. BlueXP classification discovers those buckets and displays them in the working environment you created.

## Steps

1. In the Configuration page, click **Configuration** from the Object Storage Service working environment.



2. Enable mapping-only scans, or mapping and classification scans, on your buckets.



To:	Do this:
Enable mapping-only scans on a bucket	Click <b>Map</b>
Enable full scans on a bucket	Click <b>Map &amp; Classify</b>
Disable scanning on a bucket	Click <b>Off</b>

## Result

BlueXP classification starts scanning the buckets that you enabled. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

# Manage data deprecations

## View governance details about your data using the BlueXP classification Governance dashboard

Gain control of the costs related to the data on your organizations' storage resources. BlueXP classification identifies the amount of stale data, non-business data, duplicate files, and very large files in your systems so you can decide whether you want to remove or tier some files to less expensive object storage.

Additionally, if you're planning to migrate data from on-premises locations to the cloud, you can view the size of the data and whether any of the data contains sensitive information before moving it.

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

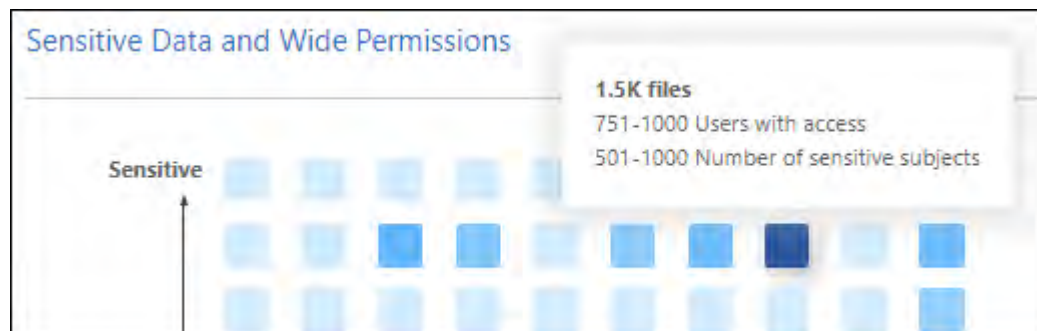
### Data listed by sensitivity and wide permissions on the Governance dashboard

The *Sensitive Data and Wide Permissions* area on the Governance dashboard provides a heatmap of files that contain sensitive data (including both sensitive and sensitive personal data) and that are overly permissive. This can help you to see where you may have some risks with sensitive data.



This applies to BlueXP classification versions 1.30 and earlier.

Files are rated based on the number of users with permission to access the files on the X axis (lowest to highest), and the number of sensitive identifiers within the files on the Y axis (lowest to highest). The blocks represent the number of files that match the items from the X and Y axes. The lighter colored blocks are good; with fewer users able to access the files, and with fewer sensitive identifiers per file. The darker blocks are the items you may want to investigate. For example, the screen below shows the tooltip text for the dark blue block. It shows that you have 1,500 files where 751-1000 users have access, and where there are 501-1000 sensitive identifiers per file.



You can click the block you are interested in to view the filtered results of the affected files in the Investigation page so that you can investigate further.

No data is displayed in this panel if you haven't integrated an identity service with BlueXP classification. [See how to integrate your Active Directory service with BlueXP classification.](#)



This panel supports files in CIFS shares, OneDrive, and SharePoint data sources. There is currently no support for databases, Google Drive, Amazon S3, and generic object storage.

### Classification area on the dashboard showing AIP labels

The *Classification* area on the dashboard provides a list of the most identified Azure Information Protection (AIP) Labels in your scanned data.

If you have subscribed to Azure Information Protection (AIP), you can classify and protect documents and files by applying labels to content. Reviewing the most used AIP labels that are assigned to files enables you to see which labels are most used in your files.

See [AIP Labels](#) for more information.

## Organize your private data with BlueXP classification

BlueXP classification provides many ways for you to manage and organize your private data. This makes it easier to see the data that is most important to you.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

- If you are subscribed to [Azure Information Protection \(AIP\)](#) to classify and protect your files, you can use BlueXP classification to manage AIP labels.
- You can add Tags to files that you want to mark for organization or for some type of follow-up.
- You can assign a BlueXP user to a specific file, or to multiple files, so that person can be responsible for managing the file.
- With the saved search functionality, you can create your own custom search queries so that you can easily see the results by clicking one button.
- You can send email alerts to BlueXP users, or any other email address, when certain critical Policies return results.



The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.

### Should I use tags or labels?

Below is a comparison of BlueXP classification tagging and Azure Information Protection labeling.

Tags	Labels
File tags are an integrated part of BlueXP classification.	Requires that you have subscribed to Azure Information Protection (AIP).
The tag is only kept in the BlueXP classification database - it is not written to the file. It does not change the file, or the file accessed or modified times.	The label is part of the file and when the label changes, the file changes. This change also changes the file accessed and modified times.
You can have multiple tags on a single file.	You can have one label on a single file.
The tag can be used for internal BlueXP classification action, such as copy, move, delete, run a policy, etc.	Other systems that can read the file can see the label - which can be used for additional automation.
Only a single API call is used to see if a file has a tag.	

### Categorize your data using AIP labels

You can manage AIP labels in the files that BlueXP classification is scanning if you have subscribed to [Azure Information Protection \(AIP\)](#). AIP enables you to classify and protect documents and files by applying labels to content. BlueXP classification enables you to view the labels that are already assigned to files, add labels to files, and change labels when a label already exists.

BlueXP classification supports AIP labels within the following file types: .DOC, .DOCX, .PDF, .PPTX, .XLS, .XLSX.





- You can't currently change labels in files larger than 30 MB. For OneDrive, SharePoint, and Google Drive accounts the maximum file size is 4 MB.
- If a file has a label which doesn't exist anymore in AIP, BlueXP classification considers it as a file without a label.
- If you've deployed BlueXP classification in a Government region, or in an on-prem location that has no internet access (also known as a dark site), then the AIP label functionality is unavailable.

### Integrate AIP labels in your project or workspace


Before you can manage AIP labels, you need to integrate the AIP label functionality into BlueXP classification by signing into your existing Azure account. Once enabled, you can manage AIP labels within files for all [data sources](#) in your BlueXP project or workspace.

### Requirements

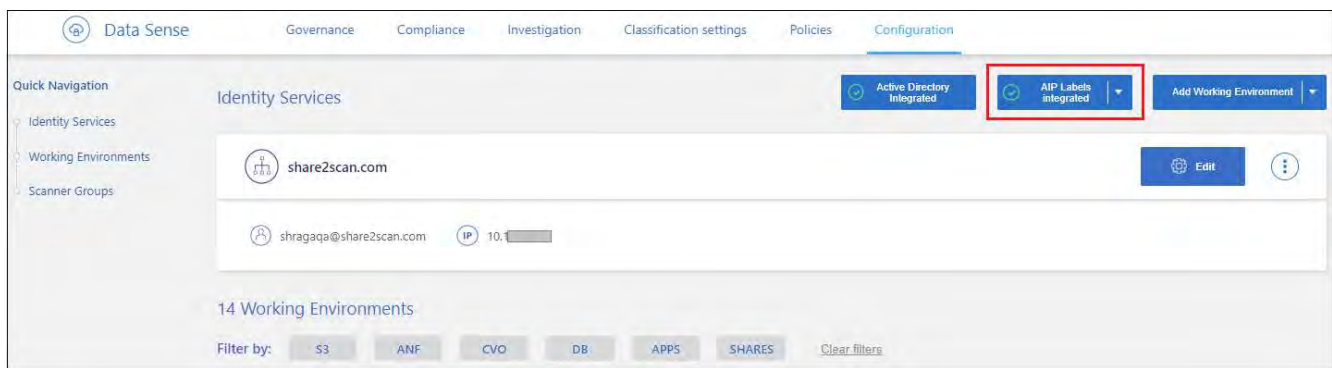
- You must have an account and an Azure Information Protection license.
- You must have the login credentials for the Azure account.
- If you plan to change labels in files that reside in Amazon S3 buckets, ensure that the permission `s3:PutObject` is included in the IAM role. See [setting up the IAM role](#).

### Steps

1. From the BlueXP classification Configuration page, click **Integrate AIP Labels**.



2. In the Integrate AIP Labels dialog, click **Sign in to Azure**.
3. In the Microsoft page that appears, select the account and enter the required credentials.
4. Return to the BlueXP classification tab and you'll see the message "AIP Labels were integrated successfully with the account <account\_name>".
5. Click **Close** and you'll see the text *AIP Labels integrated* at the top of the page.



### Result

You can view and assign AIP labels from the results pane of the Investigation page. You can also assign AIP

labels to files using Policies.

### View AIP labels in your files

You can view the current AIP label that is assigned to a file.

In the Data Investigation results pane, click  for the file to expand the file metadata details.




### Assign AIP labels manually

You can add, change, and remove AIP labels from your files using BlueXP classification.

Follow these steps to assign an AIP label to a single file.

#### Steps

1. In the Data Investigation results pane, click  for the file to expand the file metadata details.





2. Click **Assign a Label to this file** and then select the label.

The label appears in the file metadata.

Follow these steps to assign an AIP label to multiple files. Note that you can assign an AIP label to a maximum of 20 files at a time (one page in the UI).

### Steps


1. In the Data Investigation results pane, select the file, or files, that you want to label.



◦ To select individual files, check the box for each file ( Volume\_1).

◦ To select all files on the current page, check the box in the title row ( File Name).

2. From the button bar, click **Label** and select the AIP label:



The AIP label is added to the metadata for all selected files.

### Remove the AIP integration

If you no longer want the ability to manage AIP labels in files, you can remove the AIP account from the BlueXP classification interface.

Note that no changes are made to the labels you have added using BlueXP classification. The labels that exist in files will stay as they currently exist.

### Steps

1. From the *Configuration* page, click **AIP Labels integrated > Remove Integration**.



2. Click **Remove Integration** from the confirmation dialog.

### Apply tags to manage your scanned files

You can add a tag to files that you want to mark for some type of follow-up. For example, you may have found some duplicate files and you want to delete one of them, but you need to check to see which one should be deleted. You could add a tag of "Check to delete" to the file so you know this file requires some research and some type of future action.

BlueXP classification enables you to view the tags that are assigned to files, add or remove tags from files, and change the name or delete an existing tag.

Note that the tag is not added to the file in the same way as AIP Labels are part of the file metadata. The tag is just seen by BlueXP users using BlueXP classification so you can see if a file needs to be deleted or checked for some type of follow-up.

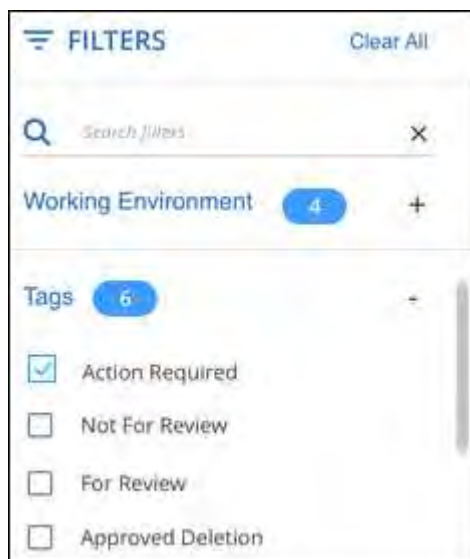


Tags assigned to files in BlueXP classification are not related to the tags you can add to resources, such as volumes or virtual machine instances. BlueXP classification tags are applied at the file level.

### View files that have certain tags applied

You can view all the files that have specific tags assigned.

1. Click the **Investigation** tab from BlueXP classification.
2. In the Data Investigation page, click **Tags** in the Filters pane and then select the required tags.




The Investigation Results pane displays all the files that have those tags assigned.


### Assign tags to files

You can add tags to a single file or to a group of files.

To add a tag to a single file:

#### Steps

1. In the Data Investigation results pane, click  for the file to expand the file metadata details.
2. Click the **Tags** field and the currently assigned tags are displayed.
3. Add the tag or tags:
  - To assign an existing tag, click in the **New Tag...** field and start typing the name of the tag. When the tag you are looking for appears, select it and press **Enter**.
  - To create a new tag and assign it to the file, click in the **New Tag...** field, enter the name of the new tag, and press **Enter**.



The tag appears in the file metadata.

To add a tag to multiple files:

#### Steps

1. In the Data Investigation results pane, select the file, or files, that you want to tag.

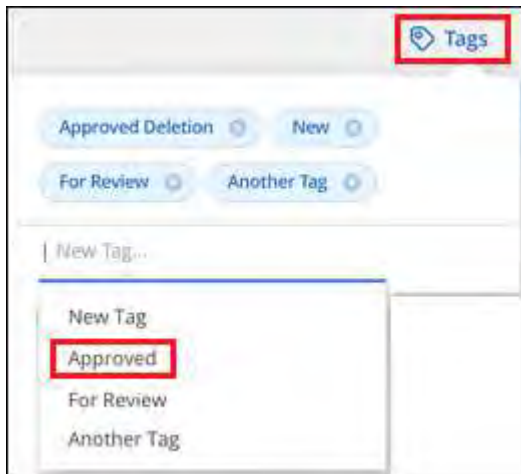
255 items 1.2 GB | 2 Selected 3 MB Tags Assign to Label Copy Move Delete

<input type="checkbox"/>	File Name	Personal	Sensitive Personal	Data Subjects	File Type	
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	16	PDF
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF

- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).
- To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 items on this page selected Select all items in list (63K items)**, click **Select all items in list (xxx items)**.

You can apply tags to a maximum of 100,000 files at a time.

- From the button bar, click **Tags** and the currently assigned tags are displayed.
- Add the tag or tags:
  - To assign an existing tag, click in the **New Tag...** field and start typing the name of the tag. When the tag you are looking for appears, select it and press **Enter**.
  - To create a new tag and assign it to the file, click in the **New Tag...** field, enter the name of the new tag, and press **Enter**.



- Approve adding the tags in the confirmation dialog and the tags are added to the metadata for all selected files.

#### Delete tags from files

You can delete a tag if you don't need to use it anymore.

Just click the **x** for an existing tag.

If you had selected multiple files, the tag is removed from all the files.

### Assign users to manage certain files

You can assign a BlueXP user to a specific file, or to multiple files, so that person can be responsible for any follow-up actions that need to be done on the file. This capability is often used with the feature to add custom Status tags to a file.


For example, you might have a file that contains certain personal data that allows too many users read and write access (open permissions). So you could assign the Status tag "Change permissions" and assign this file to user "Joan Smith" so they can decide how to fix the issue. When they have fixed the issue they could change the Status tag to "Completed".

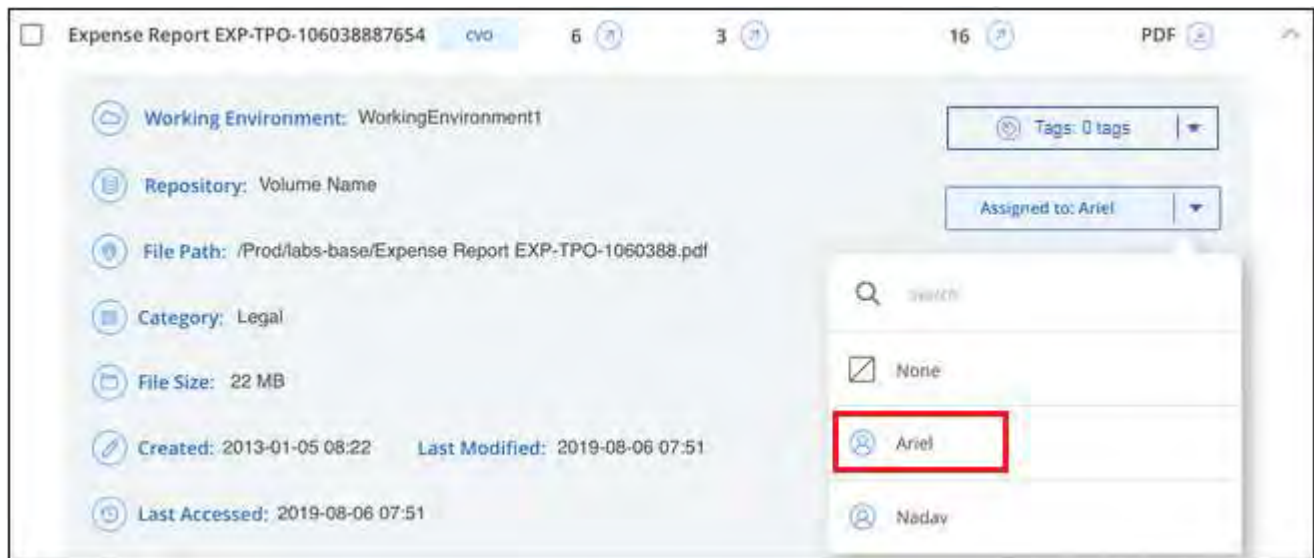
Note that the user name is not added to the file as part of the file metadata - it is just seen by BlueXP users when using BlueXP classification.

A new Filter in the Investigation page enables you to easily view all files that have the same person in the "Assigned To" field.

Follow these steps to assign a user to a single file.

#### Steps

1. In the Data Investigation results pane, click  for the file to expand the file metadata details.
2. Click the **Assigned to** field and select the user name.



The User name appears in the file metadata.

Follow these steps to assign a user to multiple files. Note that you can assign a user to a maximum of 20 files at a time (one page in the UI).

#### Steps

1. In the Data Investigation results pane, select the file, or files, that you want to assign to a user.

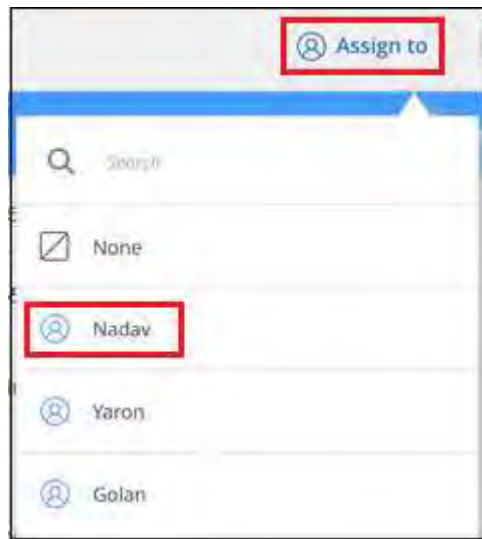
255 items 1.2 GB | 2 Selected 3 MB

Tags Assign to Label Copy Move Delete

<input type="checkbox"/>	File Name	Personal	Sensitive Personal	Data Subjects	File Type	
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	16	PDF
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF

- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).

2. From the button bar, click **Assign to** and select the user name:



The user is added to the metadata for all selected files.

## Manage your private data with BlueXP classification

BlueXP classification provides many ways for you to manage your private data. Some functionality makes it easier to prepare for migrating your data, while other functionality allows you to make changes to the data.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

- You can copy files to a destination NFS share if you want to make a copy of certain data and move it to a different NFS location.
- You can clone an ONTAP volume to a new volume, while including only selected files from the source volume in the new cloned volume. This is useful for situations where you're migrating data and you want to exclude certain files from the original volume.
- You can copy and synchronize files from a source repository to a directory in a specific destination location. This is useful for situations where you're migrating data from one source system to another while there is



still some final activity on the source files.

- You can move source files that BlueXP classification is scanning to any NFS share.
- You can delete files that seem insecure or too risky to leave in your storage system, or that you have identified as duplicate.



- The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.
- Data from Google Drive accounts can't use any of these capabilities at this time.

## Copy source files

You can copy any source files that BlueXP classification is scanning. There are three types of copy operations depending on what you're trying to accomplish:

- **Copy files** from the same, or different, volumes or data sources to a destination NFS share.

This is useful if you want to make a copy of certain data and move it to a different NFS location.

- **Clone an ONTAP volume** to a new volume in the same aggregate, but include only selected files from the source volume in the new cloned volume.

This is useful for situations where you're migrating data and you want to exclude certain files from the original volume. This action uses the [NetApp FlexClone](#) functionality to quickly duplicate the volume and then remove the files that you **didn't** select.

- **Copy and synchronize files** from a single source repository (ONTAP volume, S3 bucket, NFS share, etc.) to a directory in a specific destination (target) location.

This is useful for situations where you're migrating data from one source system to another. After the initial copy, the service syncs any changed data based on the schedule that you set. This action uses the [NetApp BlueXP copy and sync](#) functionality to copy and sync data from a source to a target.

## Copy source files to an NFS share

You can copy source files that BlueXP classification is scanning to any NFS share. The NFS share doesn't need to be integrated with BlueXP classification, you just need to know the name of the NFS share where all selected files will be copied in the format `<host_name>:/<share_path>`.



You can't copy files that reside in databases.

## Requirements

- You must have permissions to copy files. [Learn about user access to compliance information.](#)
- Copying files requires that the destination NFS share allows access from the BlueXP classification instance.
- You can copy between 1 and 100,000 files at a time.

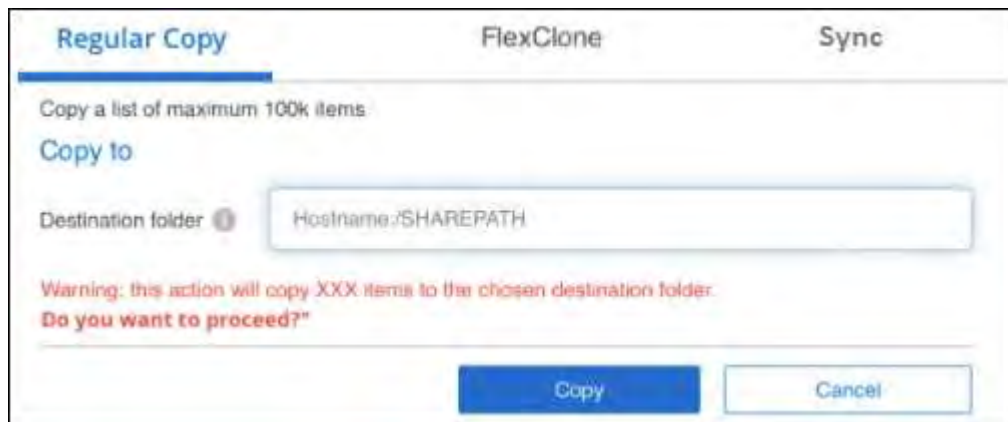
## Steps

1. In the Data Investigation results pane, select the file, or files, that you want to copy, and click **Copy**.



- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).
- To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 items on this page selected** [Select all items in list \(63K items\)](#), click **Select all items in list (xxx items)**.

2. In the *Copy Files* dialog, select the **Regular Copy** tab.



3. Enter the name of the NFS share where all selected files will be copied in the format `<host_name>:/<share_path>`, and click **Copy**.

A dialog appears with the status of the copy operation.

You can view the progress of the copy operation in the [Actions Status](#) pane.

Note that you can also copy an individual file when viewing the metadata details for a file. Just click **Copy File**.





### Clone volume data to a new volume

You can clone an existing ONTAP volume that BlueXP classification is scanning using NetApp *FlexClone* functionality. This allows you to quickly duplicate the volume while including only those files you selected. This is useful if you're migrating data and you want to exclude certain files from the original volume, or if you want to create a copy of a volume for testing.

The new volume is created in the same aggregate as the source volume. Ensure that you have enough space for this new volume in the aggregate before you start this task. Contact your storage administrator if necessary.

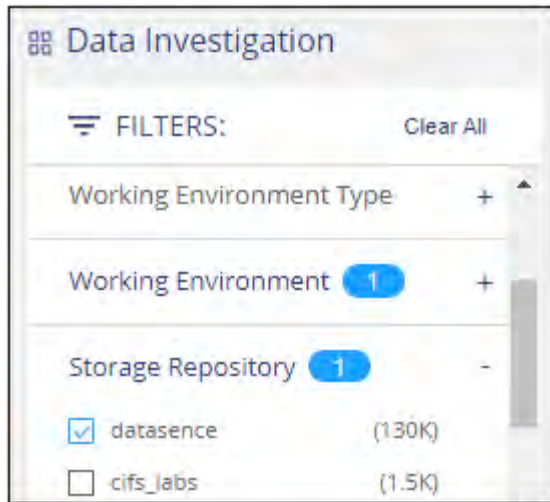
**Note:** FlexGroup volumes can't be cloned because they're not supported by FlexClone.

### Requirements

- You must have permissions to copy files. [Learn about user access to compliance information.](#)
- You must select a minimum of 20 files.
- All selected files must be from the same volume, and the volume must be online.
- The volume must be from a Cloud Volumes ONTAP or on-premises ONTAP system. No other data sources are currently supported.
- The FlexClone license must be installed on the cluster. This license is installed by default on Cloud Volumes ONTAP systems.

### Steps

1. In the Data Investigation pane, create a filter by selecting a single **Working Environment** and a single **Storage Repository** to make sure all the files are from the same ONTAP volume.




Apply any other filters so that you're seeing only the files that you want to clone to the new volume.

2. In the Investigation results pane, select the files that you want to clone and click **Copy**.



- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).
- To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 items on this page selected** [Select all items in list \(63K items\)](#), click **Select all items in list (xxx items)**.

3. In the *Copy Files* dialog, select the **FlexClone** tab. This page shows the total number of files that will be cloned from the volume (the files you selected), and the number of files that are not included/deleted (the files you didn't select) from the cloned volume.



4. Enter the name of the new volume, and click **FlexClone**.

A dialog appears with the status of the clone operation.

## Result

The new, cloned volume is created in the same aggregate as the source volume.

You can view the progress of the clone operation in the [Actions Status pane](#).

If you initially selected **Map all volumes** or **Map & Classify all volumes** when you enabled BlueXP classification for the working environment where the source volume resides, then BlueXP classification will scan the new cloned volume automatically. If you didn't use either of these selections initially, then if you want to scan this new volume, you'll need to [enable scanning on the volume manually](#).

## Copy and synchronize source files to a target system

You can copy source files that BlueXP classification is scanning from any supported unstructured data source to a directory in a specific target destination location ([target locations that are supported by BlueXP copy and sync](#)). After the initial copy, any data changed in the files are synchronized based on the schedule that you configure.

This is useful for situations where you're migrating data from one source system to another. This action uses the [NetApp BlueXP copy and sync](#) functionality to copy and sync data from a source to a target.



You can't copy and sync files that reside in databases, OneDrive accounts, or SharePoint accounts.

## Requirements


- You must have permissions to copy and sync files. [Learn about user access to compliance information](#).
- You must select a minimum of 20 files.
- All selected files must be from the same source repository (ONTAP volume, S3 bucket, NFS or CIFS share, etc.).

- You'll need to activate the BlueXP copy and sync service and configure a minimum of one data broker that can be used to transfer files between the source and target systems. Review the BlueXP copy and sync requirements beginning with the [Quick Start description](#).

Note that the BlueXP copy and sync service has separate service charges for your sync relationships, and will incur resource charges if you deploy the data broker in the cloud.

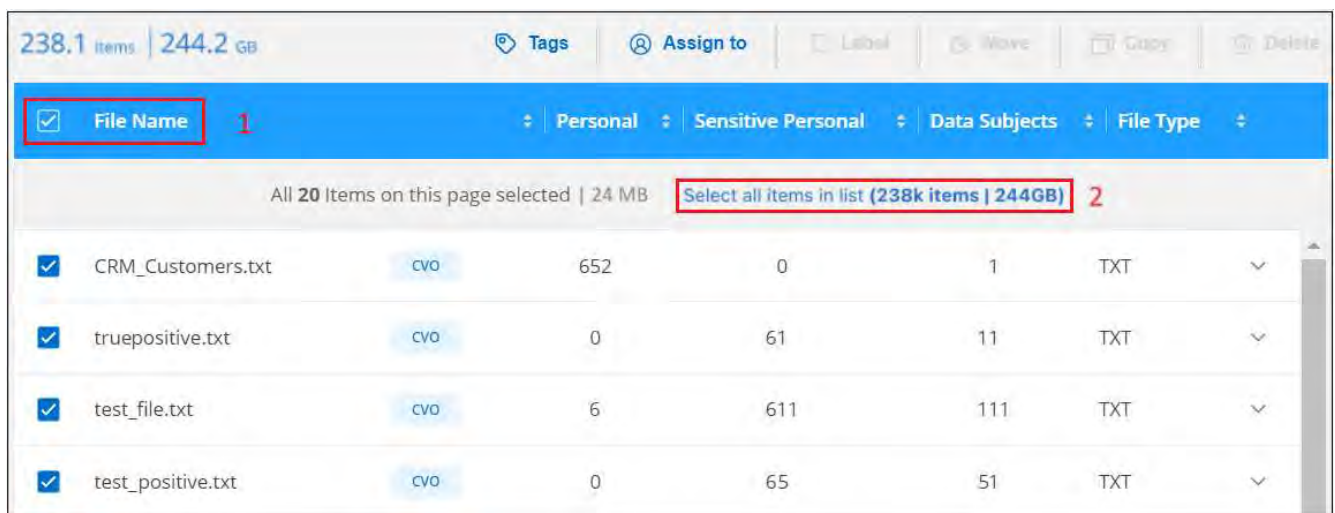
## Steps

- In the Data Investigation pane, create a filter by selecting a single **Working Environment** and a single **Storage Repository** to make sure all the files are from the same repository.

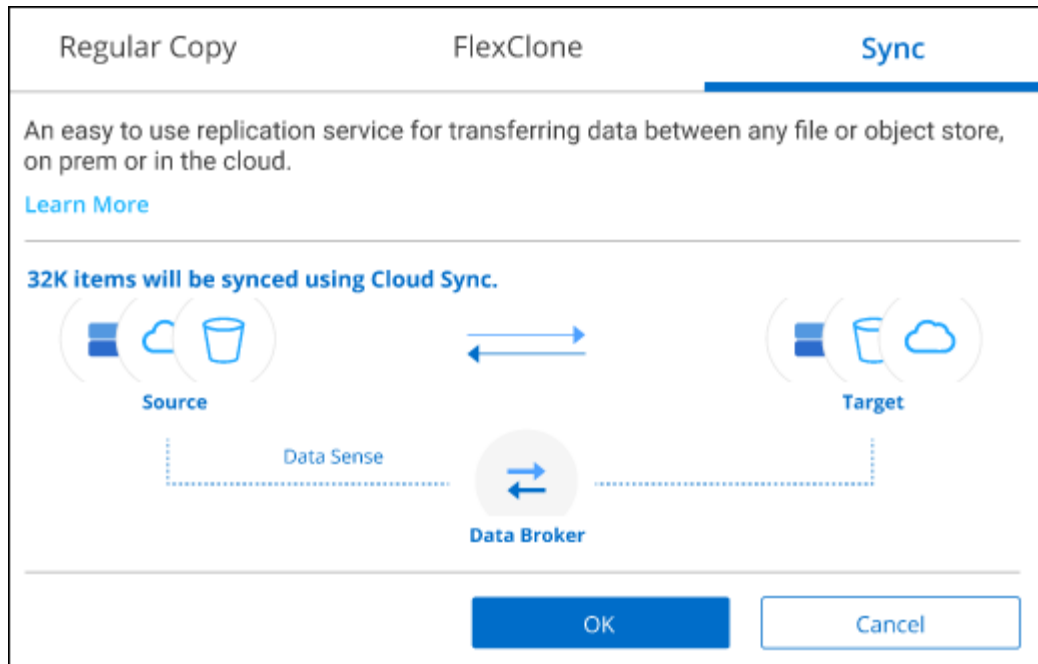


Apply any other filters so that you're seeing only the files that you want to copy and sync to the destination system.

- In the Investigation results pane, select all files on all pages by checking the box in the title row ( **File Name**), then in the pop-up message **All 20 Items on this page selected** [Select all Items in list \(63K Items\)](#) click **Select all items in list (xxx items)**, and then click **Copy**.



- In the *Copy Files* dialog, select the **Sync** tab.



- If you are sure that you want to sync the selected files to a destination location, click **OK**.

The BlueXP copy and sync UI is opened in BlueXP.

You are prompted to define the sync relationship. The Source system is pre-populated based on the repository and files you already selected in BlueXP classification.

- You'll need to select the Target system and then select (or create) the Data Broker you plan to use. Review the BlueXP copy and sync requirements beginning with the [Quick Start description](#).

## Result

The files are copied to the target system and they'll be synchronized based on the schedule you define. If you select a one-time sync then the files are copied and synchronized one time only. If you choose a periodic sync, then the files are synchronized based on the schedule. Note that if the source system adds new files that match the query you created using filters, those *new* files will be copied to the destination and synchronized in the future.

Note that some of the usual BlueXP copy and sync operations are disabled when it is invoked from BlueXP classification:

- You can't use the **Delete Files on Source** or **Delete Files on Target** buttons.
- Running a report is disabled.

## Move source files to an NFS share

You can move source files that BlueXP classification is scanning to any NFS share. The NFS share doesn't need to be integrated with BlueXP classification.

Optionally, you can leave a breadcrumb file in the location of the moved file. A breadcrumb file helps your users understand why a file was moved from its original location. For each moved file, the system creates a breadcrumb file in the source location named `<filename>-breadcrumb-<date>.txt`. You can add text in the dialog box that will be added to the breadcrumb file to indicate the location where the file was moved and the user who moved the file.

Note that the subdirectory structure from the source file is recreated on the destination share when the file is moved so it is easier to understand where the file was moved from. If a file with the same name exists in the destination location, the file will not be moved.



You can't move files that reside in databases.

## Requirements

- You must have permissions to move files. [Learn about user access to compliance information.](#)
- The source files can be located in the following data sources: On-premises ONTAP, Cloud Volumes ONTAP, Azure NetApp Files, File Shares, and SharePoint Online.
- You can move a maximum of 15 million files at a time.
- Only files which are 50 MB or smaller are moved.
- The destination NFS share must allow access from the BlueXP classification instance IP address.

## Steps

1. In the Data Investigation results pane, select the file, or files, that you want to move.

<input type="checkbox"/>	File Name	Personal	Sensitive Personal	Data Subjects	File Type	
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	16	PDF
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF

- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).
- To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 Items on this page selected** [Select all Items in list \(63K Items\)](#), click **Select all items in list (xxx items)**.

2. From the button bar, click **Move**.



### Move Files (63)

The files will be moved to the destination folder you provide and will no longer be available at their current location.

Moving files is supported only to destination folders in NFS Shares. Any NFS Share is supported, no matter where it is hosted, as long as the share's export policy allows access from the data connector instance IP address.

---

The status of this action will appear in the Action Status:

---


**Enter the NFS destination folder path to continue**

---

**Leave breadcrumb**

A breadcrumb file helps your users understand why a file was moved from its original location. For each moved file, the system creates a breadcrumb file in the source location named `<filename>-breadcrumb-<date>.txt`.

---

 Max length should be maximum 400 characters

---

3. In the *Move Files* dialog, enter the name of the NFS share where all selected files will be moved in the format `<host_name>:/<share_path>`.
4. If you want to leave a breadcrumb file, check the *Leave breadcrumb* box. You can enter text in the dialog box to indicate the location where the file was moved and the user who moved the file, and any other information, such as the reason the file was moved.
5. Click **Move Files**.

Note that you can also move an individual file when viewing the metadata details for a file. Just click **Move File**.



### Delete source files

You can permanently remove source files that seem insecure or too risky to leave in your storage system, or that you've identified as a duplicate. This action is permanent and there is no undo or restore.



You can't delete files that reside in databases. All other data sources are supported.

Deleting files requires the following permissions:

- For NFS data - the export policy needs to be defined with write permissions.
- For CIFS data - the CIFS credentials need to have write permissions.
- For S3 data - the IAM role must include the following permission: `s3:DeleteObject`.

### Delete source files manually

#### Requirements

- You must have permissions to delete files. [Learn about user access to compliance information.](#)
- You can delete a maximum of 100,000 files at a time.

#### Steps

1. In the Data Investigation results pane, select the file, or files, that you want to delete.



- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).



To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 items on this page selected** [Select all items in list \(63K items\)](#), click **Select all items in list (xxx items)**.

- From the button bar, click **Delete**.
- Because the delete operation is permanent, you must type "**permanently delete**" in the subsequent *Delete File* dialog and click **Delete File**.

You can view the progress of the delete operation in the [Actions Status](#) pane.

Note that you can also delete an individual file when viewing the metadata details for a file. Just click **Delete file**.



## Add personal data identifiers to your BlueXP classification scans


BlueXP classification provides many ways for you to add a custom list of "personal data" that BlueXP classification will identify in future scans, giving you the full picture about where potentially sensitive data resides in *all* your organizations' files.

 To create a custom classification in version 1.43 and later, see [Create a custom classification](#).

 This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

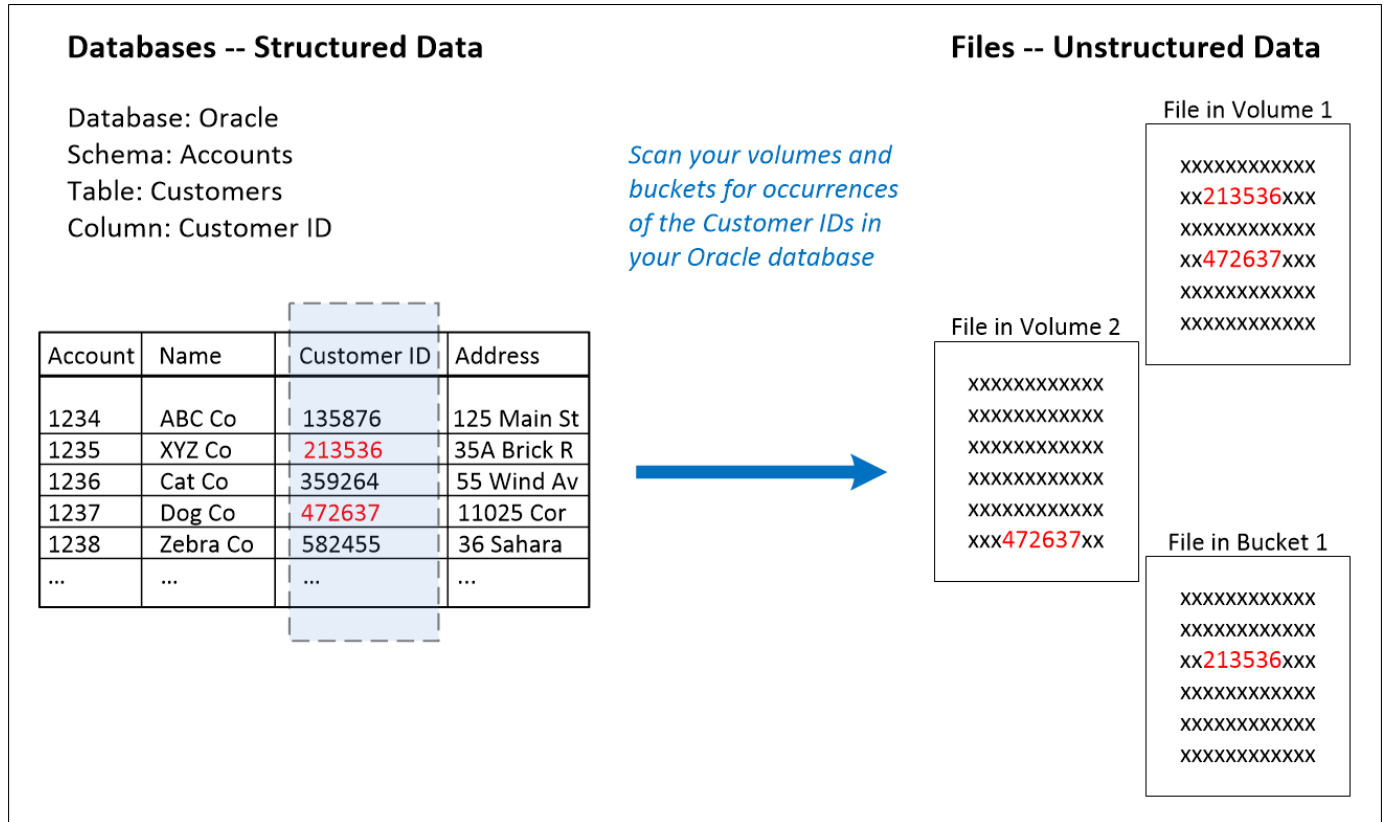
- You can add unique identifiers based on specific columns in databases you are scanning.
- You can add custom keywords from a text file — these words are identified within your data.
- You can add a personal pattern using a regular expression (regex) — the regex is added to the existing predefined patterns.
- You can add custom categories to identify where specific categories of information are found in your data.

All of these mechanisms to add custom scanning criteria are supported in all languages.

 The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.

## Add custom personal data identifiers from your databases

*Data Fusion* allows you to scan your organization's data to identify whether unique identifiers from your databases are found in any of your other data sources. You can choose the additional identifiers that BlueXP classification will look for in its scans by selecting a specific column or columns in a database table. For example, the diagram below shows how data fusion is used to scan your volumes, buckets, and databases for occurrences of all your Customer IDs from your Oracle database.



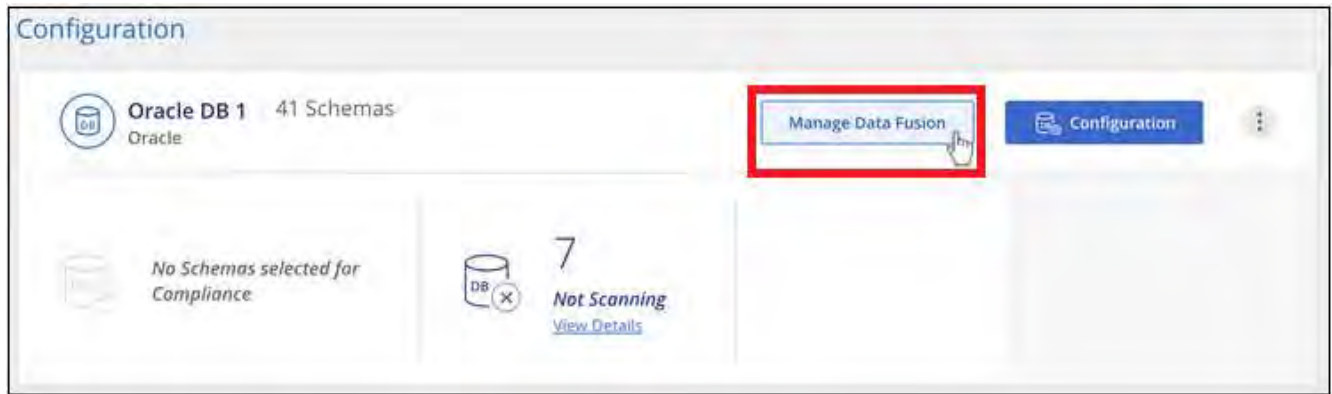
As you can see, two unique Customer IDs have been found in two volumes and in one S3 bucket. Any matches in database tables will also be identified.

Note that since you're scanning your own databases, whatever language your data is stored in will be used to identify data in future BlueXP classification scans.

### Steps

You must have [added at least one database server](#) to BlueXP classification before you can add data fusion sources.

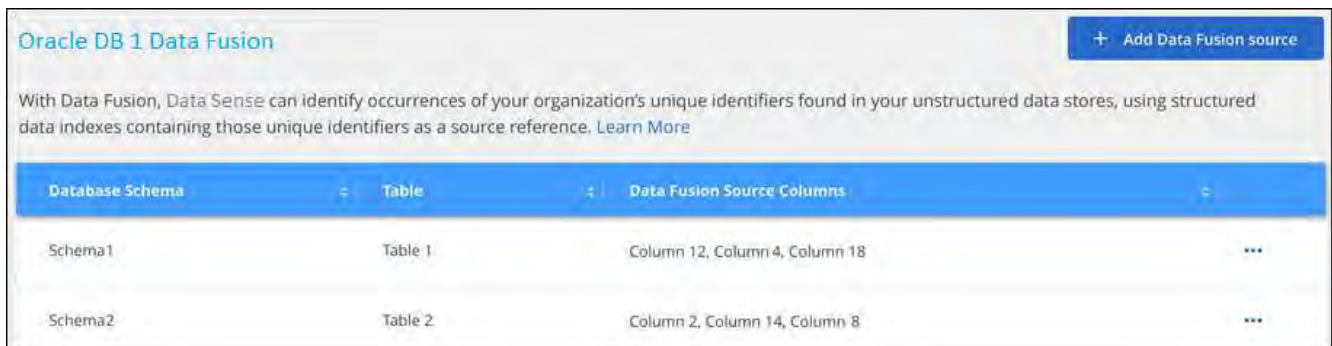
1. In the Configuration page, click **Manage Data Fusion** in the database where the source data resides.



2. Click **Add Data Fusion source** on the next page.
3. In the *Add Data Fusion Source* page:
  - a. Select the Database Schema from the drop-down menu.
  - b. Enter the Table name in that schema.
  - c. Enter the Column, or Columns, that contain the unique identifiers you want to use.

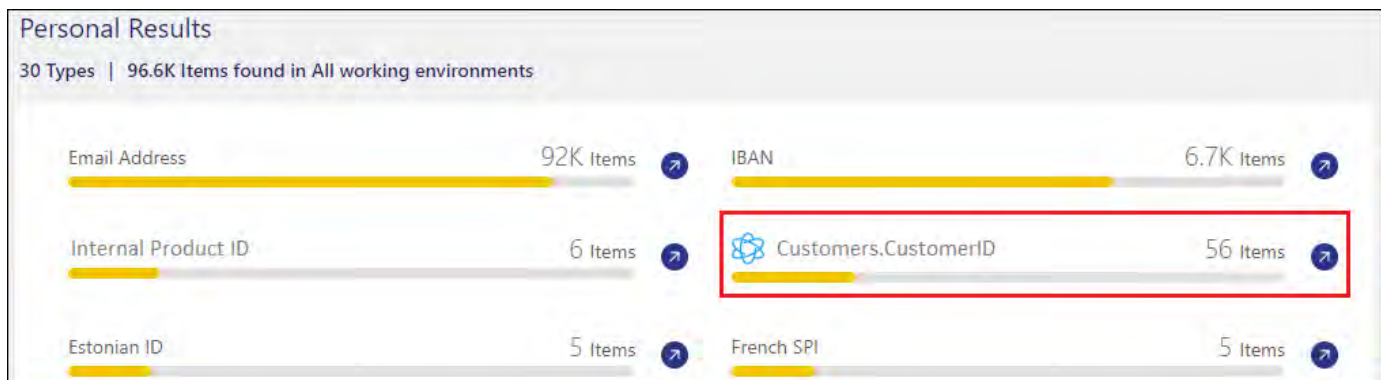
When adding multiple columns, enter each column name, or table view name, on a separate line.

4. Click **Add Data Fusion Source**.



## Results

After the next scan, the results will include this new information in the Compliance Dashboard under the "Personal Results" section, and in the Investigation page in the "Personal Data" filter. The name you used for the classifier appears in the filter list, for example `Customers.CustomerID`.



## Delete a Data Fusion source

If at some point you decide not to scan your files using a certain Data Fusion source, you can select the source row from the Data Fusion inventory page and click **Delete Data Fusion Source**.



## Add custom keywords from a list of words


You can add custom keywords to BlueXP classification so that it will identify where that information is found in your data. You add the keywords just by entering each word you want BlueXP classification to recognize. The keywords are added to the existing predefined keywords that BlueXP classification already uses, and the results will be visible under the personal patterns section.

For example, you may want to see where internal Product Names are mentioned in all of your files to make sure these names are not accessible in locations that are not secure.

After updating the custom keywords, BlueXP classification will restart scanning all data sources. After the scan has completed, the new results will appear in the BlueXP classification Compliance Dashboard under the "Personal Results" section, and in the Investigation page in the "Personal Data" filter.

## Steps

1. From the *Classification settings* tab, click **Add New Classifier** to launch the *Add Custom Classifier* wizard.



2. In the *Select type* page, enter the name of the classifier, provide a brief description, select **Personal identifier**, and then click **Next**.

The name you enter will appear in the BlueXP classification UI as the heading for scanned files that match the classifier requirements, and as the name of the filter in the Investigation page.

You can also check the box to "Mask detected results in the system" so the full result won't appear in the UI. For example, you may want to do this to hide full credit card numbers or similar personal data (the mask would appear in the UI like this: "\*\*\*\* \* 3434").

1 Select type      2 Select tool      3 Create Logic

## Select type

Select the type of classifier that you want to add to the system, and provide the name and description. Data Sense re-scans all your data sources after you add a new classifier. When the scan is complete, all matching results are displayed in the "Classification Settings" dashboard and in other Data Sense pages.

Classifier name

Internal Product Names

Description

Identify internal product names found in all files

**Personal identifier**

The classifier will be added to the system as a new personal identifier. Any matches are considered "personal data", and they are added to the results that are displayed in the Personal Results page and in the Investigation page. [See the list of personal data that Data Sense identifies by default.](#)

Mask detected results in the system

**Category**

The classifier will be added to the system as a new Category. Any matches are added to the results that are displayed in the Categories page and in the Investigation page. [See the list of categories that Data Sense identifies by default.](#)

Previous      **Next**

3. In the *Select Data Analysis Tool* page, select **Custom keywords** as the method you want to use to define the classifier, and then click **Next**.

## Select Data Analysis Tool

Select the tool that will be used to build the list of words, or patterns, that Data Sense will attempt to match in your data sources.

**Custom keywords** ⓘ  
Create a custom personal pattern based on a list of keywords that you provide.

**Custom regular expression** ⓘ  
Create a custom personal pattern based on a regular expression that you define.

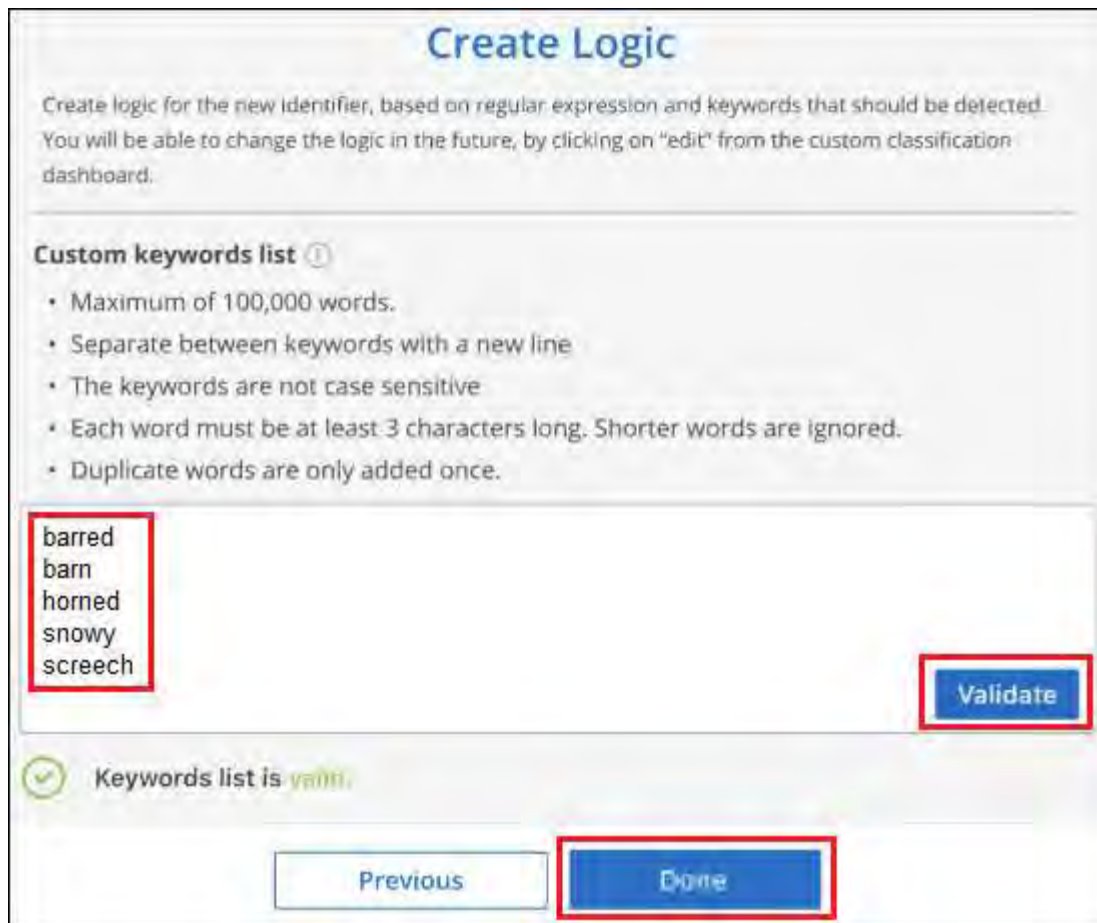
**DB fusion** ⓘ  
Create a custom personal pattern based on the values found in specific columns in your scanned database tables. This allows you to identify whether unique identifiers from your databases are found in any of your other data sources.

[Previous](#) [Next](#)

4. In the *Create Logic* page, enter the keywords you want to recognize - each word on a separate line - and click **Validate**.

The screenshot below shows internal Product Names (different types of owls). The BlueXP classification search for these items is not case sensitive.

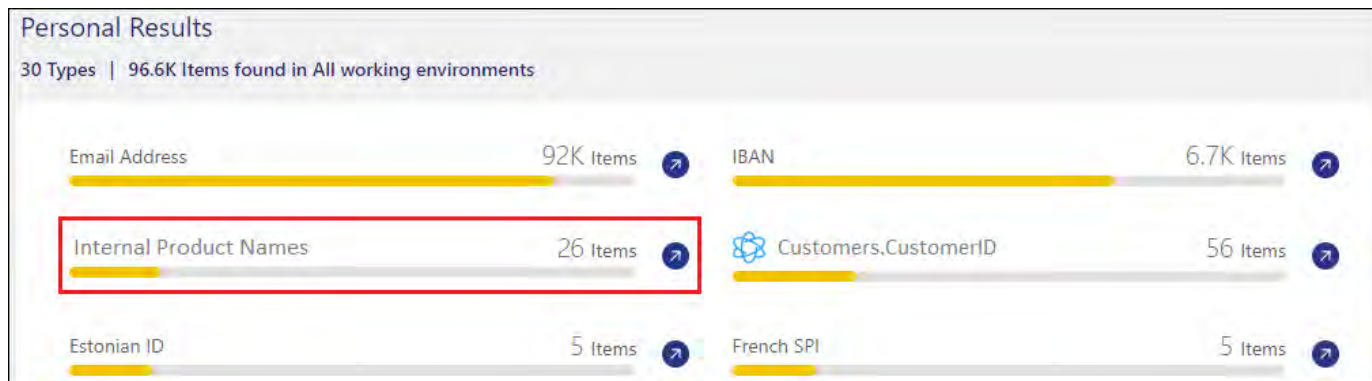




5. Click **Done** and BlueXP classification starts to rescan your data.

### Results

After the scan is complete, the results will include this new information in the Compliance Dashboard under the "Personal Results" section, and in the Investigation page in the "Personal Data" filter.



As you can see, the name of the classifier is used as the name in the Personal Results panel. In this manner you can activate many different groups of keywords and see the results for each group.

### Add custom personal data identifiers using a regex

You can add a personal pattern to identify specific information in your data using a custom regular expression (regex). This allows you to create a new custom regex to identify new personal information elements that don't yet exist in the system. The regex is added to the existing predefined patterns that BlueXP classification

already uses, and the results will be visible under the personal patterns section.

For example, you may want to see where your internal Product IDs are mentioned in all of your files. If the Product ID has a clear structure, for example, it is a 12-digit number that starts with 201, you can use the custom regex feature to search for it in your files. The regular expression for this example is `\b201\d{9}\b`.

After adding the regex, BlueXP classification will restart scanning all data sources. After the scan has completed, the new results will appear in the BlueXP classification Compliance Dashboard under the "Personal Results" section, and in the Investigation page in the "Personal Data" filter.

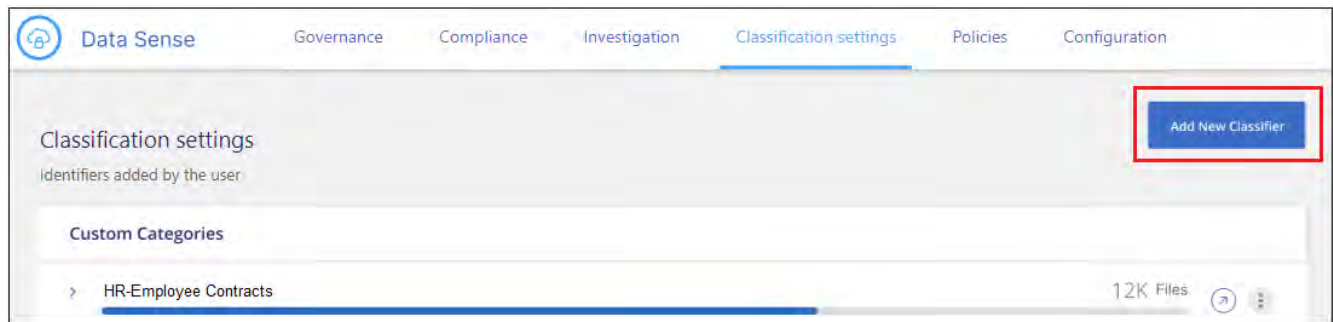
If you need assistance in building the regular expression, refer to [Regular expressions 101](#). Choose **Python** for the Flavor to see the types of results BlueXP classification will match from the regular expression. The [Python Regex Tester page](#) is also useful by displaying a graphical representation of your patterns.



BlueXP classification doesn't support pattern flags when creating a regex. This means you should not use `/`.

## Steps

1. From the *Classification settings* tab, click **Add New Classifier** to launch the *Add Custom Classifier* wizard.



2. In the *Select type* page, enter the name of the classifier, provide a brief description, select **Personal identifier**, and then click **Next**.

The name you enter will appear in the BlueXP classification UI as the heading for scanned files that match the classifier requirements, and as the name of the filter in the Investigation page. You can also check the box to "Mask detected results in the system" so the full result won't appear in the UI. For example, you may want to do this to hide full credit card numbers or similar personal data.



1 Select type      2 Select tool      3 Create Logic

## Select type

Select the type of classifier that you want to add to the system, and provide the name and description. Data Sense re-scans all your data sources after you add a new classifier. When the scan is complete, all matching results are displayed in the "Classification Settings" dashboard and in other Data Sense pages.

Classifier name

Internal Product ID

Description

Identify internal product IDs found in all files

**Personal identifier**

The classifier will be added to the system as a new personal identifier. Any matches are considered "personal data", and they are added to the results that are displayed in the Personal Results page and in the Investigation page. [See the list of personal data that Data Sense identifies by default.](#)

Mask detected results in the system

**Category**

The classifier will be added to the system as a new Category. Any matches are added to the results that are displayed in the Categories page and in the Investigation page. [See the list of categories that Data Sense identifies by default.](#)

Previous      **Next**

3. In the *Select Data Analysis Tool* page, select **Custom regular expression** as the method you want to use to define the classifier, and then click **Next**.

## Select Data Analysis Tool

Select the tool that will be used to build the list of words, or patterns, that Data Sense will attempt to match in your data sources.

**Custom keywords** ⓘ  
Create a custom personal pattern based on a list of keywords that you provide.

**Custom regular expression** ⓘ  
Create a custom personal pattern based on a regular expression that you define.

**DB fusion** ⓘ  
Create a custom personal pattern based on the values found in specific columns in your scanned database tables. This allows you to identify whether unique identifiers from your databases are found in any of your other data sources.

[Previous](#) [Next](#)

4. In the *Create Logic* page, enter the regular expression and any proximity words, and click **Done**.
  - a. You can enter any legal regular expression. Click the **Validate** button to have BlueXP classification verify that the regular expression is valid, and that it is not too broad — meaning it will return too many results.
  - b. Optionally, you can enter some proximity words to help refine the accuracy of the results. These are words that will typically be found within 300 characters of the pattern you are searching for (either before or after the found pattern). Enter each word, or phrase, on a separate line.

## Create Logic

Create logic for the new identifier, based on regular expression and keywords that should be detected.

---

**Regular expression** ⓘ

Add the pattern that should be detected to identify specific information in your data, using a custom regular expression.

Validate

✔ **Success:** Regular expression is valid.

**Proximity words** - To improve the detection accuracy, insert phrases that must appear near by the regular expression's match.

Previous
Done

### Results

The classifier is added and BlueXP classification starts to rescan all your data sources. You are returned to the Custom Classifiers page where you can view the number of files that have matched your new classifier. Results from scanning all of your data sources will take some time depending on the number of files that need to be scanned.

Data Sense
Governance
Compliance
Investigation
Classification settings
Policies
Configuration

Add New Classifier

### Classification settings

Identifiers added by the user

**Custom Categories**

> HR - Employee Contracts
7.5K Files
↻ ⋮

**Personal information**

> Internal Product ID
12K Files
↻ ⋮

### Add custom categories

BlueXP classification takes the data that it scans and divides it into different types of categories. Categories are topics based on artificial intelligence analysis of the content and metadata of each file. [See the list of](#)

[predefined categories](#).

Categories can help you understand what's happening with your data by showing you the types of information that you have. For example, a category like *resumes* or *employee contracts* may include sensitive data. When you investigate the results, you might find that employee contracts are stored in an insecure location. You can then correct that issue.

You can add custom categories to BlueXP classification so you can identify where categories of information that are unique for your data estate are found in your data. You add each category by creating "training" files that contain the categories of data that you want to identify, and then have BlueXP classification scan those files to "learn" through AI so that it can identify that data in your data sources. The categories are added to the existing predefined categories that BlueXP classification already identifies, and the results are visible under the Categories section.

For example, you may want to see where compressed installation files in .gz format are located in your files so that you can remove them, if necessary.

After updating the custom categories, BlueXP classification will restart scanning all data sources. After the scan has completed, the new results will appear in the BlueXP classification Compliance Dashboard under the "Categories" section, and in the Investigation page in the "Category" filter. [See how to view files by categories](#).

### Before you begin

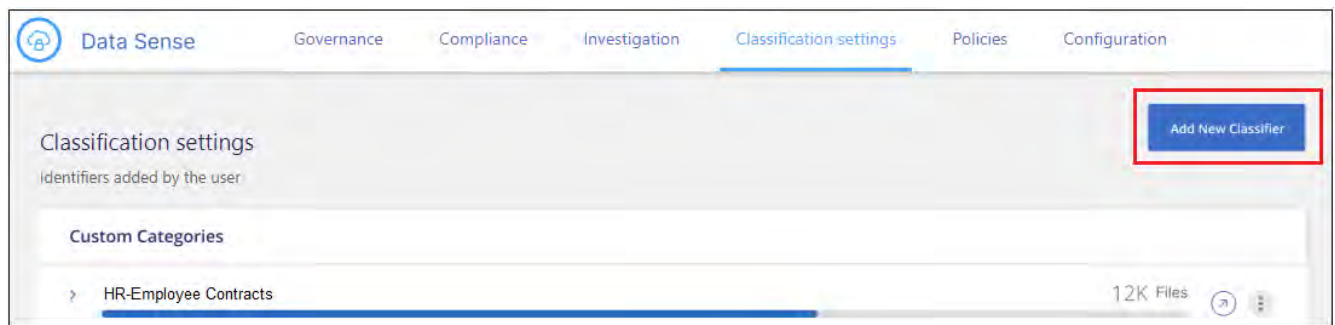
You'll need to create a minimum of 25 training files that contain samples of the categories of data that you want BlueXP classification to recognize. The following file types are supported:

.CSV, .DOC, .DOCX, .GZ, .JSON, .PDF, .PPTX, .RTF, .TXT, .XLS, .XLSX, Docs, Sheets, and Slides

The files must be a minimum of 100 bytes, and they must be located in a folder that is accessible by BlueXP classification.

### Steps

1. From the *Classification settings* tab, click **Add New Classifier** to launch the *Add Custom Classifier* wizard.



2. In the *Select type* page, enter the name of the classifier, provide a brief description, select **Category**, and then click **Next**.

The name you enter will appear in the BlueXP classification UI as the heading for scanned files that match the category of data you are defining, and as the name of the filter in the Investigation page.

1 Select type      2 Select tool      3 Create Logic

## Select type

Select the type of classifier that you want to add to the system, and provide the name and description. Data Sense re-scans all your data sources after you add a new classifier. When the scan is complete, all matching results are displayed in the "Classification Settings" dashboard and in other Data Sense pages.

Classifier name

Compressed installer files

Description

Installation files in .GZ format

Personal identifier

The classifier will be added to the system as a new personal identifier. Any matches are considered "personal data", and they are added to the results that are displayed in the Personal Results page and in the Investigation page. [See the list of personal data that Data Sense identifies by default.](#)

Mask detected results in the system

**Category**

The classifier will be added to the system as a new Category. Any matches are added to the results that are displayed in the Categories page and in the Investigation page. [See the list of categories that Data Sense identifies by default.](#)

Previous      **Next**

3. In the *Create Logic* page, make sure you have the learning files prepared, and then click **Select files**.

## Create Logic

**AI-based similarity training** ⓘ

- Insert NFS folder path
- The folder should contain minimum 25 files and maximum 1000 files that will be used for the AI training.
- Supported file types: pdf, docx, doc, pptx, xls, xlsx, csv, txt, gz, rtf, docs, sheets, slides, json
- The keywords are not case sensitive
- Minimum file size: 100B

Compressed Installer files

**Select Files**

4. Enter the IP address of the volume, and the path where the training files are located, and click **Add**.



**Insert folder path that contains at least 25 files for the training**

Enter the IP address and volume name, along with the path to the location of the training files.

IP:

Training Data - Folder path:

5. Verify that the training files were recognized by BlueXP classification. Click the **x** to remove any training files that do not meet the requirements. Then click **Done**.

### Create Logic

**AI-based similarity training**

- Insert NFS folder path
- The folder should contain minimum 25 files and maximum 1000 files that will be used for the AI training.
- Supported file types: pdf, docx, doc, pptx, xls, xlsx, csv, txt, gz, rtf, docs, sheets, slides, json
- The keywords are not case sensitive.
- Minimum file size: 100B

**Compressed Installer files**

Total uploaded files: 54

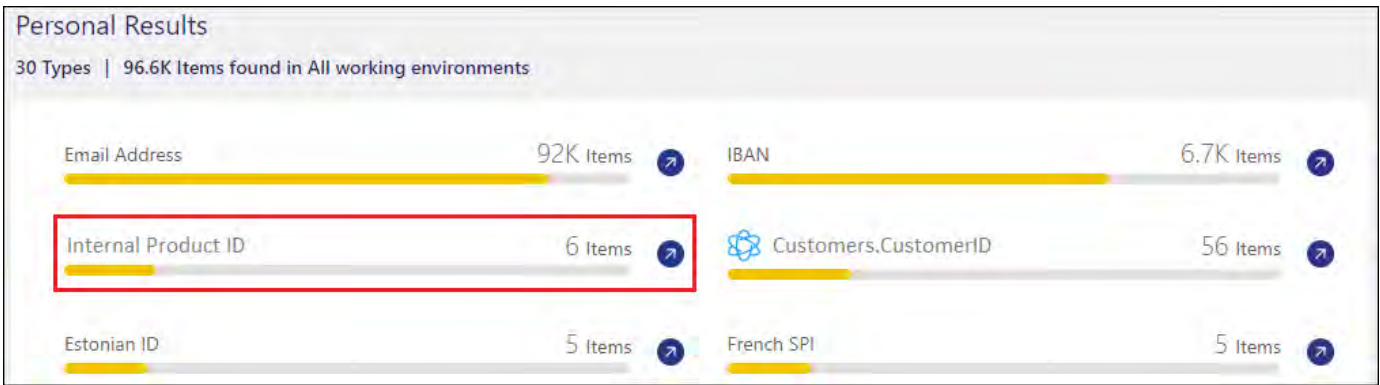
File name	File Size	File Type	Reliability	included in training
File1	56	File type	Sufficient	<input type="checkbox"/>
File2	22	File type	Sufficient	<input type="checkbox"/>
File3	43	File type	Sufficient	<input type="checkbox"/>
File4	11	File type	Sufficient	<input type="checkbox"/>

## Results

The new category is created as defined by the training files and added to BlueXP classification. Then BlueXP classification starts to rescan all your data sources to identify files that fit into this new category. You are returned to the Custom Classifiers page where you can view the number of files that have matched your new category. Results from scanning all of your data sources will take some time depending on the number of files that need to be scanned.

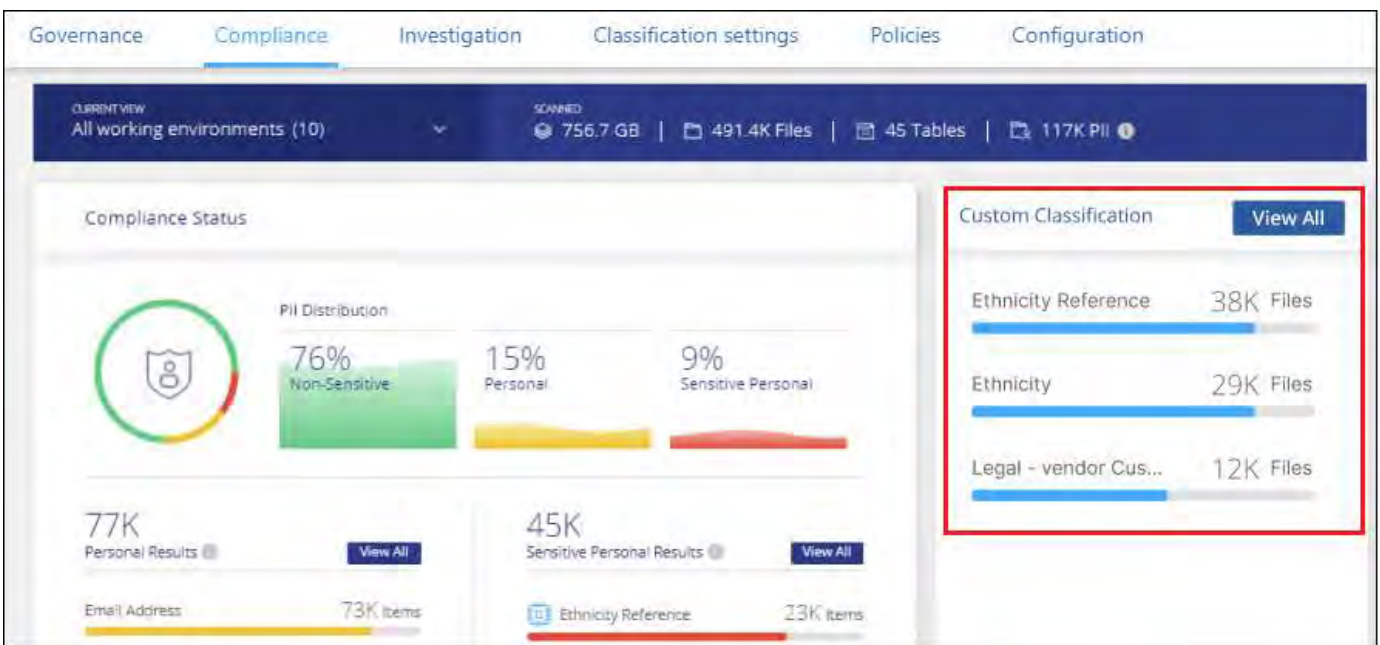
## View results from your custom classifiers

You can view the results from any of your custom classifiers in the Compliance Dashboard and in the Investigation page. For example, this screenshot shows the matched information in the Compliance Dashboard under the "Personal Results" section.



Click the button to see the detailed results in the Investigation page.

Additionally, all of your custom classifier results appear in the Custom Classifiers tab, and the top 6 custom classifier results are displayed in the Compliance Dashboard, as shown below.

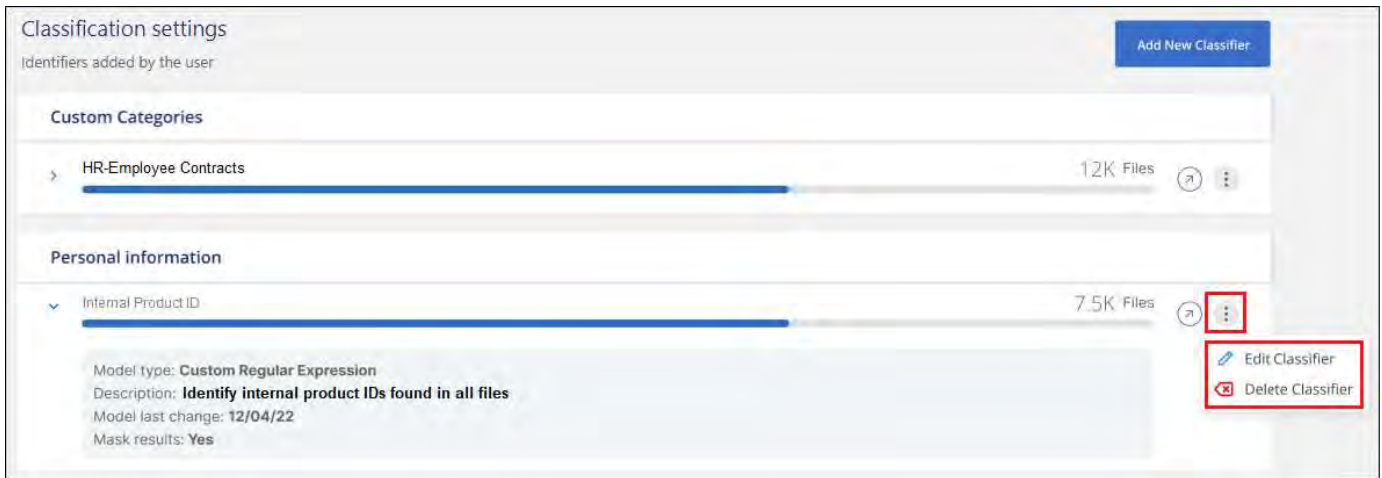


### Manage custom classifiers

You can change any of the custom classifiers that you have created by using the **Edit Classifier** button.

You can't edit Data Fusion classifiers at this time.

And if you decide at some later point that you don't need BlueXP classification to identify the custom patterns that you added, you can use the **Delete Classifier** button to remove each item.



## View the status of your compliance actions in BlueXP classification

When you run an asynchronous action from the Investigation Results pane across many files, for example, moving or deleting 100 files, the process can take some time. You can monitor the status of these actions in the *Action Status* pane so you'll know when it has been applied to all files.

This allows you to see the actions that have completed successfully, those currently in progress, and those that have failed so you can diagnose and fix any problems. Note that short operations that complete quickly, such as moving a single file, do not appear in the Actions Status pane.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

The status can be:

- Success - A BlueXP classification action is finished and all items succeeded.
- Partial Success - A BlueXP classification action is finished and some items failed and some succeeded.
- In Progress - The action is still in progress.
- Queued - The action has not started.
- Canceled - The action has been canceled.
- Failed - The action failed.

Note that you can Cancel any actions that have the "Queued" or "In Progress" status.

### Steps

1.

In the bottom-right of the BlueXP classification UI you can see the **Actions Status** button



2. Click this button and the most recent 20 actions are listed.

You can click the name of an action to view details corresponding to that operation.



## Audit the history of BlueXP classification actions

BlueXP classification logs management activities that have been performed on files from all the working environments and data sources that BlueXP classification is scanning. BlueXP classification also logs the activities when deploying the BlueXP classification instance.

You can view the contents of the BlueXP classification audit log files, or download them, to see what file changes have occurred, and when. For example, you can see what request was issued, the time of the request, and details such as source location in case a file was deleted, or source and destination location in case a file was moved.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

### Log file contents

Each line in the audit log contains information in this format:

```
<full date> | <status> | ds_audit_logger | <module> | 0 | 0 | File <full file path> deleted from device <device path> - <result>
```

- Date and time - full timestamp for the event
- Status - INFO, WARNING
- Action type (delete, copy, move, create saved search, update saved search, rescan files, download JSON report, etc.)
- File name (if the action is relevant to a file)
- Details for the action - what was done: depends on the action
  - Saved search name
  - For move - Source and destination
  - For copy - Source and destination
  - For tag - tag name
  - For assign to - user name
  - For email alert - email address / account

For example, the following lines from the log file show a successful copy operation and a failed copy operation.

```
2022-06-06 15:23:08,910 | INFO | ds_audit_logger | es_scanned_file | 237 | 49 | Copy file /CIFS_share/data/dopl/random_positives.tsv from device 10.31.133.183 (type: SMB_SHARE) to device 10.31.130.133:/export_reports (NFS_SHARE) - SUCCESS
2022-06-06 15:23:08,968 | WARNING | ds_audit_logger | es_scanned_file | 239 | 153 | Copy file /CIFS_share/data/compliance-netapp.tar.gz from device 10.31.133.183 (type: SMB_SHARE) to device 10.31.130.133:/export_reports (NFS_SHARE) - FAILURE
```

## Log file locations

The management audit log files are located on the BlueXP classification machine in:  
`/opt/netapp/audit_logs/`

The installation audit log files are written to `/opt/netapp/install_logs/`

Each log file can be a maximum of 10 MB in size. When that limit is reached, a new log file is started. The log files are named "DataSense\_audit.log", "DataSense\_audit.log.1", "DataSense\_audit.log.2", and so on. A maximum of 100 log files are retained on the system - older log files are deleted automatically after the maximum has been reached.

## Access the log files

You'll need to log into the BlueXP classification system to access the log files. See how to [log in to the BlueXP classification system](#) depending on whether you manually installed the software on a Linux machine or if you deployed the instance in the cloud.

## Reducing the BlueXP classification scan speed

Data scans have a negligible impact on your storage systems and on your data. However, if you are concerned with even a very small impact, you can configure BlueXP classification to perform "slow" scans.

When enabled, slow scanning is used on all data sources - you can't configure slow scanning for a single working environment or data source.



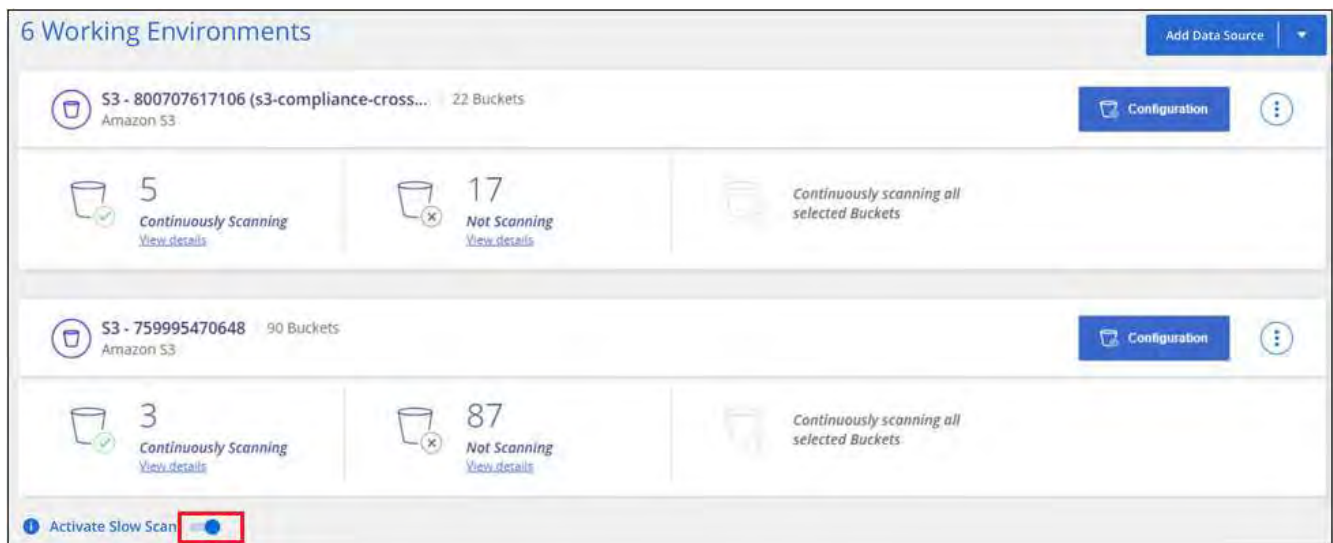
The scan speed can't be reduced when scanning databases.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

## Steps

1. From the bottom of the *Configuration* page, move the slider to the right to activate slow scanning.



The top of the Configuration page indicates that slow scanning is enabled.




2. You can disable slow scanning by clicking **Disable** from this message.

## Remove a OneDrive, SharePoint, or Google Drive account from BlueXP classification

If you no longer want to scan user files from a certain OneDrive account, from a specific SharePoint account, or from a Google Drive account, you can delete the account from the BlueXP classification interface and stop all scans.

### Steps

1. From the *Configuration* page, select the  button in the row for the OneDrive, SharePoint, or Google Drive account, then select **Remove OneDrive Account**, **Remove SharePoint Account**, or **Remove Google Drive account**.



2. Click **Delete Account** from the confirmation dialog.

# Reference

## Supported BlueXP classification instance types

BlueXP classification software must run on a host that meets specific operating system requirements, RAM requirements, software requirements, and so on. When deploying BlueXP classification in the cloud, we recommend that you use a system with the "large" characteristics for full functionality.

You can deploy BlueXP classification on a system with fewer CPUs and less RAM, but there are some limitations when using these less powerful systems. [Learn about these limitations.](#)

In the following tables, if the system marked as "default" is not available in the region where you are installing BlueXP classification, the next system in the table will be deployed.

### AWS instance types

System size	Specs	Instance type
Extra Large	32 CPUs, 128 GB RAM, 1 TiB gp3 SSD	<a href="#">m6i.8xlarge</a> (default)
Large	16 CPUs, 64 GB RAM, 500 GiB SSD	<a href="#">m6i.4xlarge</a> (default) <a href="#">m6a.4xlarge</a> <a href="#">m5a.4xlarge</a> <a href="#">m5.4xlarge</a> <a href="#">m4.4xlarge</a>
Medium	8 CPUs, 32 GB RAM, 200 GiB SSD	<a href="#">m6i.2xlarge</a> (default) <a href="#">m6a.2xlarge</a> <a href="#">m5a.2xlarge</a> <a href="#">m5.2xlarge</a> <a href="#">m4.2xlarge</a>
Small	8 CPUs, 16 GB RAM, 100 GiB SSD	<a href="#">c6a.2xlarge</a> (default) <a href="#">c5a.2xlarge</a> <a href="#">c5.2xlarge</a> <a href="#">c4.2xlarge</a>

### Azure instance types

System size	Specs	Instance type
Extra Large	32 CPUs, 128 GB RAM, OS Disk (2,048 GiB, min 250 MB/s throughput), and Data Disk (1 TiB SSD, min 750 MB/s throughput)	<a href="#">Standard_D32_v3</a> (default)
Large	16 CPUs, 64 GB RAM, 500 GiB SSD	<a href="#">Standard_D16s_v3</a> (default)

## GCP instance types

System size	Specs	Instance type
Large	16 CPUs, 64 GB RAM, 500 GiB SSD	n2-standard-16 (default) n2d-standard-16 n1-standard-16

## Metadata collected from data sources in BlueXP classification

BlueXP classification collects certain metadata when performing classification scans on the data from your data sources and working environments. BlueXP classification can access most of the metadata we need to classify your data, but there are some sources where we are unable to access the data we need.

	Metadata	CIFS	NFS
<b>Time stamps</b>	<i>Creation time</i>	Available	Not available (Unsupported in Linux)
	<i>Last access time</i>	Available	Available
	<i>Last modify time</i>	Available	Available
<b>Permissions</b>	<i>Open permissions</i>	If "EVERYONE" group has access to the file, it is considered "Open to organization"	If "Others" has access to the file, it is considered "Open to organization"
	<i>Users/group access</i>	Users and group information is taken from LDAP	Not available (NFS users are usually managed locally on the server, therefore, the same individual can have a different UID in each server)




- BlueXP classification does not extract the "last accessed time" from the database data sources.
- Older versions of the Windows OS (for example, Windows 7 and Windows 8) disable the collection of the "last accessed time" attribute by default because it can impact system performance. When this attribute is not collected, BlueXP classification analytics that are based on "last accessed time" will be impacted. You can enable the collection of the last access time on these older Windows systems if needed.

### Last access time timestamp

When BlueXP classification extracts data from file shares, the operating system considers it as accessing the data and it changes the "last access time" accordingly. After scanning, BlueXP classification attempts to revert the last access time to the original timestamp. If BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, the system can't revert the last access time to the original timestamp.

ONTAP volumes configured with SnapLock have read-only permissions and also can't revert the last access time to the original timestamp.

By default, if BlueXP classification doesn't have these permissions, the system won't scan those files in your volumes because BlueXP classification can't revert the "last access time" to the original timestamp. However, if you don't care if the last access time is reset to the original time in your files, you can click the **Scan when missing "write attributes" permissions** switch at the bottom of the Configuration page so that BlueXP classification will scan the volumes regardless of permissions.



This functionality is applicable to On-premises ONTAP systems, Cloud Volumes ONTAP, Azure NetApp Files, FSx for ONTAP, and third-party file shares.

Note that there is a filter in the Investigation page called *Scan Analysis Event* that enables you to display either the files that were not classified because BlueXP classification couldn't revert the last accessed time, or the files that were classified even though BlueXP classification couldn't revert the last access time.

The filter selections are:

- "Not classified — Cannot revert last access time" - This shows the files that were not classified due to missing write permissions.
- "Classified and updated last access time" - This shows the files that were classified and BlueXP classification was unable to reset the last access time back to the original date. This filter is relevant only for environments where you turned **Scan when missing "write attributes" permissions** ON.

If needed, you can export these results to a report so you can see which files are, or aren't, being scanned because of permissions. [Learn more about the Data Investigation Report.](#)

## Log in to the BlueXP classification system

At times you may need to log into the BlueXP classification system so you can access log files or edit configuration files.

When BlueXP classification is installed on a Linux machine on your premises, or on a Linux machine you deployed in the cloud, you can access the configuration file and script directly.

When BlueXP classification is deployed in the cloud, you need to SSH to the BlueXP classification instance. You SSH to the system by entering the user and password, or by using the SSH key you provided during the BlueXP Connector installation. The SSH command is:

```
ssh -i <path_to_the_ssh_key> <machine_user>@<datasense_ip>
```

- <path\_to\_the\_ssh\_key> = location of ssh authentication keys
- <machine\_user>:
  - For AWS: use the <ec2-user>
  - For Azure: use the user created for the BlueXP instance
  - For GCP: use the user created for the BlueXP instance
- <datasense\_ip> = IP address of the virtual machine instance

Note that you'll need to modify the security group inbound rules to access the system in the cloud. For details, see:

- [Security group rules in AWS](#)
- [Security group rules in Azure](#)
- [Firewall rules in Google Cloud](#)

## BlueXP classification APIs

The BlueXP classification capabilities that are available through the web UI are also available through the Swagger API.

There are four categories defined within BlueXP classification that correspond to the tabs in the UI:

- Investigation
- Compliance
- Governance
- Configuration

The APIs in the Swagger documentation allow you to search, aggregate data, track your scans, and create actions like copy, move, and more.

### Overview

The API enables you to perform the following functions:

- Export information
  - Everything that is available in the UI can be exported via the API (with the exception of reports)
  - Data is exported in a JSON format (easy to parse and push to 3rd party applications, like Splunk)
- Create queries using "AND" and "OR" statements, include and exclude information, and more.

For example, you can locate files *without* specific Personal Identifiable Information (PII) (functionality not available in the UI). You can also exclude specific fields for the export operation.

- Perform actions
  - Update CIFS credentials

- View and cancel actions
- Re-scan directories
- Export data

The API is secure and it uses the same authentication method as the UI. You can find information on the authentication in: [https://docs.netapp.com/us-en/bluexp-automation/platform/get\\_identifiers.html](https://docs.netapp.com/us-en/bluexp-automation/platform/get_identifiers.html)

## Accessing the Swagger API reference

To get into Swagger you'll need the IP address of the your BlueXP classification instance. In the case of a cloud deployment you'll use the public IP address. Then you'll need to get into this endpoint:

`https://<classification_ip>/documentation`

## Example using the APIs

The following example shows an API call to copy files.

### API Request

You'll initially need to get all the relevant fields and options for a working environment to view all of the filters in the investigation tab.

```
curl -X GET "http://{classification_ip}/api/{classification_version}
/search/options?data_mode=ALL_EXTRACTABLE" -H "accept: application/json"
-H "Authorization: Bearer eyJhbGciOiJSUzI1NiIsInR..... " -H "x-agent-id:
hOXsZNvnA5LsthwMILtjL9xZFYBQxAwMclients"
```

### Response

```
{
  "options": [
    {
      "active_directory_affected": false,
      "data_mode": "ALL_SCANNED",
      "field": "string",
      "is_rulable": true,
      "name": "string",
      "operators": [
        "EQUALS"
      ],
      "optional_values": [
        {}
      ],
      "secondary": {},
      "server_data": false,
      "type": "TEXT"
    }
  ]
}
```



```

]
}
{
  "options": [
    {
      "active_directory_affected": false,
      "data_mode": "ALL_EXTRACTABLE",
      "field": "POLICIES",
      "name": "Policies",
      "operators": [
        "IN",
        "NOT_IN"
      ],
      "server_data": true,
      "type": "SELECT"
    },
    {
      "active_directory_affected": false,
      "data_mode": "ALL_EXTRACTABLE",
      "field": "EXTRACTION_STATUS_RANGE",
      "name": "Scan Analysis Status",
      "operators": [
        "IN"
      ],
      "server_data": true,
      "type": "SELECT"
    },
    {
      "active_directory_affected": false,
      "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
      "field": "SCAN_ANALYSIS_ERROR",
      "name": "Scan Analysis Event",
      "operators": [
        "IN"
      ],
      "server_data": true,
      "type": "SELECT"
    },
    {
      "active_directory_affected": false,
      "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
      "field": "PUBLIC_ACCESS",
      "name": "Open Permissions",
      "operators": [
        "IN",
        "NOT_IN"
      ]
    }
  ]
}

```

```

    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": true,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "USERS_PERMISSIONS_COUNT_RANGE",
    "name": "Number of Users with Access",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": true,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "USER_GROUP_PERMISSIONS",
    "name": "User / Group Permissions",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "FILE_OWNER",
    "name": "File Owner",
    "operators": [
      "EQUALS",
      "CONTAINS"
    ],
    "server_data": true,
    "type": "TEXT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "ENVIRONMENT_TYPE",
    "name": "Working Environment Type",
    "operators": [
      "IN",

```

```

        "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
},
{
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "ENVIRONMENT",
    "name": "Working Environment",
    "operators": [
        "IN",
        "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
},
{
    "active_directory_affected": false,
    "data_mode": "ALL_SCANNED",
    "field": "SCAN_TASK",
    "name": "Storage Repository",
    "operators": [
        "IN",
        "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
},
{
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "FILE_PATH",
    "name": "File / Directory Path",
    "operators": [
        "MULTI_CONTAINS",
        "MULTI_EXCLUDE"
    ],
    "server_data": true,
    "type": "MULTI_TEXT"
},
{
    "active_directory_affected": false,
    "data_mode": "ALL_DASHBOARD_EXTRACTABLE",
    "field": "CATEGORY",
    "name": "Category",

```

```

    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "PATTERN_SENSITIVITY_LEVEL",
    "name": "Sensitivity Level",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "NUMBER_OF_IDENTIFIERS",
    "name": "Number of identifiers",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "PATTERN_PERSONAL",
    "name": "Personal Data",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "PATTERN_SENSITIVE",

```

```

    "name": "Sensitive Personal Data",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "DATA_SUBJECT",
    "name": "Data Subject",
    "operators": [
      "EQUALS",
      "CONTAINS"
    ],
    "server_data": true,
    "type": "TEXT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "DIRECTORIES",
    "field": "DIRECTORY_TYPE",
    "name": "Directory Type",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "FILE_TYPE",
    "name": "File Type",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,

```

```

    "data_mode": "ALL_EXTRACTABLE",
    "field": "FILE_SIZE_RANGE",
    "name": "File Size",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "FILE_CREATION_RANGE_RETENTION",
    "name": "Created Time",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "DISCOVERED_TIME_RANGE",
    "name": "Discovered Time",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "FILE_LAST_MODIFICATION_RETENTION",
    "name": "Last Modified",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",

```

```

    "field": "FILE_LAST_ACCESS_RANGE_RETENTION",
    "name": "Last Accessed",
    "operators": [
        "IN"
    ],
    "server_data": true,
    "type": "SELECT"
},
{
    "active_directory_affected": false,
    "data_mode": "FILES",
    "field": "IS_DUPLICATE",
    "name": "Duplicates",
    "operators": [
        "EQUALS",
        "IN"
    ],
    "server_data": true,
    "type": "SELECT"
},
{
    "active_directory_affected": false,
    "data_mode": "FILES",
    "field": "FILE_HASH",
    "name": "File Hash",
    "operators": [
        "EQUALS",
        "IN"
    ],
    "server_data": true,
    "type": "TEXT"
},
{
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "USER_DEFINED_STATUS",
    "name": "Tags",
    "operators": [
        "IN",
        "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
},
{
    "active_directory_affected": false,

```

```

    "data_mode": "ALL_EXTRACTABLE",
    "field": "ASSIGNED_TO",
    "name": "Assigned to",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  }
]
}

```

We will use that response in our request parameters to filter the desired files we want to copy.

You can apply an action on multiple items. Supported action types include: move, delete, copy, assign to, FlexClone, export data, rescan, and label.

We will create the copy action:

#### API Request

This next API is that action API and it allows you to create multiple actions.

```

curl -X POST "http://
{classification_ip}/api/{classification_version}/actions" -H "accept:
application/json" -H "Authorization: Bearer eyJhbGciOiJSUzI1NiIsInR..... "
-H "x-agent-id: hOXsZNvnA5LsthwMILtjL9xZFYBQxAwMclients " -H "Content-
Type: application/json" -d "{ \"action_type\": \"COPY\", \"data_mode\":
\"FILES\", \"policy_id\": 0, \"request_params\": { destination_nfs_path:
\"{ontap_ip}/{share_name} \" },
\"requested_query\":{\"condition\":\"AND\",\"rules\":[{\"field\":\"ENVIRONMENT_TYPE
\",\"operator\":\"IN\",\"value\":[\"ONPREM\"]},{\"field\":\"CATEGORY\",\"operator\":\"IN\",
\"value\":[\"21\"]}]}}}"

```

#### Response

The response will return the action object, so you can use the get and delete APIs to get status about the action, or to cancel it.



```
{
  "action_type": "COPY",
  "creation_time": "2023-08-08T12:37:21.705Z",
  "data_mode": "FILES",
  "end_time": "2023-08-08T12:37:21.705Z",
  "estimated_time_to_complete": 0,
  "id": 0,
  "policy_id": 0,
  "policy_name": "string",
  "priority": 0,
  "request_params": {},
  "requested_query": {},
  "result": {
    "error_message": "string",
    "failed": 0,
    "in_progress": 0,
    "succeeded": 0,
    "total": 0
  },
  "start_time": "2023-08-08T12:37:21.705Z",
  "status": "QUEUED",
  "title": "string",
  "user_id": "string"
}
```

# Knowledge and support

## Register for BlueXP 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](#)
- [Google Cloud NetApp Volumes](#)

## Support registration overview

There are two forms of registration to activate support entitlement:

- Registering your BlueXP account serial number (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 BlueXP for NetApp support

To register for support and activate support entitlement, one user in your BlueXP organization (or 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 BlueXP organization 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 organization is not registered for support. As long as one user in the organization has followed these steps, then your organization 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.



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 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 organization 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 organization is different than the NSS account that is associated with a BlueXP user login.

These NSS credentials are associated with your specific BlueXP organization ID. Users who belong to the BlueXP organization 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 for BlueXP classification

NetApp provides support for BlueXP and its cloud services in a variety of ways. Extensive free self-support options are available 24/7, 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](#)
- [Google Cloud NetApp Volumes](#)

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 organization (or account), project (or 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 Working Environment


Select 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)  Upload 

No files selected  

### 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.

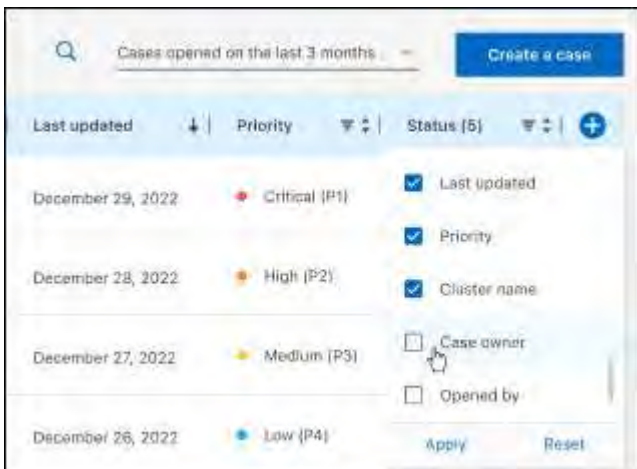
3. Optionally modify the information that displays in the table:
  - Under **Organization's cases**, select **View** to view all cases associated with your company.
  - Modify the date range by choosing an exact date range or by choosing a different time frame.



- Filter the contents of the columns.



- Change the columns that appear in the table by selecting  and then choosing the columns that you'd like to display.




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**.



# Frequently asked questions about BlueXP classification

This FAQ can help if you're just looking for a quick answer to a question.

## BlueXP classification service

The following questions provide a general understanding of BlueXP classification.

### How does BlueXP classification work?

BlueXP classification deploys another layer of AI alongside your BlueXP system and storage systems. It then scans the data on volumes, buckets, databases, and other storage accounts and indexes the data insights that are found. BlueXP classification leverages both artificial intelligence and natural language processing, as opposed to alternative solutions that are commonly built around regular expressions and pattern matching.

BlueXP classification uses AI to provide contextual understanding of data for accurate detection and classification. It is driven by AI because it is designed for modern data types and scale. It also understands data context in order to provide strong, accurate, discovery and classification.

[Learn more about how BlueXP classification works.](#)

### Does BlueXP classification have a REST API, and does it work with third-party tools?

Yes, BlueXP classification has a REST API for the supported features in the BlueXP classification version that is part of the BlueXP core platform. See [API documentation](#).

### Is BlueXP classification available through the cloud marketplaces?

BlueXP classification is part of the BlueXP core features, so you do not need to use the marketplaces for this service .

## BlueXP classification scanning and analytics

The following questions relate to BlueXP classification scanning performance and the analytics.

### How often does BlueXP classification scan my data?

While the initial scan of your data might take a little bit of time, subsequent scans only inspect the incremental changes, which reduces system scan times. BlueXP classification scans your data continuously in a round-robin fashion, six repositories at a time, so that all changed data is classified very quickly.

[Learn how scans work.](#)

BlueXP classification scans databases only once per day; databases are not continuously scanned like other data sources.

Data scans have a negligible impact on your storage systems and on your data.

## Does scan performance vary?

Scan performance can vary based on the network bandwidth and the average file size in your environment. It can also depend on the size characteristics of the host system (either in the cloud or on-premises). See [The BlueXP classification instance](#) and [Deploying BlueXP classification](#) for more information.

When initially adding new data sources, you can also choose to perform only a "mapping" (Mapping only) scan instead of a full "classification" (Map & Classify) scan. Mapping can be done on your data sources very quickly because it does not access files to see the data inside. [See the difference between a mapping and classification scan.](#)

## Can I search my data using BlueXP classification?

BlueXP classification offers extensive search capabilities that make it easy to search for a specific file or piece of data across all connected sources. BlueXP classification empowers users to search deeper than just what the metadata reflects. It is a language-agnostic service that can also read the files and analyze a multitude of sensitive data types, such as names and IDs. For example, users can search across both structured and unstructured data stores to find data that may have leaked from databases to user files, in violation of corporate policy. Searches can be saved for later, and policies can be created to search and take action on the results at a set frequency.

Once the files of interest are found, characteristics can be listed, including tags, working environment account, bucket, file path, category (from classification), file size, last modified, permission status, duplicates, sensitivity level, personal data, sensitive data types within the file, owner, file type, file size, created time, file hash, whether the data was assigned to someone seeking their attention, and more. Filters can be applied to screen out characteristics that are not pertinent.

BlueXP classification also has role-based access control (RBAC) to allow files to be moved or deleted, if the right permissions are present. If the right permissions are not present, the tasks can be assigned to someone in the organization who does have the right permissions.

## BlueXP classification management and privacy

The following questions provide information on how to manage BlueXP classification and privacy settings.

### How do I enable or disable BlueXP classification?

First you need to deploy an instance of BlueXP classification in BlueXP, or on an on-premises system. Once the instance is running, you can enable the service on existing working environments, databases, and other data sources from the **Configuration** tab or by selecting a specific working environment. [Learn how to get started.](#)



Activating BlueXP classification on a data source results in an immediate initial scan. Scan results display shortly after.

You can disable BlueXP classification from scanning an individual working environment, database, or file share group from the BlueXP classification Configuration page. See [Remove data sources from BlueXP classification.](#)

To completely remove the BlueXP classification instance, you can manually remove the BlueXP classification instance from your cloud provider's portal or on-prem location.

## Can the service exclude scanning data in certain directories?

Yes. If you want BlueXP classification to exclude scanning data that resides in certain data source directories, you can provide that list to the classification engine. After you apply that change, BlueXP classification will exclude scanning data in the specified directories. [Learn more](#).

## Are snapshots that reside on ONTAP volumes scanned?

No. BlueXP classification does not scan snapshots because the content is identical to the content in the volume.

## What happens if data tiering is enabled on your ONTAP volumes?

When BlueXP classification scans volumes that have cold data tiered to object storage using the Mapping only scans, it scans all of the data—data that's on local disks and cold data tiered to object storage. This is also true for non-NetApp products that implement tiering.

The Mapping only scan doesn't heat up the cold data—it stays cold and remains in object storage. On the other hand, if you perform the Map & Classify scan, some configurations might heat up the cold data.

## Types of source systems and data types

The following questions relate to the types of storage that can be scanned, and the types of data that is scanned.

### Are there any restrictions when deployed in a Government region?

BlueXP classification is supported when the Connector is deployed in a Government region (AWS GovCloud, Azure Gov, or Azure DoD) - also known as "Restricted mode".

### What data sources can I scan if I install BlueXP classification in a site without internet access?

BlueXP classification can only scan data from data sources that are local to the on-premises site. At this time, BlueXP classification can scan the following local data sources in "Private mode" - also known as a "dark" site:

- On-premises ONTAP systems
- Database schemas
- Object Storage that uses the Simple Storage Service (S3) protocol

See [Supported working environments and data sources](#).

### Which file types are supported?

BlueXP classification scans all files for category and metadata insights, and displays all file types in the file types section of the dashboard.

When BlueXP classification detects Personal Identifiable Information (PII), or when it performs a DSAR search, only the following file formats are supported:

.CSV, .DCM, .DOC, .DOCX, .JSON, .PDF, .PPTX, .RTF, .TXT, .XLS, .XLSX, Docs, Sheets, and Slides

## What kinds of data and metadata does BlueXP classification capture?

BlueXP classification enables you to run a general "mapping" scan or a full "classification" scan on your data sources. Mapping provides only a high-level overview of your data, whereas Classification provides deep-level scanning of your data. Mapping can be done on your data sources very quickly because it does not access files to see the data inside.

- **Data mapping scan (Mapping only scan):** BlueXP classification scans the metadata only. This is useful for overall data management and governance, quick project scoping, very large estates, and prioritization. Data mapping is based on metadata and is considered a **fast** scan.

After a fast scan, you can generate a Data Mapping Report. This report is an overview of the data stored in your corporate data sources to assist you with decisions about resource utilization, migration, backup, security, and compliance processes.

- **Data classification deep scan (Map & Classify scan):** BlueXP classification scans using standard protocols and read-only permission throughout your environments. Select files are opened and scanned for sensitive business-related data, private information, and issues related to ransomware.

After a full scan there are many additional BlueXP classification features you can apply to your data, such as view and refine data in the Data Investigation page, search for names within files, copy, move, and delete source files, and more.

BlueXP classification captures metadata such as: file name, permissions, creation time, last access, and last modification. This includes all of the metadata that appears in the Data Investigation Details page and in Data Investigation Reports.

BlueXP classification can identify many types of private data such as personal information (PII) and sensitive personal information (SPII). For details about private data, refer to [Categories of private data that BlueXP classification scans](#).

## Can I limit BlueXP classification information to specific users?

Yes, BlueXP classification is fully integrated with BlueXP. BlueXP users can only see information for the working environments they are eligible to view according to their permissions.

Additionally, if you want to allow certain users to just view BlueXP classification scan results without having the ability to manage BlueXP classification settings, you can assign those users the **Classification viewer** role (when using BlueXP in standard mode) or the **Compliance Viewer** role (when using BlueXP in restricted mode). [Learn more](#).

## Can anyone access the private data sent between my browser and BlueXP classification?

No. The private data sent between your browser and the BlueXP classification instance are secured with end-to-end encryption using TLS 1.2, which means NetApp and non-NetApp parties can't read it. BlueXP classification won't share any data or results with NetApp unless you request and approve access.

The data that is scanned stays within your environment.

## How is sensitive data handled?

NetApp does not have access to sensitive data and does not display it in the UI. Sensitive data is masked, for example, the last four numbers are displayed for credit card information.

## Where is the data stored?

Scan results are stored in Elasticsearch within your BlueXP classification instance.

## How is the data accessed?

BlueXP classification accesses data stored in Elasticsearch through API calls, which require authentication and are encrypted using AES-128. Accessing Elasticsearch directly requires root access.

## Licenses and costs

The following question relates to licensing and costs to use BlueXP classification.

### How much does BlueXP classification cost?

BlueXP classification is a BlueXP core capability and is not charged.

## Connector deployment

The following questions relate to the BlueXP Connector.

### What is the Connector?

The Connector is software running on a compute instance either within your cloud account, or on-premises, that enables BlueXP to securely manage cloud resources. You must deploy a Connector to use BlueXP classification.

### Where does the Connector need to be installed?

When scanning data, the BlueXP Connector needs to be installed in the following locations:

- For Cloud Volumes ONTAP in AWS or Amazon FSx for ONTAP: Connector is in AWS.
- For Cloud Volumes ONTAP in Azure or in Azure NetApp Files: Connector is in Azure.
- For Cloud Volumes ONTAP in GCP: Connector is in GCP.
- For on-premises ONTAP systems: Connector is on-premises.

If you have data in these locations, you may need to use [multiple Connectors](#).

### Does BlueXP classification require access to credentials?

BlueXP classification itself doesn't retrieve storage credentials. Instead, they are stored within the BlueXP Connector.

BlueXP classification uses data plane credentials, for example, CIFS credentials to mount shares before scanning.

### Does communication between the service and the Connector use HTTP?

Yes, BlueXP classification communicates with the BlueXP Connector using HTTP.



# BlueXP classification deployment

The following questions relate to the separate BlueXP classification instance.

## What deployment models does BlueXP classification support?

BlueXP allows the user to scan and report on systems virtually anywhere, including on-premises, cloud, and hybrid environments. BlueXP classification is normally deployed using a SaaS model, in which the service is enabled via the BlueXP interface and requires no hardware or software installation. Even in this click-and-run deployment mode, data management can be done regardless of whether the data stores are on premises or in the public cloud.

## What type of instance or VM is required for BlueXP classification?

When [deployed in the cloud](#):

- In AWS, BlueXP classification runs on an m6i.4xlarge instance with a 500 GiB GP2 disk. You can select a smaller instance type during deployment.
- In Azure, BlueXP classification runs on a Standard\_D16s\_v3 VM with a 500 GiB disk.
- In GCP, BlueXP classification runs on an n2-standard-16 VM with a 500 GiB Standard persistent disk.

[Learn more about how BlueXP classification works.](#)

## Can I deploy the BlueXP classification on my own host?

Yes. You can install BlueXP classification software on a Linux host that has internet access in your network or in the cloud. Everything works the same and you continue to manage your scan configuration and results through BlueXP. See [Deploying BlueXP classification on premises](#) for system requirements and installation details.

## What about secure sites without internet access?

Yes, that's also supported. You can [deploy BlueXP classification in an on-premises site that doesn't have internet access](#) for completely secure sites.

# 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 BlueXP classification](#)

## Copyright information

Copyright © 2025 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.



# Cloud Volumes ONTAP documentation

## Cloud Volumes ONTAP

NetApp  
August 19, 2025

# Table of Contents

- Cloud Volumes ONTAP documentation . . . . . 1
- Release notes . . . . . 2
  - What's new in Cloud Volumes ONTAP . . . . . 2
    - 11 August 2025 . . . . . 2
    - 14 July 2025 . . . . . 2
    - 25 June 2025 . . . . . 2
    - 29 May 2025 . . . . . 3
    - 12 May 2025 . . . . . 3
    - 16 April 2025 . . . . . 3
    - 14 April 2025 . . . . . 3
    - 3 April 2025 . . . . . 4
    - 28 March 2025 . . . . . 4
    - 12 March 2025 . . . . . 4
    - 10 March 2025 . . . . . 4
    - 6 March 2025 . . . . . 5
    - 03 March 2025 . . . . . 5
    - 18 February 2025 . . . . . 5
    - 10 February 2025 . . . . . 5
    - 9 December 2024 . . . . . 6
    - 11 November 2024 . . . . . 6
    - 25 October 2024 . . . . . 7
    - 7 October 2024 . . . . . 8
    - 9 September 2024 . . . . . 8
    - 23 August 2024 . . . . . 8
    - 22 August 2024 . . . . . 8
    - 8 August 2024 . . . . . 9
    - 10 June 2024 . . . . . 9
    - 17 May 2024 . . . . . 9
    - 23 April 2024 . . . . . 9
    - 8 March 2024 . . . . . 10
    - 5 March 2024 . . . . . 10
    - 2 February 2024 . . . . . 10
    - 16 January 2024 . . . . . 11
    - 8 January 2024 . . . . . 11
    - 6 December 2023 . . . . . 11
    - 5 December 2023 . . . . . 11
    - 10 November 2023 . . . . . 12
    - 8 November 2023 . . . . . 12
    - 1 November 2023 . . . . . 12
    - 23 October 2023 . . . . . 13
    - 6 October 2023 . . . . . 13
    - 10 September 2023 . . . . . 13
    - 30 July 2023 . . . . . 13

26 July 2023	14
2 July 2023	14
26 June 2023	15
4 June 2023	15
7 May 2023	15
4 April 2023	16
3 April 2023	16
13 March 2023	18
5 March 2023	18
5 February 2023	19
1 January 2023	19
15 December 2022	20
8 December 2022	20
4 December 2022	20
15 November 2022	20
6 November 2022	20
18 September 2022	21
31 July 2022	22
18 July 2022	23
3 July 2022	24
7 June 2022	25
2 May 2022	26
3 April 2022	28
27 February 2022	28
9 February 2022	28
6 February 2022	29
30 January 2022	29
2 January 2022	29
28 November 2021	31
4 October 2021	32
2 September 2021	32
7 July 2021	33
30 May 2021	35
24 May 2021	36
11 Apr 2021	36
8 Mar 2021	36
4 Jan 2021	37
3 Nov 2020	39
Known limitations	39
BlueXP doesn't support FlexGroup volumes creation	39
BlueXP doesn't support S3 with Cloud Volumes ONTAP	39
BlueXP doesn't support disaster recovery for storage VMs	39
Cloud Volumes ONTAP Release Notes	39
Get started	41
Learn about Cloud Volumes ONTAP	41

Supported ONTAP versions for Cloud Volumes ONTAP deployments	42
AWS	42
Azure	43
Google Cloud	44
Get started in Amazon Web Services	45
Quick start for Cloud Volumes ONTAP in AWS	45
Plan your Cloud Volumes ONTAP configuration in AWS	46
Set up your networking	50
Set up Cloud Volumes ONTAP to use a customer-managed key in AWS	73
Set up AWS IAM roles for Cloud Volumes ONTAP nodes	76
Set up licensing for Cloud Volumes ONTAP in AWS	85
Launch Cloud Volumes ONTAP in AWS	92
Deploy Cloud Volumes ONTAP in AWS Secret Cloud or AWS Top Secret Cloud	105
Get started in Microsoft Azure	122
Learn about Cloud Volumes ONTAP deployment options in Azure	122
Get started in BlueXP	123
Deploy Cloud Volumes ONTAP from the Azure marketplace	172
Get started in Google Cloud	175
Quick start for Cloud Volumes ONTAP in Google Cloud	175
Plan your Cloud Volumes ONTAP configuration in Google Cloud	177
Set up Google Cloud networking for Cloud Volumes ONTAP	180
Set up VPC Service Controls to deploy Cloud Volumes ONTAP in Google Cloud	191
Create a Google Cloud service account for Cloud Volumes ONTAP	193
Using customer-managed encryption keys with Cloud Volumes ONTAP	196
Set up licensing for Cloud Volumes ONTAP in Google Cloud	197
Launch Cloud Volumes ONTAP in Google Cloud	202
Google Cloud Platform Image Verification	215
Use Cloud Volumes ONTAP	227
License management	227
Manage capacity-based licensing for Cloud Volumes ONTAP	227
Manage Keystone subscriptions through BlueXP	232
Manage node-based licensing for Cloud Volumes ONTAP	235
Volume and LUN administration	240
Create a FlexVol volume on a Cloud Volumes ONTAP system	240
Manage volumes on Cloud Volumes ONTAP systems	247
Tier inactive Cloud Volumes ONTAP data to a low-cost object storage	256
Connect to a LUN on Cloud Volumes ONTAP from your host system	262
Accelerate data access with FlexCache volumes on a Cloud Volumes ONTAP system	263
Aggregate administration	265
Create an aggregate for Cloud Volumes ONTAP systems	265
Manage aggregates for Cloud Volumes ONTAP clusters	267
Manage the Cloud Volumes ONTAP aggregate capacity on a Connector	268
Storage VM administration	270
Manage storage VMs for Cloud Volumes ONTAP	270
Manage data-serving storage VMs for Cloud Volumes ONTAP in AWS	271

Manage data-serving storage VMs for Cloud Volumes ONTAP in Azure	278
Manage data-serving storage VMs for Cloud Volumes ONTAP in Google Cloud	281
Set up storage VM disaster recovery for Cloud Volumes ONTAP	284
Security and data encryption	284
Encrypt volumes on Cloud Volumes ONTAP with NetApp encryption solutions	284
Manage Cloud Volumes ONTAP encryption keys with AWS Key Management Service	285
Manage Cloud Volumes ONTAP encryption keys with Azure Key Vault	286
Manage Cloud Volumes ONTAP encryption keys with Google Cloud KMS	293
Enable NetApp ransomware protection solutions for Cloud Volumes ONTAP	295
Create tamperproof Snapshot copies of WORM files on Cloud Volumes ONTAP	298
System administration	299
Upgrade Cloud Volumes ONTAP software	299
Register Cloud Volumes ONTAP pay-as-you-go systems	308
Convert a Cloud Volumes ONTAP node-based license to a capacity-based license	309
Start and stop a Cloud Volumes ONTAP system	311
Synchronize Cloud Volumes ONTAP system time using the NTP server	314
Modify system write speed	314
Change the Cloud Volumes ONTAP cluster admin password	315
Add, remove, or delete systems	316
AWS administration	320
Azure administration	322
Google Cloud administration	334
Administer Cloud Volumes ONTAP using System Manager	335
Administer Cloud Volumes ONTAP from the CLI	337
System health and events	338
Verify AutoSupport setup for Cloud Volumes ONTAP	338
Configure EMS for Cloud Volumes ONTAP systems	342
Concepts	343
Licensing	343
Licensing for Cloud Volumes ONTAP	343
Learn more about capacity-based licenses for Cloud Volumes ONTAP	347
Storage	351
Supported client protocols for Cloud Volumes ONTAP	351
Disks and aggregates used for Cloud Volumes ONTAP clusters	352
Learn about support for AWS Elastic Volumes with Cloud Volumes ONTAP	354
Learn about data tiering with Cloud Volumes ONTAP in AWS, Azure, or Google Cloud	360
Cloud Volumes ONTAP storage management	365
Write speed	367
Flash Cache	370
Learn about WORM storage on Cloud Volumes ONTAP	370
High-availability pairs	372
Learn about Cloud Volumes ONTAP HA pairs in AWS	372
Learn about Cloud Volumes ONTAP HA pairs in Azure	379
Learn about Cloud Volumes ONTAP HA pairs in Google Cloud	385
Operations unavailable when a node in Cloud Volumes ONTAP HA pair is offline	389



Learn about Cloud Volumes ONTAP data encryption and ransomware protection . . . . .	390
Encryption of data at rest . . . . .	390
ONTAP virus scanning . . . . .	391
Ransomware protection . . . . .	391
Learn about performance monitoring for Cloud Volumes ONTAP workloads . . . . .	392
Performance technical reports . . . . .	392
CPU performance . . . . .	392
License management for node-based BYOL . . . . .	393
BYOL system licenses . . . . .	393
License management for a new system . . . . .	393
License expiration . . . . .	393
License renewal . . . . .	395
License transfer to a new system . . . . .	395
Learn how AutoSupport and Digital Advisor are used for Cloud Volumes ONTAP . . . . .	395
Supported default configurations for Cloud Volumes ONTAP . . . . .	396
Default setup . . . . .	396
Internal disks for system data . . . . .	398
Knowledge and support . . . . .	401
Register for support . . . . .	401
Support registration overview . . . . .	401
Register BlueXP for NetApp support . . . . .	401
Associate NSS credentials for Cloud Volumes ONTAP support . . . . .	403
Get help . . . . .	405
Get support for a cloud provider file service . . . . .	405
Use self-support options . . . . .	405
Create a case with NetApp support . . . . .	405
Manage your support cases (Preview) . . . . .	408
Legal notices . . . . .	411
Copyright . . . . .	411
Trademarks . . . . .	411
Patents . . . . .	411
Privacy policy . . . . .	411
Open source . . . . .	411

# Cloud Volumes ONTAP documentation

# Release notes

## What's new in Cloud Volumes ONTAP

Learn what's new with Cloud Volumes ONTAP management in BlueXP.

The enhancements described on this page are specific to BlueXP features that enable management of Cloud Volumes ONTAP. To learn what's new with the Cloud Volumes ONTAP software itself, [go to the Cloud Volumes ONTAP Release Notes](#).

### 11 August 2025

#### End of availability of Optimized licenses

Beginning on August 11, 2025, the Cloud Volumes ONTAP Optimized license will be deprecated and will no longer be available for purchase or renewal in the Azure and Google Cloud marketplaces for pay-as-you-go (PAYGO) subscriptions. If you have an existing annual contract with an Optimized license, you can continue to use the license until the end of your contract. When your Optimized license expires, you can opt for Cloud Volumes ONTAP Essentials or Professional licenses in BlueXP.

However, the ability to add or renew Optimized licenses will be available through the APIs.

For information about licensing packages, refer to [Licensing for Cloud Volumes ONTAP](#).

For information about switching to a different charging method, refer to [Manage capacity-based licensing](#).

### 14 July 2025

#### Support for transparent proxy

BlueXP now supports transparent proxy servers in addition to the existing explicit proxy connections. When creating or modifying the BlueXP Connector, you can configure a transparent proxy server to securely manage network traffic to and from Cloud Volumes ONTAP.

For more information about the use of proxy servers in Cloud Volumes ONTAP, refer to:

- [Network configurations to support Connector proxy in AWS](#)
- [Network configurations to support Connector proxy in Azure](#)
- [Network configurations to support Connector proxy in Google Cloud](#)

#### New VM type supported for Cloud Volumes ONTAP in Azure

Beginning with Cloud Volumes ONTAP 9.13.1, L8s\_v3 is supported as a VM type in Azure single and multiple availability zones, for both new and existing high-availability (HA) pair deployments.

For more information, refer to [Supported configurations in Azure](#).

### 25 June 2025

## Restricted availability of BYOL licensing for Cloud Volumes ONTAP

Beginning June 25, 2025, NetApp has restricted the bring your own license (BYOL) licensing model for Cloud Volumes ONTAP. The restriction applies to all customers and Cloud Volumes ONTAP deployments in AWS, Azure, and Google Cloud. The only exemptions are the U.S. Public Sector customers and China region deployments.

NetApp support and services will continue until your BYOL contract expires, but your expired licenses will not be renewed or extended. When your BYOL licenses expire, you must replace them with capacity-based licenses purchased through your cloud marketplace subscriptions. A capacity-based licensing model through hyperscaler marketplaces streamlines the licensing experience and delivers greater business benefits. Contact your NetApp accounts team or customer success representatives to discuss your options of conversion.

For more information, refer to this customer communiqué: [CPC-00661: Changes to Cloud Volumes ONTAP BYOL Policy](#).

## 29 May 2025

### Private mode deployments enabled for Cloud Volumes ONTAP 9.15.1

You can now deploy Cloud Volumes ONTAP 9.15.1 in private mode in AWS, Azure, and Google Cloud. Private mode is enabled for both single-node and high-availability (HA) deployments of Cloud Volumes ONTAP 9.15.1.

For more information about private mode deployments refer to [Learn about BlueXP deployment modes](#).

## 12 May 2025

### Discovery of deployments made through the Azure marketplace in BlueXP

BlueXP now has the capability of discovering the Cloud Volumes ONTAP systems deployed directly through the Azure marketplace. This means that you can now add and manage these systems as working environments in BlueXP, just like any other Cloud Volumes ONTAP system.

[Deploy Cloud Volumes ONTAP from the Azure marketplace](#)

## 16 April 2025

### New regions supported in Azure

You can now deploy Cloud Volumes ONTAP 9.12.1 GA and later in single and multiple availability zones in Azure in the following regions. This includes support for both single-node and high-availability (HA) deployments.

- Spain Central
- Mexico Central

For a list of all regions, refer to the [Global Regions Map under Azure](#).

## 14 April 2025

## **Storage VM creation automated through the APIs in Google Cloud**

You can now use the BlueXP APIs to automate the storage VM creation in Google Cloud. You have been using this feature in Cloud Volumes ONTAP high-availability (HA) configurations, and now you can also use it in single node deployments. By using the BlueXP APIs, you can easily create, rename, and delete additional data-serving storage VMs in your Google Cloud environment, without the need to manually configure the required network interfaces, LIFs, and management LIFs. This automation simplifies the process of managing storage VMs.

[Manage data-serving storage VMs for Cloud Volumes ONTAP in Google Cloud](#)

## **3 April 2025**

### **Support for China regions for Cloud Volumes ONTAP 9.13.1 in AWS**

You can now deploy Cloud Volumes ONTAP 9.13.1 in AWS in China regions. This includes support for both single-node and high-availability (HA) deployments. After you have deployed Cloud Volumes ONTAP 9.13.1, you can upgrade it to later versions. Only licenses purchased directly from NetApp are supported.

For regional availability, refer to the [Global Regions Maps for Cloud Volumes ONTAP](#).

## **28 March 2025**

### **Private mode deployments enabled for Cloud Volumes ONTAP 9.14.1**

You can now deploy Cloud Volumes ONTAP 9.14.1 in private mode in AWS, Azure, and Google Cloud. Private mode is enabled for both single-node and high-availability (HA) deployments of Cloud Volumes ONTAP 9.14.1.

For more information about private mode deployments refer to [Learn about BlueXP deployment modes](#).

## **12 March 2025**

### **New regions supported for multiple availability zone deployments in Azure**

The following regions now support HA multiple availability zone deployments in Azure for Cloud Volumes ONTAP 9.12.1 GA and later:

- Central US
- US Gov Virginia (US Government Region - Virginia)

For a list of all regions, refer to the [Global Regions Map under Azure](#).

## **10 March 2025**

### **Storage VM creation automated through the APIs in Azure**

You can now use the BlueXP APIs to create, rename, and delete additional data-serving storage VMs for Cloud Volumes ONTAP in Azure. Using the APIs automates the process of storage VM creation, including the configuration of the required network interfaces, LIFs, and a management LIF, if you need to use a storage VM for management purposes.

[Manage data-serving storage VMs for Cloud Volumes ONTAP in Azure](#)

## 6 March 2025

### Cloud Volumes ONTAP 9.16.1 GA

You can now use BlueXP to deploy and manage the Cloud Volumes ONTAP 9.16.1 General Availability release in Azure and Google Cloud. However, this version is not available for deployment and upgrade in AWS.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

## 03 March 2025

### Support for New Zealand North region in Azure

The New Zealand North region is now supported in Azure for single node and high-availability (HA) configurations of Cloud Volumes ONTAP 9.12.1 GA and later. Note that the Lsv3 instance type is not supported in this region.

For a list of all supported regions, refer to the [Global Regions Map under Azure](#).

## 18 February 2025

### Introducing Azure marketplace direct deployment

You can now take advantage of Azure marketplace direct deployment to easily and quickly deploy Cloud Volumes ONTAP directly from the Azure marketplace. Using this streamlined method, you can explore the core features and capabilities of Cloud Volumes ONTAP in your environment without the need to set up the BlueXP Connector or meet other onboarding criteria required for deploying Cloud Volumes ONTAP through BlueXP.

- [Learn about Cloud Volumes ONTAP deployment options in Azure](#)
- [Deploy Cloud Volumes ONTAP from the Azure marketplace](#)

## 10 February 2025

### User authentication enabled for accessing System Manager from BlueXP

As a BlueXP administrator, you can now activate authentication for ONTAP users accessing ONTAP System Manager from BlueXP. You can enable this option by editing the BlueXP Connector settings. This option is available for standard and private modes.

[Administer Cloud Volumes ONTAP using System Manager.](#)

### BlueXP Advanced View renamed to System Manager

The option for advanced management of Cloud Volumes ONTAP from BlueXP through ONTAP System Manager has been renamed from **Advanced View** to **System Manager**.

[Administer Cloud Volumes ONTAP using System Manager.](#)

### Introducing a simpler way to manage licenses with the BlueXP digital wallet

Now, you can experience simplified management of Cloud Volumes ONTAP licenses by using improved navigation points within the BlueXP digital wallet:

- Access your Cloud Volumes ONTAP license information easily through the **Governance > Digital wallet > Overview/Direct Licenses** tabs.
- Click **View** on the Cloud Volume ONTAP panel in the **Overview** tab to gain a comprehensive understanding of your capacity-based licenses. This advanced view offers detailed insight into your licenses and subscriptions.
- If you prefer the previous interface, you can click the **Switch to legacy view** button to view license details by type and modify charging methods for your licenses.

[Manage capacity-based licenses.](#)

## 9 December 2024

### List of supported VMs updated for Azure to align with the best practices

The DS\_v2 and Es\_v3 machine families are no longer available for selection on BlueXP when deploying new instances of Cloud Volumes ONTAP in Azure. These families will be retained and supported only in older, existing systems. New deployments of Cloud Volumes ONTAP are supported in Azure only from the 9.12.1 release. We recommend that you switch to either Es\_v4 or any other series compatible with Cloud Volumes ONTAP 9.12.1 and later. The DS\_v2 and Es\_v3 series machines, however, will be available for new deployments made through the API.

[Supported configurations in Azure](#)

## 11 November 2024

### End of availability for node-based licenses

NetApp has planned the end of availability (EOA) and end of support (EOS) of Cloud Volumes ONTAP node-based licensing. Beginning with 11 November, 2024, the limited availability of node-based licenses has been terminated. The support for node-based licensing ends on 31 December, 2024. After the EOA of your node-based licenses, you should transition to capacity-based licensing by using the BlueXP license conversion tool.

For annual or longer-term commitments, NetApp recommends that you contact your NetApp representative prior to the EOA date or license expiration date to ensure that the prerequisites for the transition are in place. If you don't have a long-term contract for a Cloud Volumes ONTAP node and run your system against an on-demand pay-as-you-go (PAYGO) subscription, it is important to plan your conversion before the EOS date. For both long-term contracts and PAYGO subscriptions, you can use the BlueXP license conversion tool for a seamless conversion.

[End of availability of node-based licenses](#)

[Convert a Cloud Volumes ONTAP node-based license to a capacity-based license](#)

### Removal of node-based deployments from BlueXP

The option to deploy Cloud Volumes ONTAP systems by using node-based licenses is deprecated on BlueXP. Except for a few special cases, you cannot use node-based licenses for Cloud Volumes ONTAP deployments for any cloud provider.

NetApp recognizes the following unique licensing requirements in compliance with contractual obligations and operational needs, and will continue to support node-based licenses in these situations:

- U.S. Public Sector customers

- Deployments in private mode
- China region deployments of Cloud Volumes ONTAP in AWS
- If you have a valid, non-expired by-node bring your own license (BYOL license)

#### [End of availability of node-based licenses](#)

### **Addition of a cold tier for Cloud Volumes ONTAP data on Azure Blob storage**

BlueXP now enables you to select a cold tier to store the inactive capacity tier data on Azure Blob storage. Adding the cold tier to the existing hot and cool tiers provides you with a more affordable storage option and improved cost efficiency.

#### [Data tiering in Azure](#)

### **Option to restrict public access to storage account for Azure**

You now have the option to restrict public access to your storage account for Cloud Volumes ONTAP systems in Azure. By disabling access, you can secure your private IP address from exposure even within the same VNet, should there be a need to comply with your organization's security policies. This option also disables data tiering for your Cloud Volumes ONTAP systems, and is applicable to both single node and high-availability pairs.

#### [Security group rules.](#)

### **WORM enablement after deploying Cloud Volumes ONTAP**

You now have the ability to activate write once, read many (WORM) storage on an existing Cloud Volumes ONTAP system using BlueXP. This functionality provides you with the flexibility of enabling WORM on a working environment, even if WORM was not enabled on it during its creation. Once enabled, you cannot disable WORM.

#### [Enabling WORM on a Cloud Volumes ONTAP working environment](#)

## **25 October 2024**

### **List of supported VMs updated for Google Cloud to align with the best practices**

The n1 series machines are no longer available for selection on BlueXP when deploying new instances of Cloud Volumes ONTAP in Google Cloud. The n1 series machines will be retained and supported only in older, existing systems. New deployments of Cloud Volumes ONTAP are supported in Google Cloud only from the 9.8 release. We recommend that you switch to the n2 series machine types that are compatible with Cloud Volumes ONTAP 9.8 and later. The n1 series machines, however, will be available for new deployments performed through the API.

#### [Supported configurations in Google Cloud.](#)

### **Local Zones support for Amazon Web Services in private mode**

BlueXP now supports AWS Local Zones for Cloud Volumes ONTAP high availability (HA) deployments in private mode. The support that was earlier limited to only standard mode has now been extended to include private mode.



AWS Local Zones are not supported when using BlueXP in restricted mode.



For more information on AWS Local Zones with HA Deployments, refer to [AWS Local Zones](#).

## 7 October 2024

### Enhanced user experience in version selection for upgrade

Beginning with this release, when you try to upgrade Cloud Volumes ONTAP using the BlueXP notification, you will receive guidance on the default, latest, and compatible versions to use. Also, now you can select the latest patch or major version compatible with your Cloud Volumes ONTAP instance, or manually enter a version for upgrade.

[Upgrade Cloud Volumes ONTAP software](#)

## 9 September 2024

### WORM and ARP functionalities are no longer chargeable

The built-in data protection and security features of WORM (Write Once Read Many) and ARP (Autonomous Ransomware Protection) will be offered with Cloud Volumes ONTAP licenses at no extra cost. The new pricing model applies to both new and existing BYOL and PAYGO/marketplace subscriptions of AWS, Azure, and Google Cloud. Both capacity-based and node-based licenses will contain ARP and WORM for all configurations, including single node and high-availability (HA) pairs, at no additional cost.

The simplified pricing brings you these benefits:

- Accounts that currently include WORM and ARP will no longer incur charges for these features. Going forward, your billing will only have charges for capacity usage, as it was before this change. WORM and ARP will no longer be included in your future bills.
- If your current accounts do not include these features, you can now opt for WORM and ARP at no additional cost.
- All Cloud Volumes ONTAP offerings for any new accounts will exclude charges for WORM and ARP.

Learn more about these features:

- [Enable NetApp ransomware protection solutions for Cloud Volumes ONTAP](#)
- [WORM storage](#)

## 23 August 2024

### Canada West region now supported in AWS

The Canada West region is now supported in AWS for Cloud Volumes ONTAP 9.12.1 GA and later.

For a list of all regions, see the [Global Regions Map under AWS](#).

## 22 August 2024

### Cloud Volumes ONTAP 9.15.1 GA

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.15.1 General Availability release in AWS, Azure, and Google Cloud.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

## 8 August 2024

### Edge Cache licensing packages deprecated

Edge Cache capacity-based licensing packages will no longer be available for future deployments of Cloud Volumes ONTAP. However, you can use the API to avail this functionality.

### Minimum version support for Flash Cache in Azure

The minimum Cloud Volumes ONTAP version required for configuring Flash Cache in Azure is 9.13.1 GA. You can only use ONTAP 9.13.1 GA and later versions for deploying Flash Cache on Cloud Volumes ONTAP systems in Azure.

For supported configurations, see [Supported configurations in Azure](#).

### Free trials for marketplace subscriptions deprecated

The 30-day automatic free trial or evaluation license for pay-as-you-go subscriptions in cloud provider's marketplace will no longer be available in Cloud Volumes ONTAP. The charging for any type of marketplace subscription (PAYGO or annual contract) will be activated from the first use, without any free trial period.

## 10 June 2024

### Cloud Volumes ONTAP 9.15.0

BlueXP can now deploy and manage the Cloud Volumes ONTAP 9.15.0 in AWS, Azure, and Google Cloud.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

## 17 May 2024

### Amazon Web Services Local Zones support

Support for AWS Local Zones is now available for Cloud Volumes ONTAP HA deployments. AWS Local Zones are an infrastructure deployment where storage, compute, database, and other select AWS services are located close to large cities and industry areas.



AWS Local Zones are supported when using BlueXP in standard mode. At this time, AWS Local Zones are not supported when using BlueXP in restricted mode or private mode.

For more information on AWS Local Zones with HA Deployments, refer to [AWS Local Zones](#).

## 23 April 2024

### New regions supported for multiple availability zone deployments in Azure

The following regions now support HA multiple availability zone deployments in Azure for Cloud Volumes ONTAP 9.12.1 GA and later:

- Germany West Central

- Poland Central
- West US 3
- Israel Central
- Italy North
- Canada Central

For a list of all regions, refer to the [Global Regions Map under Azure](#).

### **Johannesburg region now supported in Google Cloud**

The Johannesburg region (`africa-south1` region) is now supported in Google Cloud for Cloud Volumes ONTAP 9.12.1 GA and later.

For a list of all regions, refer to the [Global Regions Map under Google Cloud](#).

### **Volume templates and tags no longer supported**

You can no longer create a volume from a template or edit a volume's tags. These actions were associated with the BlueXP remediation service, which is no longer available.

## **8 March 2024**

### **Amazon Instant Metadata Service v2 support**

In AWS, Cloud Volumes ONTAP, the Mediator, and the Connector now support Amazon Instant Metadata Service v2 (IMDSv2) for all functions. IMDSv2 provides enhanced protection against vulnerabilities. Only IMDSv1 was previously supported.

If required by your security policies, you can configure your EC2 instances to use IMDSv2. For instructions, refer to [BlueXP setup and administration documentation for managing existing Connectors](#).

## **5 March 2024**

### **Cloud Volumes ONTAP 9.14.1 GA**

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.14.1 General Availability release in AWS, Azure, and Google Cloud.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

## **2 February 2024**

### **Support for Edv5-series VMs in Azure**

Cloud Volumes ONTAP now supports the following Edv5-series VMs starting with the 9.14.1 release.

- E4ds\_v5
- E8ds\_v5
- E20s\_v5
- E32ds\_v5

- [E48ds\\_v5](#)
- [E64ds\\_v5](#)

[Supported configurations in Azure](#)

## 16 January 2024

### Patch releases in BlueXP

Patch releases are available in BlueXP only for the latest three versions of Cloud Volumes ONTAP.

[Upgrade Cloud Volumes ONTAP](#)

## 8 January 2024

### New VMs for Azure multiple availability zones

Starting from Cloud Volumes ONTAP 9.13.1, the following VM types support Azure multiple availability zones for new and existing high-availability pair deployments:

- [L16s\\_v3](#)
- [L32s\\_v3](#)
- [L48s\\_v3](#)
- [L64s\\_v3](#)

[Supported configurations in Azure](#)

## 6 December 2023

### Cloud Volumes ONTAP 9.14.1 RC1

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.14.1 in AWS, Azure, and Google Cloud.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

### 300 TiB FlexVol volume max limit

You can now create a FlexVol volume up to the maximum size of 300 TiB with System Manager and the ONTAP CLI starting from Cloud Volumes ONTAP 9.12.1 P2 and 9.13.0 P2, and in BlueXP starting from Cloud Volumes ONTAP 9.13.1.

- [Storage limits in AWS](#)
- [Storage limits in Azure](#)
- [Storage limits in Google Cloud](#)

## 5 December 2023

The following changes were introduced.

## **New region support in Azure**

### **Single availability zone region support**

The following regions now support highly-available single availability zone deployments in Azure for Cloud Volumes ONTAP 9.12.1 GA and later:

- Tel Aviv
- Milan

### **Multiple availability zone region support**

The following regions now support highly-available multiple availability zone deployments in Azure for Cloud Volumes ONTAP 9.12.1 GA and later:

- Central India
- Norway East
- Switzerland North
- South Africa North
- United Arab Emirates North

For a list of all regions, refer to the [Global Regions Map under Azure](#).

## **10 November 2023**

The following change was introduced with the 3.9.35 release of the Connector.

### **Berlin region now supported in Google Cloud**

The Berlin region is now supported in Google Cloud for Cloud Volumes ONTAP 9.12.1 GA and later.

For a list of all regions, refer to the [Global Regions Map under Google Cloud](#).

## **8 November 2023**

The following change was introduced with the 3.9.35 release of the Connector.

### **Tel Aviv region now supported in AWS**

The Tel Aviv region is now supported in AWS for Cloud Volumes ONTAP 9.12.1 GA and later.

For a list of all regions, refer to the [Global Regions Map under AWS](#).

## **1 November 2023**

The following change was introduced with the 3.9.34 release of the Connector.

### **Saudi Arabia region now supported in Google Cloud**

The Saudi Arabia region is now supported in Google Cloud for Cloud Volumes ONTAP and the Connector for Cloud Volumes ONTAP 9.12.1 GA and later.

For a list of all regions, refer to the [Global Regions Map under Google Cloud](#).

## 23 October 2023

The following change was introduced with the 3.9.34 release of the Connector.

### **New regions supported for HA multiple availability zone deployments in Azure**

The following regions in Azure now support highly-available multiple availability zone deployments for Cloud Volumes ONTAP 9.12.1 GA and later:

- Australia East
- East Asia
- France Central
- North Europe
- Qatar Central
- Sweden Central
- West Europe
- West US 2

For a list of all regions that support multiple availability zones, refer to the [Global Regions Map under Azure](#).

## 6 October 2023

The following change was introduced with the 3.9.34 release of the Connector.

### **Cloud Volumes ONTAP 9.14.0**

BlueXP can now deploy and manage the Cloud Volumes ONTAP 9.14.0 General Availability release in AWS, Azure, and Google Cloud.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

## 10 September 2023

The following change was introduced with the 3.9.33 release of the Connector.

### **Support for Lsv3-series VMs in Azure**

The L48s\_v3 and L64s\_v3 instance types are now supported with Cloud Volumes ONTAP in Azure for single node and high-availability pair deployments with shared managed disks in single and multiple availability zones, starting with the 9.13.1 release. These instance types support Flash Cache.

[View supported configurations for Cloud Volumes ONTAP in Azure](#)

[View storage limits for Cloud Volumes ONTAP in Azure](#)

## 30 July 2023

The following changes were introduced with the 3.9.32 release of the Connector.

## Flash Cache and high write speed support in Google Cloud

Flash Cache and high write speed can be enabled separately in Google Cloud for Cloud Volumes ONTAP 9.13.1 and later. High write speed is available on all supported instance types. Flash Cache is supported on the following instance types:

- n2-standard-16
- n2-standard-32
- n2-standard-48
- n2-standard-64

You can use these features separately or together on both single node and high-availability pair deployments.

[Launch Cloud Volumes ONTAP in Google Cloud](#)

## Usage reports enhancements

Various improvements to the displayed information within the usage reports are now available. The following are enhancements to the usage reports:

- The TiB unit is now included in the name of columns.
- A new "node(s)" field for serial numbers is now included.
- A new "Workload Type" column is now included under the Storage VMs usage report.
- Working environment names now included in Storage VMs and Volume usage reports.
- Volume type "file" is now labeled "Primary (Read/Write)".
- Volume type "secondary" is now labeled "Secondary (DP)".

For more information on usage reports, refer to [Download usage reports](#).

## 26 July 2023

The following changes were introduced with the 3.9.31 release of the Connector.

### Cloud Volumes ONTAP 9.13.1 GA

BlueXP can now deploy and manage the Cloud Volumes ONTAP 9.13.1 General Availability release in AWS, Azure, and Google Cloud.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

## 2 July 2023

The following changes were introduced with the 3.9.31 release of the Connector.

### Support for HA multiple availability zone deployments in Azure

The Japan East and Korea Central in Azure now supports HA multiple availability zone deployments for Cloud Volumes ONTAP 9.12.1 GA and later.

For a list of all regions that support multiple availability zones, refer to the [Global Regions Map under Azure](#).

## **Autonomous Ransomware Protection support**

Autonomous Ransomware Protection (ARP) is now supported on Cloud Volumes ONTAP. ARP support is available on Cloud Volumes ONTAP version 9.12.1 and higher.

To learn more about ARP with Cloud Volumes ONTAP, refer to [Autonomous Ransomware Protection](#).

## **26 June 2023**

The following change was introduced with the 3.9.30 release of the Connector.

### **Cloud Volumes ONTAP 9.13.1 RC1**

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.13.1 in AWS, Azure, and Google Cloud.

[Learn about the new features included in this release of Cloud Volumes ONTAP](#).

## **4 June 2023**

The following change was introduced with the 3.9.30 release of the Connector.

### **Cloud Volumes ONTAP upgrade version selector update**

Through the Upgrade Cloud Volumes ONTAP page, you can now choose to upgrade to the latest available version of Cloud Volumes ONTAP or an older version.

To learn more about upgrading Cloud Volumes ONTAP through BlueXP, refer to [Upgrade Cloud Volumes ONTAP](#).

## **7 May 2023**

The following changes were introduced with the 3.9.29 release of the Connector.

### **Qatar region now supported in Google Cloud**

The Qatar region is now supported in Google Cloud for Cloud Volumes ONTAP and the Connector for Cloud Volumes ONTAP 9.12.1 GA and later.

### **Sweden Central region now supported in Azure**

The Sweden Central region is now supported in Azure for Cloud Volumes ONTAP and the Connector for Cloud Volumes ONTAP 9.12.1 GA and later.

### **Support for HA multiple availability zone deployments in Azure Australia East**

The Australia East region in Azure now supports HA multiple availability zone deployments for Cloud Volumes ONTAP 9.12.1 GA and later.

## **Charging usage breakdown**

Now you can find out what you're being charged for when you're subscribed to capacity-based licenses. The following types of usage reports are available for download from the digital wallet in BlueXP. The usage reports provide capacity details of your subscriptions and tell you how you're being charged for the resources in your Cloud Volumes ONTAP subscriptions. The downloadable reports can be easily shared with others.



- Cloud Volumes ONTAP package usage
- High-level usage
- Storage VMs usage
- Volumes usage

For more information, refer to [Manage capacity-based licenses](#).

### **Notification now displays when accessing BlueXP without a marketplace subscription**

A notification now displays whenever you access Cloud Volumes ONTAP in BlueXP without a marketplace subscription. The notification states "a marketplace subscription for this working environment is required to be compliant with Cloud Volumes ONTAP terms and conditions."

## **4 April 2023**

### **Support for China regions for AWS**

Starting with Cloud Volumes ONTAP 9.12.1 GA, China regions are now supported in AWS as follows.

- Single node systems are supported.
- Licenses purchased directly from NetApp are supported.

For regional availability, refer to the [Global Regions Maps for Cloud Volumes ONTAP](#).

## **3 April 2023**

The following changes were introduced with the 3.9.28 release of the Connector.

### **Turin region now supported in Google Cloud**

The Turin region is now supported in Google Cloud for Cloud Volumes ONTAP and the Connector for Cloud Volumes ONTAP 9.12.1 GA and later.

### **BlueXP digital wallet enhancement**

The BlueXP digital wallet now shows the licensed capacity that you purchased with marketplace private offers.

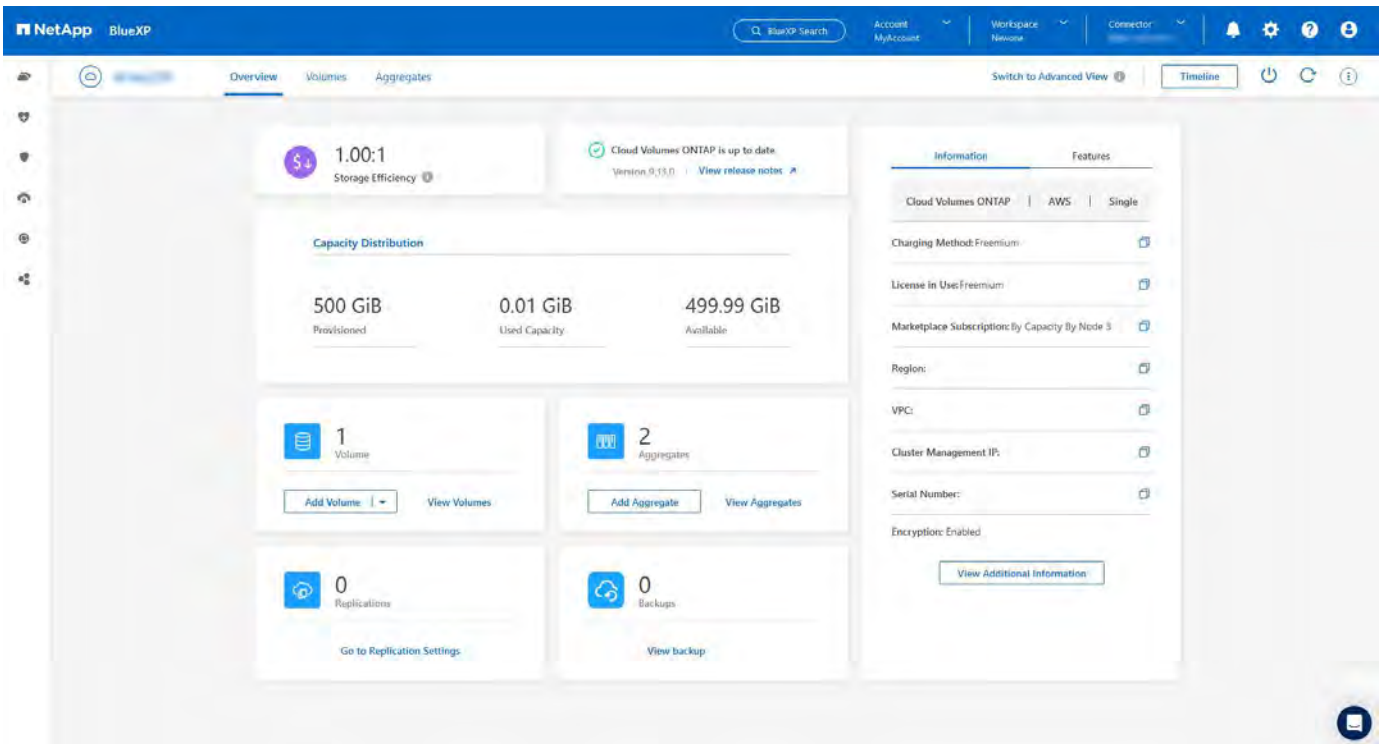
[Learn how to view the consumed capacity in your account.](#)

### **Support for comments during volume creation**

This release enables you to make comments when creating an Cloud Volumes ONTAP FlexGroup volume or FlexVol volume when using the API.

### **BlueXP user interface redesign for Cloud Volumes ONTAP Overview, Volumes, and Aggregates pages**

BlueXP now has a redesigned user interface for Cloud Volumes ONTAP Overview, Volumes, and Aggregates pages. The tile-based design presents more comprehensive information in each tile for a better user experience.

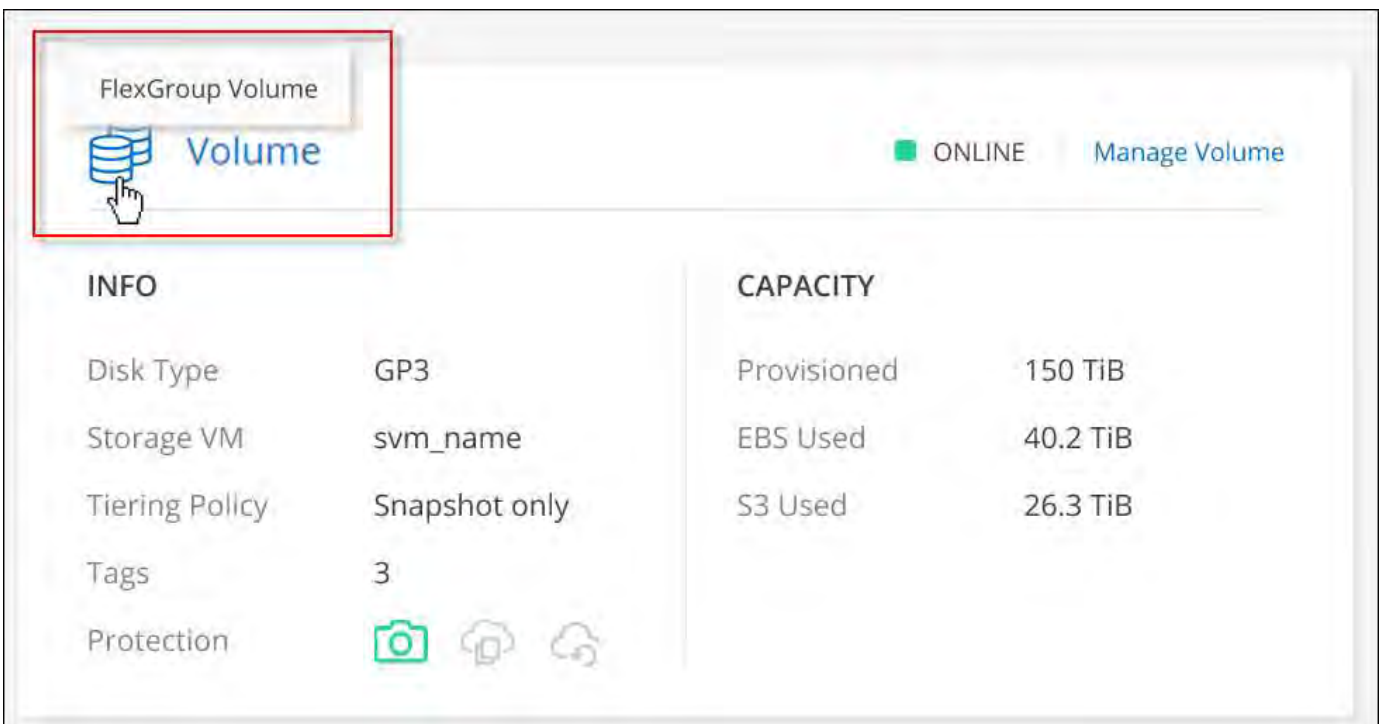


## FlexGroup Volumes viewable through Cloud Volumes ONTAP

FlexGroup volumes created through ONTAP System Manager or the ONTAP CLI directly are now viewable through the redesigned Volumes tile in BlueXP. Identical to the information provided for FlexVol volumes, BlueXP provides detailed information for created FlexGroup volumes through a dedicated Volumes tile.



Currently, you can only view existing FlexGroup volumes under BlueXP. The ability to create FlexGroup volumes in BlueXP is not available but planned for a future release.



[Learn more about viewing created FlexGroup volumes.](#)

## 13 March 2023

### Support for China regions in Azure

China North 3 region is now supported for single node deployments of Cloud Volumes ONTAP 9.12.1 GA and 9.13.0 GA in Azure. Only licenses purchased directly from NetApp (BYOL licenses) are supported in these regions.



Fresh deployments of Cloud Volumes ONTAP in China regions are supported only in 9.12.1 GA and 9.13.0 GA. You can upgrade these versions to later patches and releases of Cloud Volumes ONTAP. If you want to deploy later Cloud Volumes ONTAP versions in China regions, contact NetApp Support.

For regional availability, refer to the [Global Regions Maps for Cloud Volumes ONTAP](#).

## 5 March 2023

The following changes were introduced with the 3.9.27 release of the Connector.

### Cloud Volumes ONTAP 9.13.0

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.13.0 in AWS, Azure, and Google Cloud.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

### 16 TiB and 32 Tib support in Azure

Cloud Volumes ONTAP now supports 16 TiB and 32 TiB disk sizes for high-availability deployments running on managed disks in Azure.

Learn more about [supported disk sizes in Azure](#).

### MTEKM license

The Multi-tenant Encryption Key Management (MTEKM) license is now included with new and existing Cloud Volumes ONTAP systems running version 9.12.1 GA or later.

Multi-tenant external key management enables individual storage VMs (SVMs) to maintain their own keys through a KMIP server when using NetApp Volume Encryption.

[Learn how to encrypt volumes with NetApp encryption solutions.](#)

### Support for environments without internet

Cloud Volumes ONTAP is now supported in any cloud environment that has complete isolation from the internet. Only node-based licensing (BYOL) is supported in these environments. Capacity-based licensing is not supported. To get started, manually install the Connector software, log in to the BlueXP console that's running on the Connector, add your BYOL license to the BlueXP digital wallet, and then deploy Cloud Volumes ONTAP.

- [Install the Connector in a location without internet access](#)

- [Access the BlueXP console on the Connector](#)
- [Add an unassigned license](#)

## Flash Cache and high write speed in Google Cloud

Support for Flash Cache, high write speed, and a high maximum transmission unit (MTU) of 8,896 bytes is now available for select instances with the Cloud Volumes ONTAP 9.13.0 release.

Learn more about [supported configurations by license for Google Cloud](#).

## 5 February 2023

The following changes were introduced with the 3.9.26 release of the Connector.

### Placement group creation in AWS

A new configuration setting is now available for placement group creation with AWS HA single availability zone (AZ) deployments. Now you can choose to bypass failed placement group creations and allow AWS HA single AZ deployments to complete successfully.

For detailed information on how to configure the placement group creation setting, refer to [Configure placement group creation for AWS HA Single AZ](#).

### Private DNS zone configuration update

A new configuration setting is now available so that you can avoid creating a link between a private DNS zone and a virtual network when using Azure Private Links. Creation is enabled by default.

[Provide BlueXP with details about your Azure Private DNS](#)

### WORM storage and data tiering

You can now enable both data tiering and WORM storage together when you create a Cloud Volumes ONTAP 9.8 system or later. Enabling data tiering with WORM storage allows you to tier the data to an object store in the cloud.

[Learn about WORM storage](#).

## 1 January 2023

The following changes were introduced with the 3.9.25 release of the Connector.

### Licensing packages available in Google Cloud

Optimized and Edge Cache capacity-based licensing packages are available for Cloud Volumes ONTAP in the Google Cloud Marketplace as a pay-as-you-go offering or as an annual contract.

Refer to [Cloud Volumes ONTAP licensing](#).

### Default configuration for Cloud Volumes ONTAP

The Multi-tenant Encryption Key Management (MTEKM) license is no longer included in new Cloud Volumes ONTAP deployments.

For more information on the ONTAP feature licenses automatically installed with Cloud Volumes ONTAP, refer to [Default Configuration for Cloud Volumes ONTAP](#).

## 15 December 2022

### Cloud Volumes ONTAP 9.12.0

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.12.0 in AWS and Google Cloud.

[Learn about the new features included in this release of Cloud Volumes ONTAP](#).

## 8 December 2022

### Cloud Volumes ONTAP 9.12.1

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.12.1, which includes support for new features and additional cloud provider regions.

[Learn about the new features included in this release of Cloud Volumes ONTAP](#)

## 4 December 2022

The following changes were introduced with the 3.9.24 release of the Connector.

### **WORM + Cloud Backup now available during Cloud Volumes ONTAP creation**

The ability to activate both write once, read many (WORM) and Cloud Backup features is now available during the Cloud Volumes ONTAP creation process.

### **Israel region now supported in Google Cloud**

The Israel region is now supported in Google Cloud for Cloud Volumes ONTAP and the Connector for Cloud Volumes ONTAP 9.11.1 P3 and later.

## 15 November 2022

The following changes were introduced with the 3.9.23 release of the Connector.

### **ONTAP S3 license in Google Cloud**

An ONTAP S3 license is now included on new and existing Cloud Volumes ONTAP systems running version 9.12.1 or later in Google Cloud Platform.

[ONTAP documentation: Learn how to configure and manage S3 object storage services](#)

## 6 November 2022

The following changes were introduced with the 3.9.23 release of the Connector.

### **Moving resource groups in Azure**

You can now move a working environment from one resource group to a different resource group in Azure within the same Azure subscription.

For more information, refer to [Moving resource groups](#).

### **NDMP-copy certification**

NDMP-copy is now certified for use with Cloud Volume ONTAP.

For information on how to configure and use NDMP, refer to the [ONTAP documentation: NDMP configuration overview](#).

### **Managed disk encryption support for Azure**

A new Azure permission has been added that now allows you to encrypt all managed disks upon creation.

For more information on this new functionality, refer to [Set up Cloud Volumes ONTAP to use a customer-managed key in Azure](#).

## **18 September 2022**

The following changes were introduced with the 3.9.22 release of the Connector.

### **Digital Wallet enhancements**

- The Digital Wallet now shows a summary of the Optimized I/O licensing package and the provisioned WORM capacity for Cloud Volumes ONTAP systems across your account.

These details can help you better understand how you're being charged and whether you need to purchase additional capacity.

[Learn how to view the consumed capacity in your account.](#)

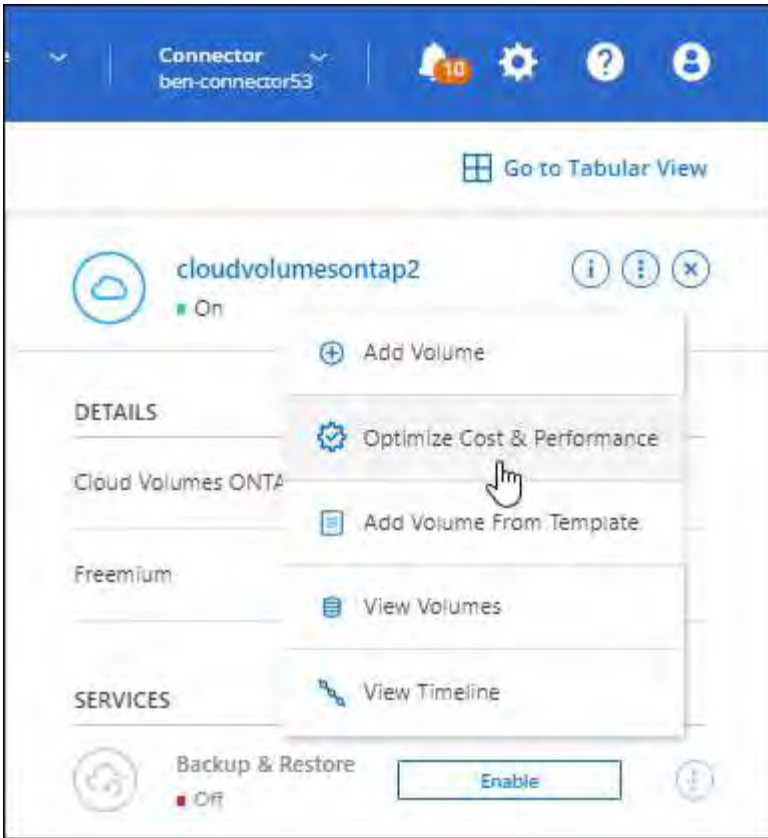
- You can now change from one charging method to the Optimized charging method.

[Learn how to change charging methods.](#)

### **Optimize cost and performance**

You can now optimize the cost and performance of a Cloud Volumes ONTAP system directly from the Canvas.

After you select a working environment, you can choose the **Optimize Cost & Performance** option to change the instance type for Cloud Volumes ONTAP. Choosing a smaller-sized instance can help you reduce costs, while changing to a larger-sized instance can help you optimize performance.



### AutoSupport notifications

BlueXP will now generate a notification if a Cloud Volumes ONTAP system is unable to send AutoSupport messages. The notification includes a link to instructions that you can use to troubleshoot networking issues.

### 31 July 2022

The following changes were introduced with the 3.9.21 release of the Connector.

### MTEKM license

The Multi-tenant Encryption Key Management (MTEKM) license is now included with new and existing Cloud Volumes ONTAP systems running version 9.11.1 or later.

Multi-tenant external key management enables individual storage VMs (SVMs) to maintain their own keys through a KMIP server when using NetApp Volume Encryption.

[Learn how to encrypt volumes with NetApp encryption solutions.](#)

### Proxy server

BlueXP now automatically configures your Cloud Volumes ONTAP systems to use the Connector as a proxy server, if an outbound internet connection isn't available to send AutoSupport messages.

AutoSupport proactively monitors the health of your system and sends messages to NetApp technical support.

The only requirement is to ensure that the Connector's security group allows *inbound* connections over port 3128. You'll need to open this port after you deploy the Connector.

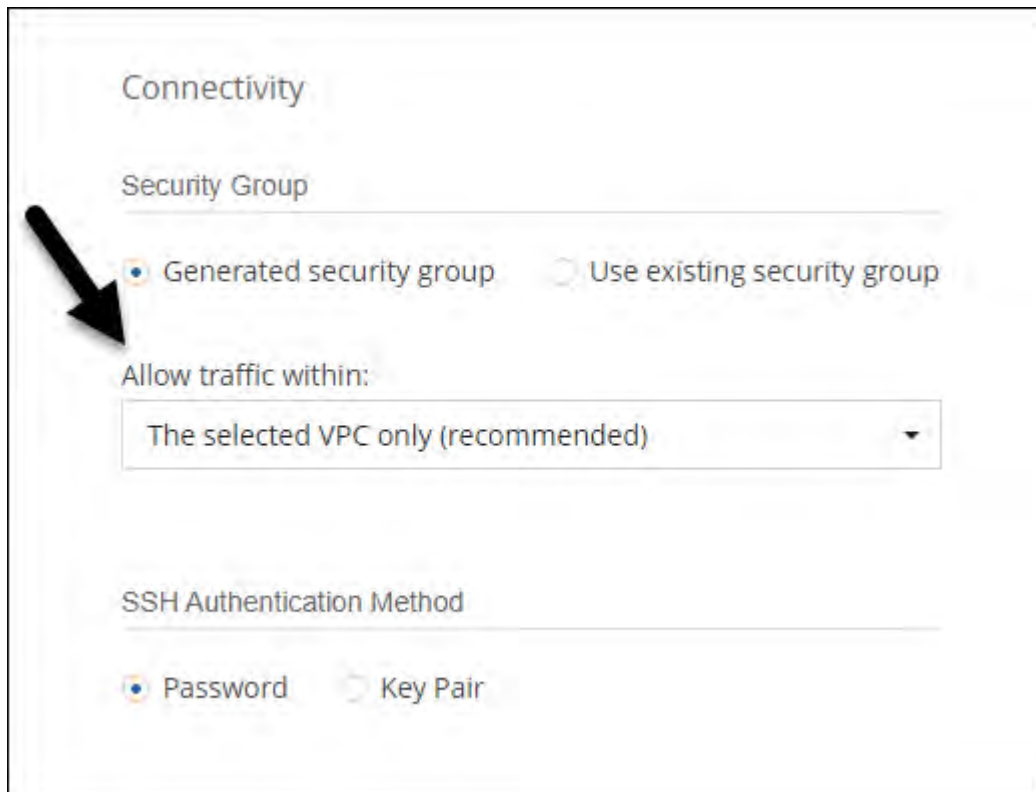
## Change charging method

You can now change the charging method for a Cloud Volumes ONTAP system that uses capacity-based licensing. For example, if you deployed a Cloud Volumes ONTAP system with the Essentials package, you can change it to the Professional package if your business needs changed. This feature is available from the Digital Wallet.

[Learn how to change charging methods.](#)

## Security group enhancement

When you create a Cloud Volumes ONTAP working environment, the user interface now enables you to choose whether you want the predefined security group to allow traffic within the selected network only (recommended) or all networks.



Connectivity

Security Group

Generated security group  Use existing security group

Allow traffic within:

The selected VPC only (recommended)

SSH Authentication Method

Password  Key Pair

## 18 July 2022

### New licensing packages in Azure

Two new capacity-based licensing packages are available for Cloud Volumes ONTAP in Azure when you pay through an Azure Marketplace subscription:

- **Optimized:** Pay for provisioned capacity and I/O operations separately
- **Edge Cache:** Licensing for [Cloud Volumes Edge Cache](#)

[Learn more about these licensing packages.](#)

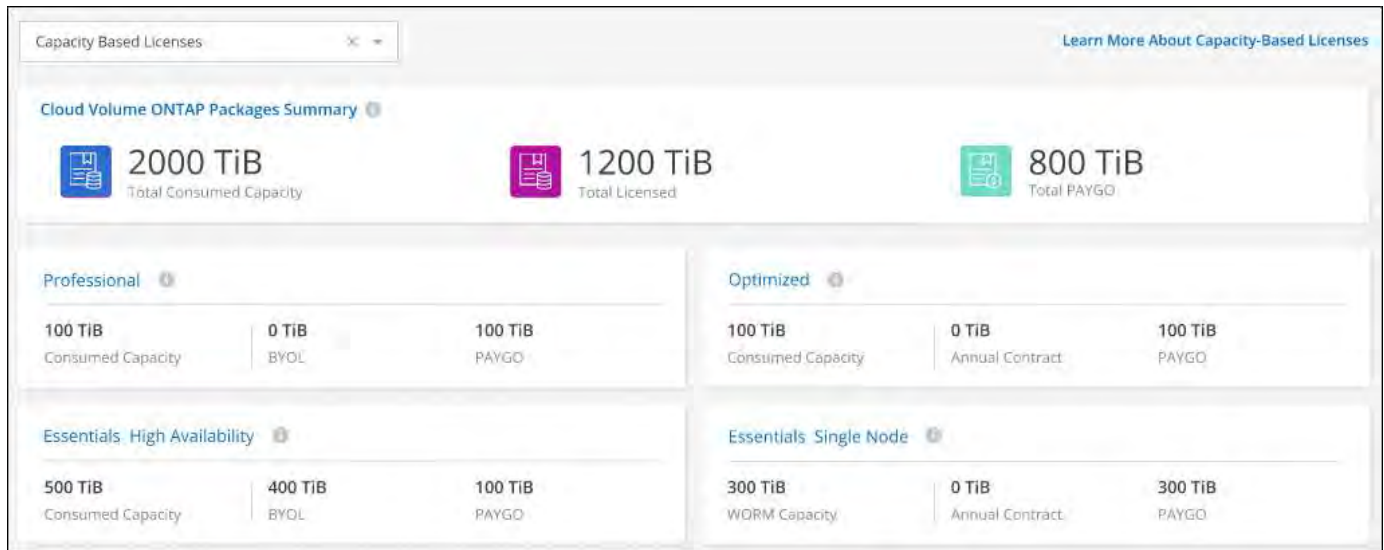


## 3 July 2022

The following changes were introduced with the 3.9.20 release of the Connector.

### Digital Wallet

The Digital Wallet now shows you the total consumed capacity in your account and the consumed capacity by licensing package. This can help you understand how you're being charged and whether you need to purchase additional capacity.



### Elastic Volumes enhancement

BlueXP now supports the Amazon EBS Elastic Volumes feature when creating a Cloud Volumes ONTAP working environment from the user interface. The Elastic Volumes feature is enabled by default when using gp3 or io1 disks. You can choose the initial capacity based on your storage needs and revise it after Cloud Volumes ONTAP is deployed.

[Learn more about support for Elastic Volumes in AWS.](#)

### ONTAP S3 license in AWS

An ONTAP S3 license is now included on new and existing Cloud Volumes ONTAP systems running version 9.11.0 or later in AWS.

[ONTAP documentation: Learn how to configure and manage S3 object storage services](#)

### New Azure Cloud region support

Starting with the 9.10.1 release, Cloud Volumes ONTAP is now supported in the Azure West US 3 region.

[View the full list of supported regions for Cloud Volumes ONTAP](#)

### ONTAP S3 license in Azure

An ONTAP S3 license is now included on new and existing Cloud Volumes ONTAP systems running version 9.9.1 or later in Azure.

## 7 June 2022

The following changes were introduced with the 3.9.19 release of the Connector.

### Cloud Volumes ONTAP 9.11.1

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.11.1, which includes support for new features and additional cloud provider regions.

[Learn about the new features included in this release of Cloud Volumes ONTAP](#)

### New Advanced View

If you need to perform advanced management of Cloud Volumes ONTAP, you can do so using ONTAP System Manager, which is a management interface that's provided with an ONTAP system. We have included the System Manager interface directly inside BlueXP so that you don't need to leave BlueXP for advanced management.

This Advanced View is available as a Preview with Cloud Volumes ONTAP 9.10.0 and later. We plan to refine this experience and add enhancements in upcoming releases. Please send us feedback by using the in-product chat.

[Learn more about the Advanced View.](#)

### Support for Amazon EBS Elastic Volumes

Support for the Amazon EBS Elastic Volumes feature with a Cloud Volumes ONTAP aggregate provides better performance and additional capacity, while enabling BlueXP to automatically increase the underlying disk capacity as needed.

Support for Elastic Volumes is available starting with *new* Cloud Volumes ONTAP 9.11.0 systems and with gp3 and io1 EBS disk types.

[Learn more about support for Elastic Volumes.](#)

Note that support for Elastic Volumes requires new AWS permissions for the Connector:

```
"ec2:DescribeVolumesModifications",  
"ec2:ModifyVolume",
```

Be sure to provide these permissions to each set of AWS credentials that you've added to BlueXP. [View the latest Connector policy for AWS.](#)

### Support for deploying HA pairs in shared AWS subnets

Cloud Volumes ONTAP 9.11.1 includes support for AWS VPC sharing. This release of the Connector enables you to deploy an HA pair in an AWS shared subnet when using the API.

[Learn how to deploy an HA pair in a shared subnet.](#)

## Limited network access when using service endpoints

BlueXP now limits network access when using a VNet service endpoint for connections between Cloud Volumes ONTAP and storage accounts. BlueXP uses a service endpoint if you disable Azure Private Link connections.

[Learn more about Azure Private Link connections with Cloud Volumes ONTAP.](#)

## Support for creating storage VMs in Google Cloud

Multiple storage VMs are now supported with Cloud Volumes ONTAP in Google Cloud, starting with the 9.11.1 release. Starting with this release of the Connector, BlueXP enables you to create storage VMs on Cloud Volumes ONTAP HA pairs in Google Cloud by using the API.

Support for creating storage VMs requires new Google Cloud permissions for the Connector:

- `compute.instanceGroups.get`
- `compute.addresses.get`

Note that you must use the ONTAP CLI or System Manager to create a storage VM on a single node system.

- [Learn more about storage VM limits in Google Cloud](#)
- [Learn how to create data-serving storage VMs for Cloud Volumes ONTAP in Google Cloud](#)

## 2 May 2022

The following changes were introduced with the 3.9.18 release of the Connector.

### Cloud Volumes ONTAP 9.11.0

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.11.0.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

### Enhancement to mediator upgrades

When BlueXP upgrades the mediator for an HA pair, it now validates that a new mediator image is available before it deletes the boot disk. This change ensures that the mediator can continue to operate successfully if the upgrade process is unsuccessful.

### K8s tab has been removed

The K8s tab was deprecated in a previous release, and has now been removed.

### Annual contract in Azure

The Essentials and Professional packages are now available in Azure through an annual contract. You can contact your NetApp sales representative to purchase an annual contract. The contract is available as a private offer in the Azure Marketplace.

After NetApp shares the private offer with you, you can select the annual plan when you subscribe from the Azure Marketplace during working environment creation.

[Learn more about licensing.](#)

### **S3 Glacier Instant Retrieval**

You can now store tiered data in the Amazon S3 Glacier Instant Retrieval storage class.

[Learn how to change the storage class for tiered data.](#)

### **New AWS permissions required for the Connector**

The following permissions are now required to create an AWS spread placement group when deploying an HA pair in a single Availability Zone (AZ):

```
"ec2:DescribePlacementGroups",  
"iam:GetRolePolicy",
```

These permissions are now required to optimize how BlueXP creates the placement group.

Be sure to provide these permissions to each set of AWS credentials that you've added to BlueXP. [View the latest Connector policy for AWS.](#)

### **New Google Cloud region support**

Cloud Volumes ONTAP is now supported in the following Google Cloud regions starting with the 9.10.1 release:

- Delhi (asia-south2)
- Melbourne (australia-southeast2)
- Milan (europe-west8) - single node only
- Santiago (southamerica-west1) - single node only

[View the full list of supported regions for Cloud Volumes ONTAP](#)

### **Support for n2-standard-16 in Google Cloud**

The n2-standard-16 machine type is now supported with Cloud Volumes ONTAP in Google Cloud, starting with the 9.10.1 release.

[View supported configurations for Cloud Volumes ONTAP in Google Cloud](#)

### **Enhancements to Google Cloud firewall policies**

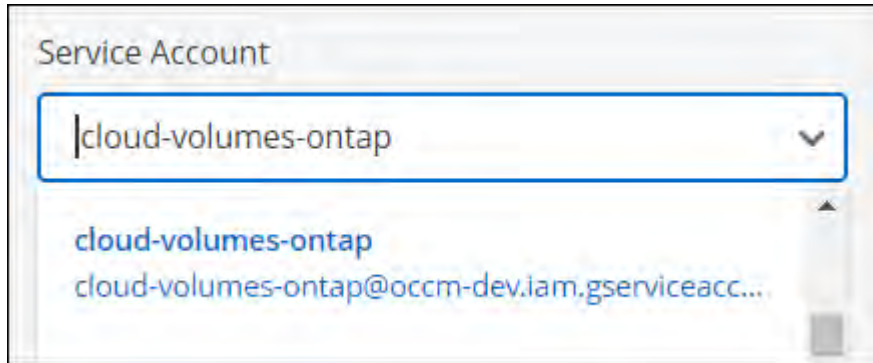
- When you create a Cloud Volumes ONTAP HA pair in Google Cloud, BlueXP will now display all existing firewall policies in a VPC.

Previously, BlueXP wouldn't display any policies in VPC-1, VPC-2, or VPC-3 that didn't have a target tag.

- When you create a Cloud Volumes ONTAP single node system in Google Cloud, you can now choose whether you want the predefined firewall policy to allow traffic within the selected VPC only (recommended) or all VPCs.

## Enhancement to Google Cloud service accounts

When you select the Google Cloud service account to use with Cloud Volumes ONTAP, BlueXP now displays the email address that's associated with each service account. Viewing the email address can make it easier to distinguish between service accounts that share the same name.



## 3 April 2022

### System Manager link has been removed

We have removed the System Manager link that was previously available from within a Cloud Volumes ONTAP working environment.

You can still connect to System Manager by entering the cluster management IP address in a web browser that has a connection to the Cloud Volumes ONTAP system. [Learn more about connecting to System Manager.](#)

### Charging for WORM storage

Now that the introductory special rate has expired, you will now be charged for using WORM storage. Charging is hourly, according to the total provisioned capacity of WORM volumes. This applies to new and existing Cloud Volumes ONTAP systems.

[Learn about pricing for WORM storage.](#)

## 27 February 2022

The following changes were introduced with the 3.9.16 release of the Connector.

### Redesigned volume wizard

The create new volume wizard that we recently introduced is now available when creating a volume on a specific aggregate from the **Advanced allocation** option.

[Learn how to create volumes on a specific aggregate.](#)

## 9 February 2022

### Marketplace updates

- The Essentials package and Professional package are now available in all cloud provider marketplaces.

These by-capacity charging methods enable you to pay by the hour or to purchase an annual contract

directly from your cloud provider. You still have the option to purchase a by-capacity license directly from NetApp.

If you have an existing subscription in a cloud marketplace, you're automatically subscribed to these new offerings as well. You can choose by-capacity charging when you deploy a new Cloud Volumes ONTAP working environment.

If you're a new customer, BlueXP will prompt you to subscribe when you create a new working environment.

- By-node licensing from all cloud provider marketplaces is deprecated and no longer available for new subscribers. This includes annual contracts and hourly subscriptions (Explore, Standard, and Premium).

This charging method is still available for existing customers who have an active subscription.

[Learn more about the licensing options for Cloud Volumes ONTAP.](#)

## 6 February 2022

### Exchange unassigned licenses

If you have an unassigned node-based license for Cloud Volumes ONTAP that you haven't used, you can now exchange the license by converting it to a Cloud Backup license, Cloud Data Sense license, or Cloud Tiering license.

This action revokes the Cloud Volumes ONTAP license and creates a dollar-equivalent license for the service with the same expiry date.

[Learn how to exchange unassigned node-based licenses.](#)

## 30 January 2022

The following changes were introduced with the 3.9.15 release of the Connector.

### Redesigned licensing selection

We redesigned the licensing selection screen when creating a new Cloud Volumes ONTAP working environment. The changes highlight the by-capacity charging methods that were introduced in July 2021 and support upcoming offerings through the cloud provider marketplaces.

### Digital Wallet update

We updated the **Digital Wallet** by consolidating Cloud Volumes ONTAP licenses in a single tab.

## 2 January 2022

The following changes were introduced with the 3.9.14 release of the Connector.

### Support for additional Azure VM types

Cloud Volumes ONTAP is now supported with the following VM types in Microsoft Azure, starting with the 9.10.1 release:

- E4ds\_v4

- E8ds\_v4
- E32ds\_v4
- E48ds\_v4

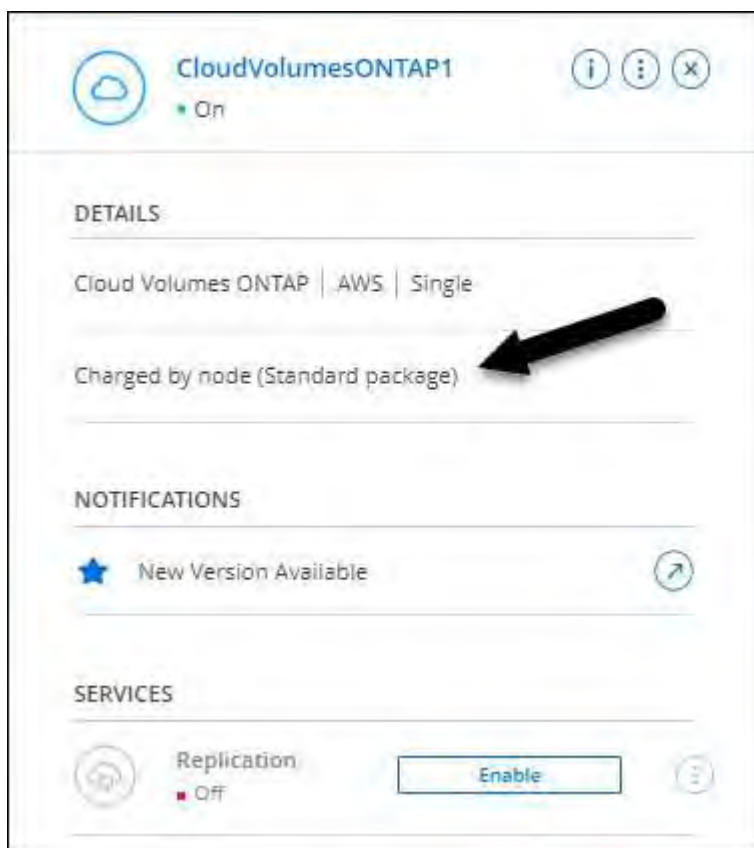
Go to the [Cloud Volumes ONTAP Release Notes](#) for more details about supported configurations.

### FlexClone charging update

If you use a [capacity-based license](#) for Cloud Volumes ONTAP, you are no longer charged for the capacity used by FlexClone volumes.

### Charging method now displayed

BlueXP now shows the charging method for each Cloud Volumes ONTAP working environment in the right panel of the Canvas.



### Choose your user name

When you create a Cloud Volumes ONTAP working environment, you now have the option to enter your preferred user name, instead of the default admin user name.

Credentials

User Name

Password

Confirm Password

### Volume creation enhancements

We made a few enhancements to volume creation:

- We redesigned the create volume wizard for ease of use.
- You can now choose a custom export policy for NFS.

Details, Protection & Tags  
  **2 Protocol**  
  3 Disk Type  
  4 Usage Profile & Tiering Policy  
  5 Review

Volumes Protocol

Select the volume's protocol:  
 NFS Protocol  
 CIFS Protocol  
 iSCSI Protocol

Access Control

Export Policy (1 rule defined)

[Manage volume's export policy](#)

## 28 November 2021

The following changes were introduced with the 3.9.13 release of the Connector.

### Cloud Volumes ONTAP 9.10.1

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.10.1.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)



## NetApp Keystone Subscriptions

You can now use Keystone Subscriptions to pay for Cloud Volumes ONTAP HA pairs.

A Keystone Subscription is a pay-as-you-grow subscription-based service that delivers a seamless hybrid cloud experience for those preferring OpEx consumption models to upfront CapEx or leasing.

A Keystone Subscription is supported with all new versions of Cloud Volumes ONTAP that you can deploy from BlueXP.

- [Learn more about NetApp Keystone Subscriptions.](#)
- [Learn how to get started with Keystone Subscriptions in BlueXP.](#)

## New AWS region support

Cloud Volumes ONTAP is now supported in the AWS Asia Pacific (Osaka) region (ap-northeast-3).

## Port reduction

Ports 8023 and 49000 are no longer open on Cloud Volumes ONTAP systems in Azure for both single node systems and HA pairs.

This change applies to *new* Cloud Volumes ONTAP systems starting with the 3.9.13 release of the Connector.

## 4 October 2021

The following changes were introduced with the 3.9.11 release of the Connector.

### Cloud Volumes ONTAP 9.10.0

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.10.0.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

### Reduced deployment time

We reduced the amount of time that it takes to deploy a Cloud Volumes ONTAP working environment in Microsoft Azure or in Google Cloud when normal write speed is enabled. The deployment time is now 3-4 minutes shorter on average.

## 2 September 2021

The following changes were introduced with the 3.9.10 release of the Connector.

### Customer-managed encryption key in Azure

Data is automatically encrypted on Cloud Volumes ONTAP in Azure using [Azure Storage Service Encryption](#) with a Microsoft-managed key. But you can now use your own customer-managed encryption key instead by completing the following steps:

1. From Azure, create a key vault and then generate a key in that vault.
2. From BlueXP, use the API to create a Cloud Volumes ONTAP working environment that uses the key.

[Learn more about these steps.](#)

## 7 July 2021

The following changes were introduced with the 3.9.8 release of the Connector.

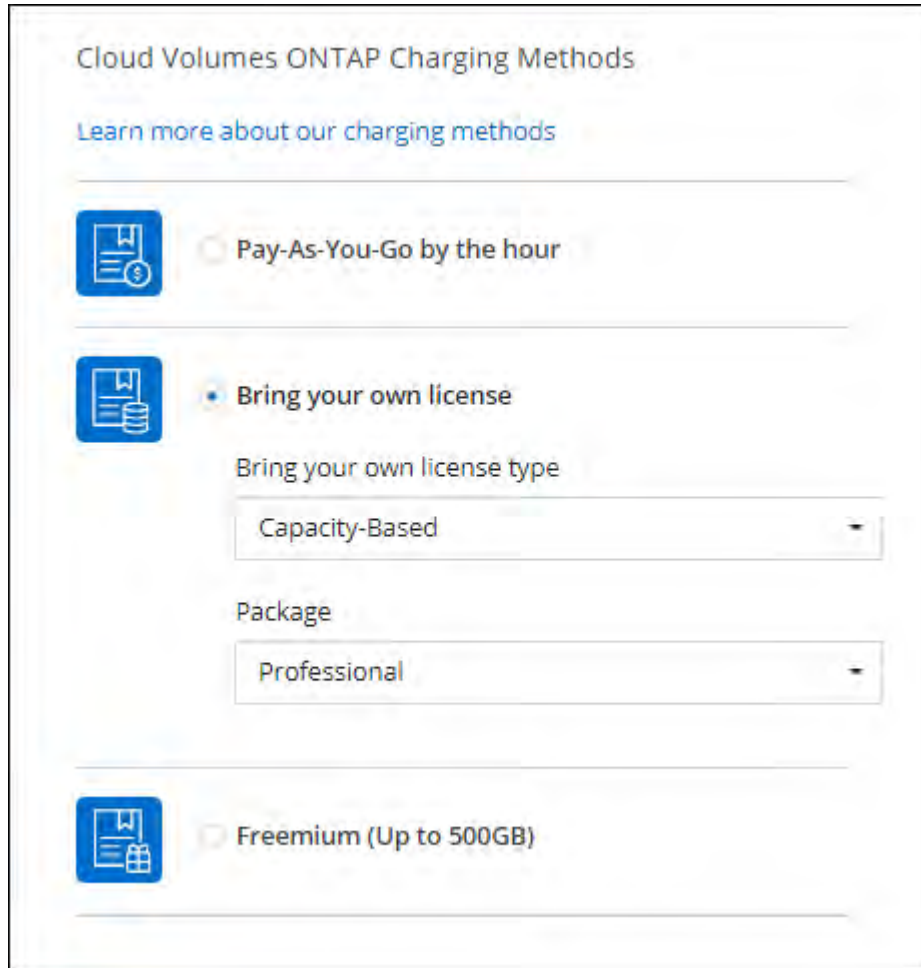
### New charging methods

New charging methods are available for Cloud Volumes ONTAP.

- **Capacity-based BYOL:** A capacity-based license enables you to pay for Cloud Volumes ONTAP per TiB of capacity. The license is associated with your NetApp account and enables you to create as multiple Cloud Volumes ONTAP systems, as long as enough capacity is available through your license. Capacity-based licensing is available in the form of a package, either *Essentials* or *Professional*.
- **Freemium offering:** Freemium enables you to use all Cloud Volumes ONTAP features free of charge from NetApp (cloud provider charges still apply). You're limited to 500 GiB of provisioned capacity per system and there's no support contract. You can have up to 10 Freemium systems.

[Learn more about these licensing options.](#)

Here's an example of the charging methods that you can choose from:



The screenshot shows a configuration page titled "Cloud Volumes ONTAP Charging Methods". At the top, there is a link "Learn more about our charging methods". Below this, there are three radio button options, each with a blue icon of a document with a dollar sign:

- Pay-As-You-Go by the hour
- Bring your own license
  - Bring your own license type
    - Capacity-Based
  - Package
    - Professional
- Freemium (Up to 500GB)

### WORM storage available for general use

Write once, read many (WORM) storage is no longer in Preview and is now available for general use with Cloud Volumes ONTAP. [Learn more about WORM storage.](#)

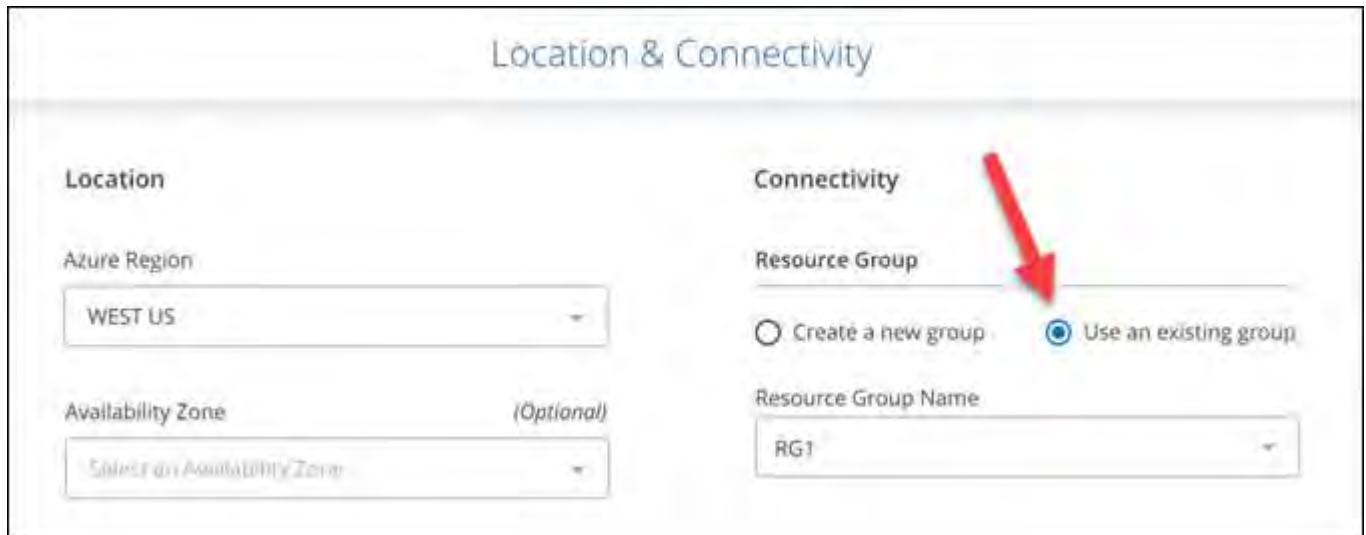
## Support for m5dn.24xlarge in AWS

Starting with the 9.9.1 release, Cloud Volumes ONTAP now supports the m5dn.24xlarge instance type with the following charging methods: PAYGO Premium, bring your own license (BYOL), and Freemium.

[View supported configurations for Cloud Volumes ONTAP in AWS.](#)

## Select existing Azure resource groups

When creating a Cloud Volumes ONTAP system in Azure, you now have the option to select an existing resource group for the VM and its associated resources.



The screenshot shows the 'Location & Connectivity' configuration page. Under the 'Connectivity' section, the 'Resource Group' field has two radio button options: 'Create a new group' and 'Use an existing group'. The 'Use an existing group' option is selected, indicated by a blue dot and a red arrow pointing to it. Below this, the 'Resource Group Name' field contains the text 'RG1'.

The following permissions enable BlueXP to remove Cloud Volumes ONTAP resources from a resource group, in case of deployment failure or deletion:

```
"Microsoft.Network/privateEndpoints/delete",  
"Microsoft.Compute/availabilitySets/delete",
```

Be sure to provide these permissions to each set of Azure credentials that you've added to BlueXP. [View the latest Connector policy for Azure.](#)

## Blob public access now disabled in Azure

As a security enhancement, BlueXP now disables **Blob public access** when creating a storage account for Cloud Volumes ONTAP.

## Azure Private Link enhancement

By default, BlueXP now enables an Azure Private Link connection on the boot diagnostics storage account for new Cloud Volumes ONTAP systems.

This means *all* storage accounts for Cloud Volumes ONTAP will now use a private link.

[Learn more about using an Azure Private Link with Cloud Volumes ONTAP.](#)

## Balanced persistent disks in Google Cloud

Starting with the 9.9.1 release, Cloud Volumes ONTAP now supports Balanced persistent disks (pd-balanced).

These SSDs balance performance and cost by providing lower IOPS per GiB.

## custom-4-16384 no longer supported in Google Cloud

The custom-4-16384 machine type is no longer supported with new Cloud Volumes ONTAP systems.

If you have an existing system running on this machine type, you can keep using it, but we recommend switching to the n2-standard-4 machine type.

[View supported configurations for Cloud Volumes ONTAP in GCP.](#)

## 30 May 2021

The following changes were introduced with the 3.9.7 release of the Connector.

### New Professional Package in AWS

A new Professional Package enables you to bundle Cloud Volumes ONTAP and Cloud Backup Service by using an annual contract from the AWS Marketplace. Payment is per TiB. This subscription doesn't enable you to back up on-premises data.

If you choose this payment option, you can provision up to 2 PiB per Cloud Volumes ONTAP system through EBS disks and tiering to S3 object storage (single node or HA).

Go to the [AWS Marketplace page](#) to view pricing details and go to the [Cloud Volumes ONTAP Release Notes](#) to learn more about this licensing option.

### Tags on EBS volumes in AWS

BlueXP now adds tags to EBS volumes when it creates a new Cloud Volumes ONTAP working environment. The tags were previously created after Cloud Volumes ONTAP was deployed.

This change can help if your organization uses service control policies (SCPs) to manage permissions.

### Minimum cooling period for auto tiering policy

If you enabled data tiering on a volume using the *auto* tiering policy, you can now adjust the minimum cooling period using the API.

[Learn how to adjust the minimum cooling period.](#)

### Enhancement to custom export policies

When you create a new NFS volume, BlueXP now displays custom export policies in ascending order, making it easier for you to find the export policy that you need.

### Deletion of old cloud snapshots

BlueXP now deletes older cloud snapshots of root and boot disks that are created when a Cloud Volumes ONTAP system is deployed and every time its powered down. Only the two most recent snapshots are retained for both the root and boot volumes.

This enhancement helps reduce cloud provider costs by removing snapshots that are no longer needed.

Note that a Connector requires a new permission to delete Azure snapshots. [View the latest Connector policy for Azure.](#)

```
"Microsoft.Compute/snapshots/delete"
```

## 24 May 2021

### Cloud Volumes ONTAP 9.9.1

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.9.1.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

## 11 Apr 2021

The following changes were introduced with the 3.9.5 release of the Connector.

### Logical space reporting

BlueXP now enables logical space reporting on the initial storage VM that it creates for Cloud Volumes ONTAP.

When space is reported logically, ONTAP reports the volume space such that all the physical space saved by the storage efficiency features are also reported as used.

### Support for gp3 disks in AWS

Cloud Volumes ONTAP now supports *General Purpose SSD (gp3)* disks, starting with the 9.7 release. gp3 disks are the lowest-cost SSDs that balance cost and performance for a broad range of workloads.

[Learn more about using gp3 disks with Cloud Volumes ONTAP.](#)

### Cold HDD disks no longer supported in AWS

Cloud Volumes ONTAP no longer supports Cold HDD (sc1) disks.

### TLS 1.2 for Azure storage accounts

When BlueXP creates storage accounts in Azure for Cloud Volumes ONTAP, the TLS version for the storage account is now version 1.2.

## 8 Mar 2021

The following changes were introduced with the 3.9.4 release of the Connector.

### Cloud Volumes ONTAP 9.9.0

BlueXP can now deploy and manage Cloud Volumes ONTAP 9.9.0.

[Learn about the new features included in this release of Cloud Volumes ONTAP.](#)

## Support for the AWS C2S environment

You can now deploy Cloud Volumes ONTAP 9.8 in the AWS Commercial Cloud Services (C2S) environment.

[Learn how to get started in C2S.](#)

## AWS encryption with customer-managed CMKs

BlueXP has always enabled you to encrypt Cloud Volumes ONTAP data using the AWS Key Management Service (KMS). Starting with Cloud Volumes ONTAP 9.9.0, data on EBS disks and data tiered to S3 are encrypted if you select a customer-managed CMK. Previously, only EBS data would be encrypted.

Note that you'll need to provide the Cloud Volumes ONTAP IAM role with access to use the CMK.

[Learn more about setting up the AWS KMS with Cloud Volumes ONTAP.](#)

## Support for Azure DoD

You can now deploy Cloud Volumes ONTAP 9.8 in the Azure Department of Defense (DoD) Impact Level 6 (IL6).

## IP address reduction in Google Cloud

We've reduced the number of IP addresses that are required for Cloud Volumes ONTAP 9.8 and later in Google Cloud. By default, one less IP address is required (we unified the intercluster LIF with the node management LIF). You also have the option to skip the creation of the SVM management LIF when using the API, which would reduce the need for an additional IP address.

[Learn more about IP address requirements in Google Cloud.](#)

## Shared VPC support in Google Cloud

When you deploy a Cloud Volumes ONTAP HA pair in Google Cloud, you can now choose shared VPCs for VPC-1, VPC-2, and VPC-3. Previously, only VPC-0 could be a shared VPC. This change is supported with Cloud Volumes ONTAP 9.8 and later.

[Learn more about Google Cloud networking requirements.](#)

## 4 Jan 2021

The following changes were introduced with the 3.9.2 release of the Connector.

### AWS Outposts

A few months ago, we announced that Cloud Volumes ONTAP had achieved the Amazon Web Services (AWS) Outposts Ready designation. Today, we're pleased to announce that we've validated BlueXP and Cloud Volumes ONTAP with AWS Outposts.

If you have an AWS Outpost, you can deploy Cloud Volumes ONTAP in that Outpost by selecting the Outpost VPC in the Working Environment wizard. The experience is the same as any other VPC that resides in AWS. Note that you will need to first deploy a Connector in your AWS Outpost.

There are a few limitations to point out:

- Only single node Cloud Volumes ONTAP systems are supported at this time

- The EC2 instances that you can use with Cloud Volumes ONTAP are limited to what's available in your Outpost
- Only General Purpose SSDs (gp2) are supported at this time

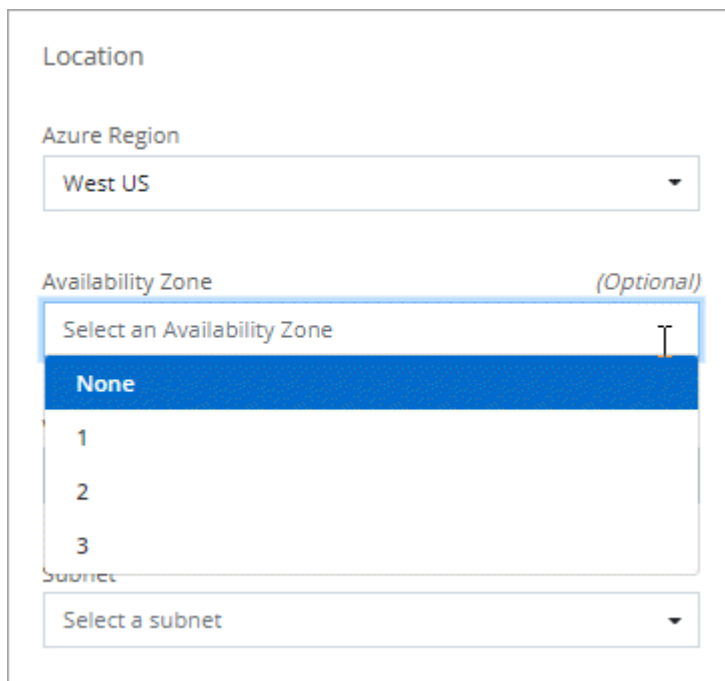
### Ultra SSD VNV RAM in supported Azure regions

Cloud Volumes ONTAP can now use an Ultra SSD as VNV RAM when you use the E32s\_v3 VM type with a single node system [in any supported Azure region](#).

VNV RAM provides better write performance.

### Choose an Availability Zone in Azure

You can now choose the Availability Zone in which you'd like to deploy a single node Cloud Volumes ONTAP system. If you don't select an AZ, BlueXP will select one for you.



The screenshot shows a configuration form for an Azure resource. It includes a 'Location' section with an 'Azure Region' dropdown set to 'West US'. Below that is an 'Availability Zone' dropdown, which is currently open and shows options: 'None' (highlighted in blue), '1', '2', and '3'. The 'Availability Zone' label is marked as '(Optional)'. At the bottom, there is a 'Subnet' dropdown set to 'Select a subnet'.

### Larger disks in Google Cloud

Cloud Volumes ONTAP now supports 64 TB disks in GCP.



The maximum system capacity with disks alone remains at 256 TB due to GCP limits.

### New machine types in Google Cloud

Cloud Volumes ONTAP now supports the following machine types:

- n2-standard-4 with the Explore license and with BYOL
- n2-standard-8 with the Standard license and with BYOL
- n2-standard-32 with the Premium license and with BYOL

## 3 Nov 2020

The following changes were introduced with the 3.9.0 release of the Connector.

### Azure Private Link for Cloud Volumes ONTAP

By default, BlueXP now enables an Azure Private Link connection between Cloud Volumes ONTAP and its associated storage accounts. A Private Link secures connections between endpoints in Azure.

- [Learn more about Azure Private Links](#)
- [Learn more about using an Azure Private Link with Cloud Volumes ONTAP](#)

## Known limitations

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

These limitations are specific to Cloud Volumes ONTAP management in BlueXP. To view limitations with the Cloud Volumes ONTAP software itself, [go to the Cloud Volumes ONTAP Release Notes](#)

### BlueXP doesn't support FlexGroup volumes creation

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



The ability to create FlexGroup volumes in BlueXP is planned for a future release.

### BlueXP doesn't support S3 with Cloud Volumes ONTAP

While Cloud Volumes ONTAP supports S3 as an option for scale-out storage, BlueXP doesn't provide any management capabilities for this feature. Using the CLI is the best practice to configure S3 client access from Cloud Volumes ONTAP. For details, refer to the [S3 Configuration Power Guide](#).

[Learn more about Cloud Volumes ONTAP support for S3 and other client protocols.](#)

### BlueXP doesn't support disaster recovery for storage VMs

BlueXP doesn't provide any setup or orchestration support for storage VM (SVM) disaster recovery. You must use ONTAP System Manager or the ONTAP CLI.

[Learn more about SVM disaster recovery.](#)

## Cloud Volumes ONTAP Release Notes

The Release Notes for Cloud Volumes ONTAP provide release-specific information. What's new in the release, supported configurations, storage limits, and any known limitations or issues that can affect product functionality.



[Go to the Cloud Volumes ONTAP Release Notes](#)

# Get started

## Learn about Cloud Volumes ONTAP

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

Cloud Volumes ONTAP is a software-only storage appliance that runs ONTAP data management software in the cloud. It provides enterprise-grade storage with the following key features:

- Storage efficiencies

Leverage built-in data deduplication, data compression, thin provisioning, and cloning to minimize storage costs.

- High availability

Ensure enterprise reliability and continuous operations in case of failures in your cloud environment.

- Data protection

Cloud Volumes ONTAP leverages SnapMirror, NetApp's industry-leading replication technology, to replicate on-premises data to the cloud so it's easy to have secondary copies available for multiple use cases.

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

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

- Data tiering

Switch between high and low-performance storage pools on-demand without taking applications offline.

- Application consistency

Ensure consistency of NetApp Snapshot copies using NetApp SnapCenter.

[Learn more about SnapCenter](#)

- Data security

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

- Privacy compliance controls

Integration with BlueXP classification helps you understand data context and identify sensitive data.

[Learn more about BlueXP classification](#)



Licenses for ONTAP features are included with Cloud Volumes ONTAP.

[View supported Cloud Volumes ONTAP configurations](#)

## Supported ONTAP versions for Cloud Volumes ONTAP deployments

BlueXP enables you to choose from several different ONTAP versions when you create a new Cloud Volumes ONTAP working environment.

Cloud Volumes ONTAP versions other than those listed here are not available for new deployments. For information on upgrade, refer to [Supported upgrade paths](#).

### AWS

#### Single node

- 9.15.1 GA
- 9.15.0 P1
- 9.14.1 GA
- 9.14.1 RC1
- 9.14.0 GA
- 9.13.1 GA
- 9.12.1 GA
- 9.12.1 RC1
- 9.12.0 P1
- 9.11.1 P3
- 9.10.1
- 9.9.1 P6
- 9.8
- 9.7 P5
- 9.5 P6

#### HA pair

- 9.15.1 GA
- 9.15.0 P1
- 9.14.1 GA
- 9.14.1 RC1
- 9.14.0 GA
- 9.13.1 GA
- 9.12.1 GA
- 9.12.1 RC1
- 9.12.0 P1
- 9.11.1 P3

- 9.10.1
- 9.9.1 P6
- 9.8
- 9.7 P5
- 9.5 P6

## **Azure**

### **Single node**

- 9.16.1 GA
- 9.15.1 GA
- 9.15.0 P1
- 9.14.1 GA
- 9.14.1 RC1
- 9.14.0 GA
- 9.13.1 GA
- 9.12.1 GA
- 9.12.1 RC1
- 9.11.1 P3
- 9.10.1 P3
- 9.9.1 P8
- 9.9.1 P7
- 9.8 P10
- 9.7 P6
- 9.5 P6

### **HA pair**

- 9.16.1 GA
- 9.15.1 GA
- 9.15.0 P1
- 9.14.1 GA
- 9.14.1 RC1
- 9.14.0 GA
- 9.13.1 GA
- 9.12.1 GA
- 9.12.1 RC1
- 9.11.1 P3
- 9.10.1 P3
- 9.9.1 P8

- 9.9.1 P7
- 9.8 P10
- 9.7 P6

## **Google Cloud**

### **Single node**

- 9.16.1 GA
- 9.15.1 GA
- 9.15.0 P1
- 9.14.1 GA
- 9.14.1 RC1
- 9.14.0 GA
- 9.13.1 GA
- 9.12.1 GA
- 9.12.1 RC1
- 9.12.0 P1
- 9.11.1 P3
- 9.10.1
- 9.9.1 P6
- 9.8
- 9.7 P5

### **HA pair**

- 9.16.1 GA
- 9.15.1 GA
- 9.15.0 P1
- 9.14.1 GA
- 9.14.1 RC1
- 9.14.0 GA
- 9.13.1 GA
- 9.12.1 GA
- 9.12.1 RC1
- 9.12.0 P1
- 9.11.1 P3
- 9.10.1
- 9.9.1 P6
- 9.8

# Get started in Amazon Web Services

## Quick start for Cloud Volumes ONTAP in AWS

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

1

### Create a Connector

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

Note that if you want to deploy Cloud Volumes ONTAP in a subnet where no internet access is available, then you need to manually install the Connector and access the BlueXP user interface that's running on that Connector. [Learn how to manually install the Connector in a location without internet access](#)

2

### Plan your configuration

BlueXP offers preconfigured packages that match your workload requirements, or you can create your own configuration. If you choose your own configuration, you should understand the options available to you. [Learn more](#).

3

### Set up your networking

- a. Ensure that your VPC and subnets will support connectivity between the Connector and Cloud Volumes ONTAP.
- b. Enable outbound internet access from the target VPC for NetApp AutoSupport.

This step isn't required if you're deploying Cloud Volumes ONTAP in a location where no internet access is available.

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

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

[Learn more about networking requirements.](#)

4

### Set up the AWS KMS

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

5

### Launch Cloud Volumes ONTAP using BlueXP

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

#### Related links

- [Create a Connector in AWS from BlueXP](#)
- [Create a Connector from the AWS Marketplace](#)
- [Install and set up a Connector on premises](#)
- [AWS permissions for the Connector](#)

## Plan your Cloud Volumes ONTAP configuration in AWS

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

### Choose a Cloud Volumes ONTAP license

Several licensing options are available for Cloud Volumes ONTAP. Each option enables you to choose a consumption model that meets your needs.

- [Learn about licensing options for Cloud Volumes ONTAP](#)
- [Learn how to set up licensing](#)

### Choose a supported region

Cloud Volumes ONTAP is supported in most AWS regions. [View the full list of supported regions.](#)

Newer AWS regions must be enabled before you can create and manage resources in those regions. [AWS documentation: Learn how to enable a region.](#)

### Choose a supported Local Zone

Selecting a Local Zone is optional. Cloud Volumes ONTAP is supported in some AWS Local Zones including Singapore. Cloud Volumes ONTAP in AWS supports only high availability (HA) mode in a single availability zone. Single node deployments are not supported.



Cloud Volumes ONTAP does not have support for data tiering and cloud tiering in AWS Local Zones. Additionally, Local Zones with instances that have not been qualified for Cloud Volumes ONTAP are not supported. An example of this is Miami, that is not available as a Local Zone, because it has only Gen6 instances that are unsupported and unqualified.

[AWS Documentation: View the full list of Local Zones.](#)

Local Zones must be enabled before you can create and manage resources in those zones.

[AWS Documentation: Getting started with AWS Local Zones.](#)

### Choose a supported instance

Cloud Volumes ONTAP supports several instance types, depending on the license type that you choose.

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

## Understand storage limits

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

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

## Size your system in AWS

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

### Instance type

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

### EBS disk type

At a high level, the differences between EBS disk types are as follows. To learn more about the use cases for EBS disks, refer to [AWS Documentation: EBS Volume Types](#).

- *General Purpose SSD (gp3)* disks are the lowest-cost SSDs that balance cost and performance for a broad range of workloads. Performance is defined in terms of IOPS and throughput. gp3 disks are supported with Cloud Volumes ONTAP 9.7 and later.

When you select a gp3 disk, BlueXP fills in default IOPS and throughput values that provide performance that is equivalent to a gp2 disk based on the selected disk size. You can increase the values to get better performance at a higher cost, but we do not support lower values because it can result in inferior performance. In short, stick with the default values or increase them. Don't lower them. [AWS Documentation: Learn more about gp3 disks and their performance](#).

Note that Cloud Volumes ONTAP supports the Amazon EBS Elastic Volumes feature with gp3 disks. [Learn more about Elastic Volumes support](#).

- *General Purpose SSD (gp2)* disks balance cost and performance for a broad range of workloads. Performance is defined in terms of IOPS.
- *Provisioned IOPS SSD (io1)* disks are for critical applications that require the highest performance at a higher cost.

Note that Cloud Volumes ONTAP supports the Amazon EBS Elastic Volumes feature with io1 disks. [Learn more about Elastic Volumes support](#).

- *Throughput Optimized HDD (st1)* disks are for frequently accessed workloads that require fast and consistent throughput at a lower price.



Data tiering to AWS S3 is not available in AWS Local Zones due to lack of connectivity.



## EBS disk size

If you choose a configuration that doesn't support the [Amazon EBS Elastic Volumes feature](#), then you need to choose an initial disk size when you launch a Cloud Volumes ONTAP system. After that, you can [let BlueXP manage a system's capacity for you](#), but if you want to [create aggregates yourself](#), be aware of the following:

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

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

As noted above, choosing a disk size is not supported with Cloud Volumes ONTAP configurations that support the Amazon EBS Elastic Volumes feature. [Learn more about Elastic Volumes support](#).

## View default system disks

In addition to the storage for user data, BlueXP also purchases cloud storage for Cloud Volumes ONTAP system data (boot data, root data, core data, and NVRAM). For planning purposes, it might help for you to review these details before you deploy Cloud Volumes ONTAP.

[View the default disks for Cloud Volumes ONTAP system data in AWS](#).



The Connector also requires a system disk. [View details about the Connector's default configuration](#).

## Prepare to deploy Cloud Volumes ONTAP in an AWS Outpost

If you have an AWS Outpost, you can deploy Cloud Volumes ONTAP in that Outpost by selecting the Outpost VPC in the Working Environment wizard. The experience is the same as any other VPC that resides in AWS. Note that you will need to first deploy a Connector in your AWS Outpost.

There are a few limitations to point out:

- Only single node Cloud Volumes ONTAP systems are supported at this time
- The EC2 instances that you can use with Cloud Volumes ONTAP are limited to what's available in your Outpost
- Only General Purpose SSDs (gp2) are supported at this time

## Collect networking information

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

## Single node or HA pair in a single AZ

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

#### HA pair in multiple AZs

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

#### Choose a write speed

BlueXP enables you to choose a write speed setting for Cloud Volumes ONTAP. Before you choose a write speed, you should understand the differences between the normal and high settings and risks and recommendations when using high write speed. [Learn more about write speed.](#)

#### Choose a volume usage profile

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

NetApp storage efficiency features provide the following benefits:

## Thin provisioning

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

## Deduplication

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

## Compression

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

## Set up your networking

### Set up AWS networking for Cloud Volumes ONTAP

BlueXP handles the set up of networking components for Cloud Volumes ONTAP, such as IP addresses, netmasks, and routes. You need to make sure that outbound internet access is available, that enough private IP addresses are available, that the right connections are in place, and more.

#### General requirements

Ensure that you have fulfilled the following requirements in AWS.

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

Cloud Volumes ONTAP systems require outbound internet access for accessing external endpoints for various functions. Cloud Volumes ONTAP can't operate properly if these endpoints are blocked in environments with strict security requirements.

The BlueXP Connector also contacts several endpoints for day-to-day operations, as well as the BlueXP web-based console. For information about the BlueXP endpoints, refer to [View endpoints contacted from the Connector](#) and [Prepare networking for using the BlueXP console](#).

#### Cloud Volumes ONTAP endpoints

Cloud Volumes ONTAP uses these endpoints to communicate with various services.

Endpoints	Applicable for	Purpose	BlueXP deployment modes	Impact if endpoint is not available
https://netapp-cloud-account.auth0.com	Authentication	Used for BlueXP authentication.	Standard and restricted modes.	User authentication fails and the following services remain unavailable: <ul style="list-style-type: none"> <li>• Cloud Volumes ONTAP services</li> <li>• ONTAP services</li> <li>• Protocols and proxy services</li> </ul>
https://cloudmanager.cloud.netapp.com/tenancy	Tenancy	Used to retrieve the Cloud Volumes ONTAP resources from BlueXP tenancy to authorize resources and users.	Standard and restricted modes.	Cloud Volumes ONTAP resources and the users are not authorized.
https://support.netapp.com/aods/asupmessage https://support.netapp.com/asupprod/post/1.0/postAsup	AutoSupport	Used to send AutoSupport telemetry data to NetApp support.	Standard and restricted modes.	AutoSupport information remains undelivered.
The exact commercial endpoint for AWS service (suffixed with amazonaws.com) depends on the AWS region that you are using. Refer to <a href="#">AWS documentation for details</a> .	<ul style="list-style-type: none"> <li>• CloudFormation</li> <li>• Elastic Compute Cloud (EC2)</li> <li>• Identity and Access Management (IAM)</li> <li>• Key Management Service (KMS)</li> <li>• Security Token Service (STS)</li> <li>• Simple Storage Service (S3)</li> </ul>	Communication with AWS services.	Standard and private modes.	Cloud Volumes ONTAP cannot communicate with AWS service to perform specific BlueXP operations in AWS.

Endpoints	Applicable for	Purpose	BlueXP deployment modes	Impact if endpoint is not available
The exact government endpoint for AWS service depends on the AWS region that you are using. The endpoints are suffixed with <code>amazonaws.com</code> and <code>c2s.ic.gov</code> . Refer to <a href="#">AWS SDK</a> and <a href="#">AWS Documentation</a> for more information.	<ul style="list-style-type: none"> <li>• CloudFormation</li> <li>• Elastic Compute Cloud (EC2)</li> <li>• Identity and Access Management (IAM)</li> <li>• Key Management Service (KMS)</li> <li>• Security Token Service (STS)</li> <li>• Simple Storage Service (S3)</li> </ul>	Communication with AWS services.	Restricted mode.	Cloud Volumes ONTAP cannot communicate with AWS service to perform specific BlueXP operations in AWS.

### Outbound internet access for the HA mediator

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

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

### Network configurations to support Connector proxy

You can use the proxy servers configured for the BlueXP Connector to enable outbound internet access from Cloud Volumes ONTAP. BlueXP supports two types of proxies:

- **Explicit proxy:** The outbound traffic from Cloud Volumes ONTAP uses the HTTP address of the proxy server specified during the Connector proxy configuration. The Connector administrator might also have configured user credentials and root CA certificates for additional authentication. If a root CA certificate is available for the explicit proxy, make sure to obtain and upload the same certificate to your Cloud Volumes ONTAP working environment using the [ONTAP CLI: security certificate install](#) command.
- **Transparent proxy:** The network is configured to automatically route outbound traffic from Cloud Volumes ONTAP through the Connector proxy. When setting up a transparent proxy, the Connector administrator needs to provide only a root CA certificate for connectivity from Cloud Volumes ONTAP, not the HTTP address of the proxy server. Make sure that you obtain and upload the same root CA certificate to your Cloud Volumes ONTAP working environment using the [ONTAP CLI: security certificate install](#) command.

For information about configuring proxy servers for the BlueXP Connector, refer to the [Configure a Connector to use a proxy server](#).

### Private IP addresses

BlueXP automatically allocates the required number of private IP addresses to Cloud Volumes ONTAP. You need to ensure that your networking has enough private IP addresses available.

The number of LIFs that BlueXP allocates for Cloud Volumes ONTAP depends on whether you deploy a single node system or an HA pair. A LIF is an IP address associated with a physical port.

### IP addresses for a single node system

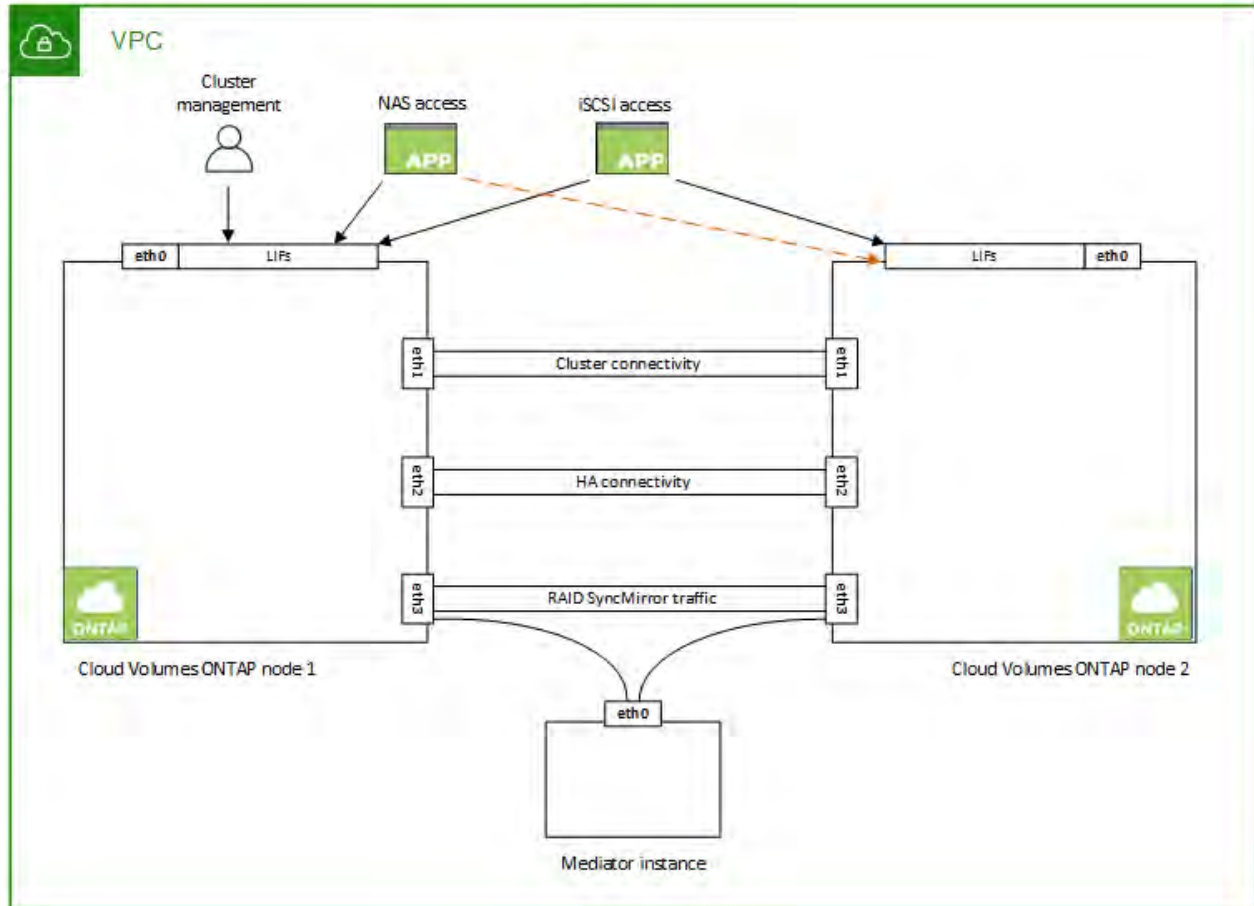
BlueXP allocates 6 IP addresses to a single node system.

The following table provides details about the LIFs that are associated with each private IP address.

LIF	Purpose
Cluster management	Administrative management of the entire cluster (HA pair).
Node management	Administrative management of a node.
Intercluster	Cross-cluster communication, backup, and replication.
NAS data	Client access over NAS protocols.
iSCSI data	Client access over the iSCSI protocol. Also used by the system for other important networking workflows. This LIF is required and should not be deleted.
Storage VM management	A storage VM management LIF is used with management tools like SnapCenter.

### IP addresses for HA pairs

HA pairs require more IP addresses than a single node system does. These IP addresses are spread across different ethernet interfaces, as shown in the following image:



The number of private IP addresses required for an HA pair depends on which deployment model you choose. An HA pair deployed in a *single* AWS Availability Zone (AZ) requires 15 private IP addresses, while an HA pair deployed in *multiple* AZs requires 13 private IP addresses.

The following tables provide details about the LIFs that are associated with each private IP address.

**Table 1. LIFs for HA pairs in a single AZ**

LIF	Interface	Node	Purpose
Cluster management	eth0	node 1	Administrative management of the entire cluster (HA pair).
Node management	eth0	node 1 and node 2	Administrative management of a node.
Intercluster	eth0	node 1 and node 2	Cross-cluster communication, backup, and replication.
NAS data	eth0	node 1	Client access over NAS protocols.
iSCSI data	eth0	node 1 and node 2	Client access over the iSCSI protocol. Also used by the system for other important networking workflows. These LIFs are required and should not be deleted.

LIF	Interface	Node	Purpose
Cluster connectivity	eth1	node 1 and node 2	Enables the nodes to communicate with each other and to move data within the cluster.
HA connectivity	eth2	node 1 and node 2	Communication between the two nodes in case of failover.
RSM iSCSI traffic	eth3	node 1 and node 2	RAID SyncMirror iSCSI traffic, as well as communication between the two Cloud Volumes ONTAP nodes and the mediator.
Mediator	eth0	Mediator	A communication channel between the nodes and the mediator to assist in storage takeover and giveback processes.

**Table 2. LIFs for HA pairs in multiple AZs**

LIF	Interface	Node	Purpose
Node management	eth0	node 1 and node 2	Administrative management of a node.
Intercluster	eth0	node 1 and node 2	Cross-cluster communication, backup, and replication.
iSCSI data	eth0	node 1 and node 2	Client access over the iSCSI protocol. These LIFs also manage the migration of floating IP addresses between nodes. These LIFs are required and should not be deleted.
Cluster connectivity	eth1	node 1 and node 2	Enables the nodes to communicate with each other and to move data within the cluster.
HA connectivity	eth2	node 1 and node 2	Communication between the two nodes in case of failover.
RSM iSCSI traffic	eth3	node 1 and node 2	RAID SyncMirror iSCSI traffic, as well as communication between the two Cloud Volumes ONTAP nodes and the mediator.
Mediator	eth0	Mediator	A communication channel between the nodes and the mediator to assist in storage takeover and giveback processes.



When deployed in multiple Availability Zones, several LIFs are associated with [floating IP addresses](#), which don't count against the AWS private IP limit.

## Security groups

You don't need to create security groups because BlueXP does that for you. If you need to use your own, refer to [Security group rules](#).



Looking for information about the Connector? [View security group rules for the Connector](#)



## Connection for data tiering

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

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

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

## Connections to ONTAP systems

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

## DNS and Active Directory for CIFS

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

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

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

## VPC sharing

Starting with the 9.11.1 release, Cloud Volumes ONTAP HA pairs are supported in AWS with VPC sharing. VPC sharing enables your organization to share subnets with other AWS accounts. To use this configuration, you must set up your AWS environment and then deploy the HA pair using the API.

[Learn how to deploy an HA pair in a shared subnet.](#)

## Requirements for HA pairs in multiple AZs

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

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

## Availability Zones

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

A subnet should be available in each Availability Zone.

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

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

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



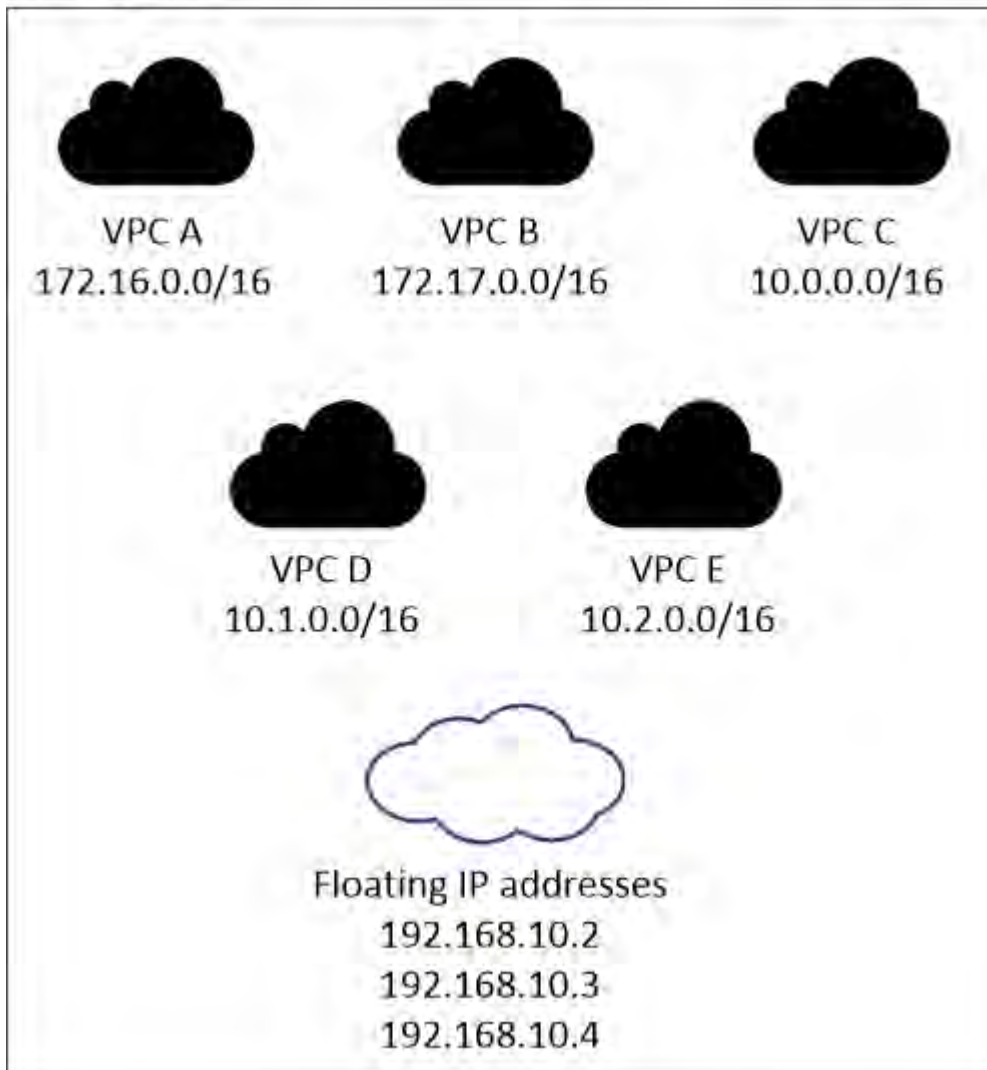
A floating IP address is required for the SVM management LIF if you use SnapDrive for Windows or SnapCenter with the HA pair.

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

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

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

### AWS region





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

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

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

### Route tables

After you specify the floating IP addresses in BlueXP, you are then prompted to select the route tables that should include routes to the floating IP addresses. This enables client access to the HA pair.

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

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

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

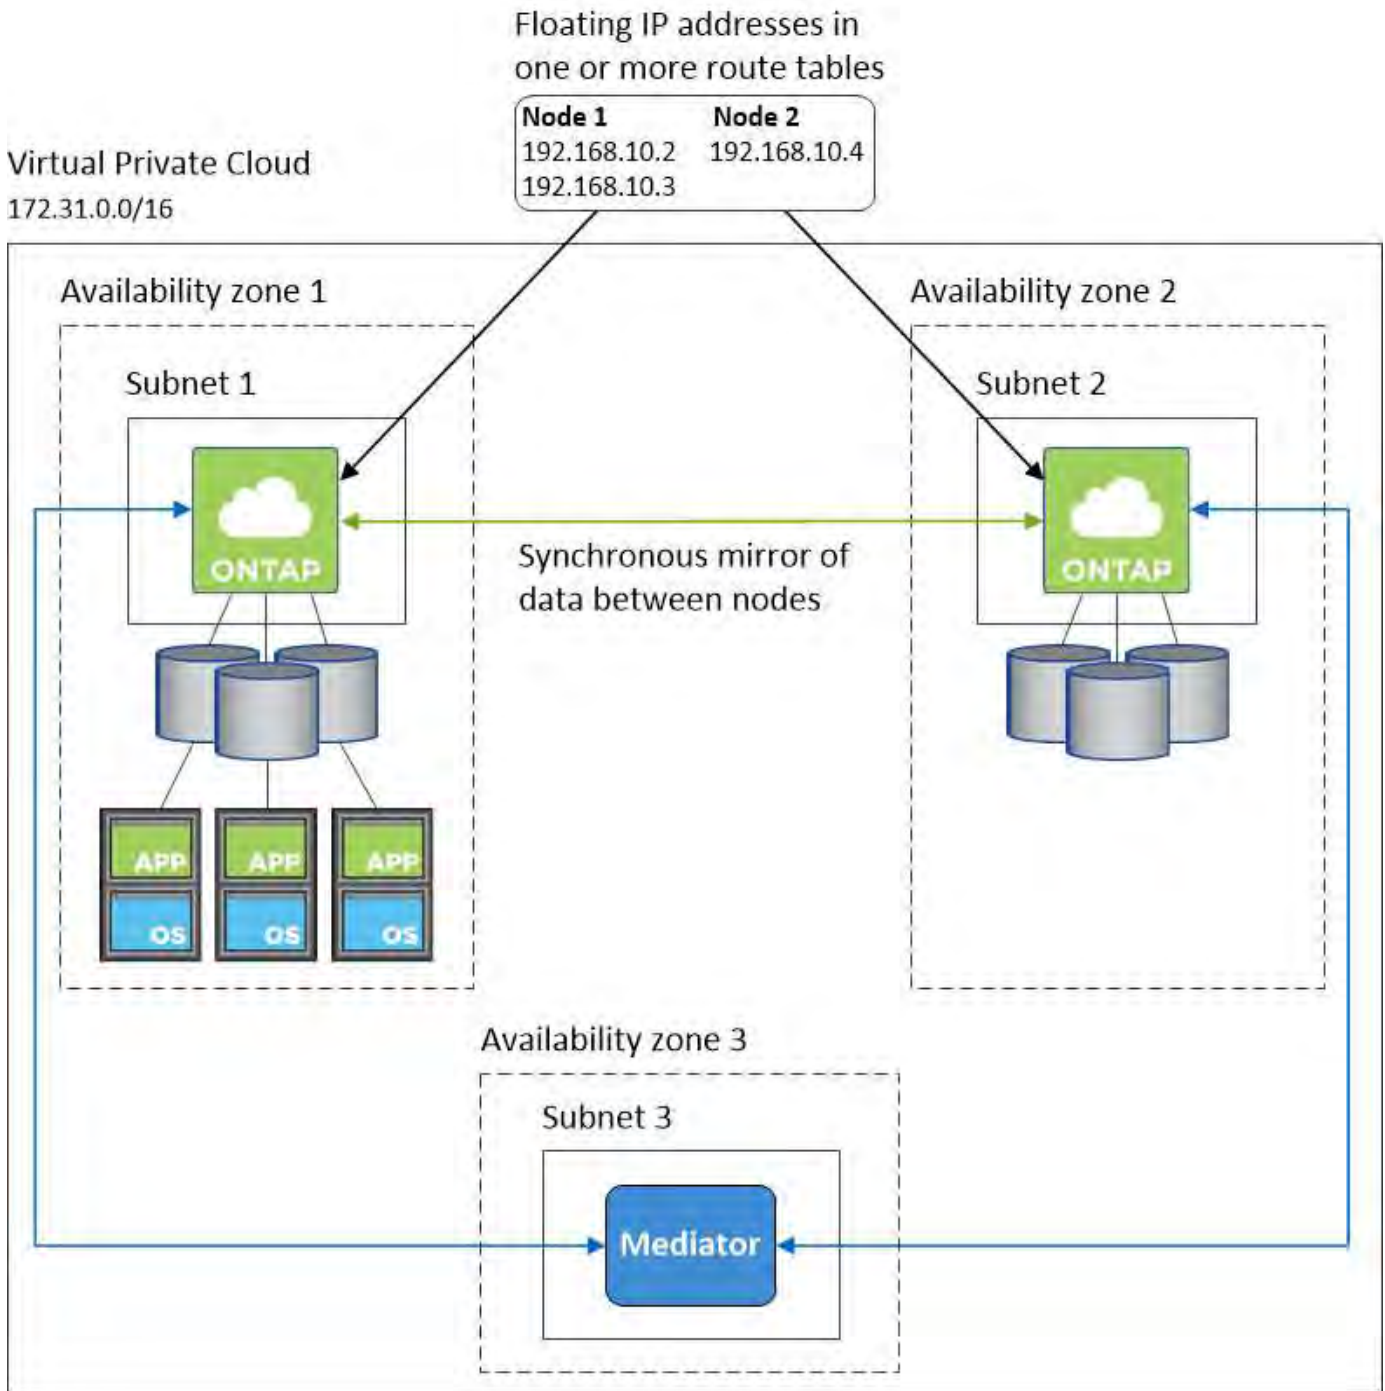
### Connection to NetApp management tools

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

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

### Example HA configuration

The following image illustrates the networking components specific to an HA pair in multiple AZs: three Availability Zones, three subnets, floating IP addresses, and a route table.



### Requirements for the Connector

If you haven't created a Connector yet, you should review networking requirements for the Connector as well.

- [View networking requirements for the Connector](#)
- [Security group rules in AWS](#)

### Related topics

- [Verify AutoSupport setup for Cloud Volumes ONTAP](#)
- [Learn about ONTAP internal ports.](#)

## Set up an AWS transit gateway for Cloud Volumes ONTAP HA pairs

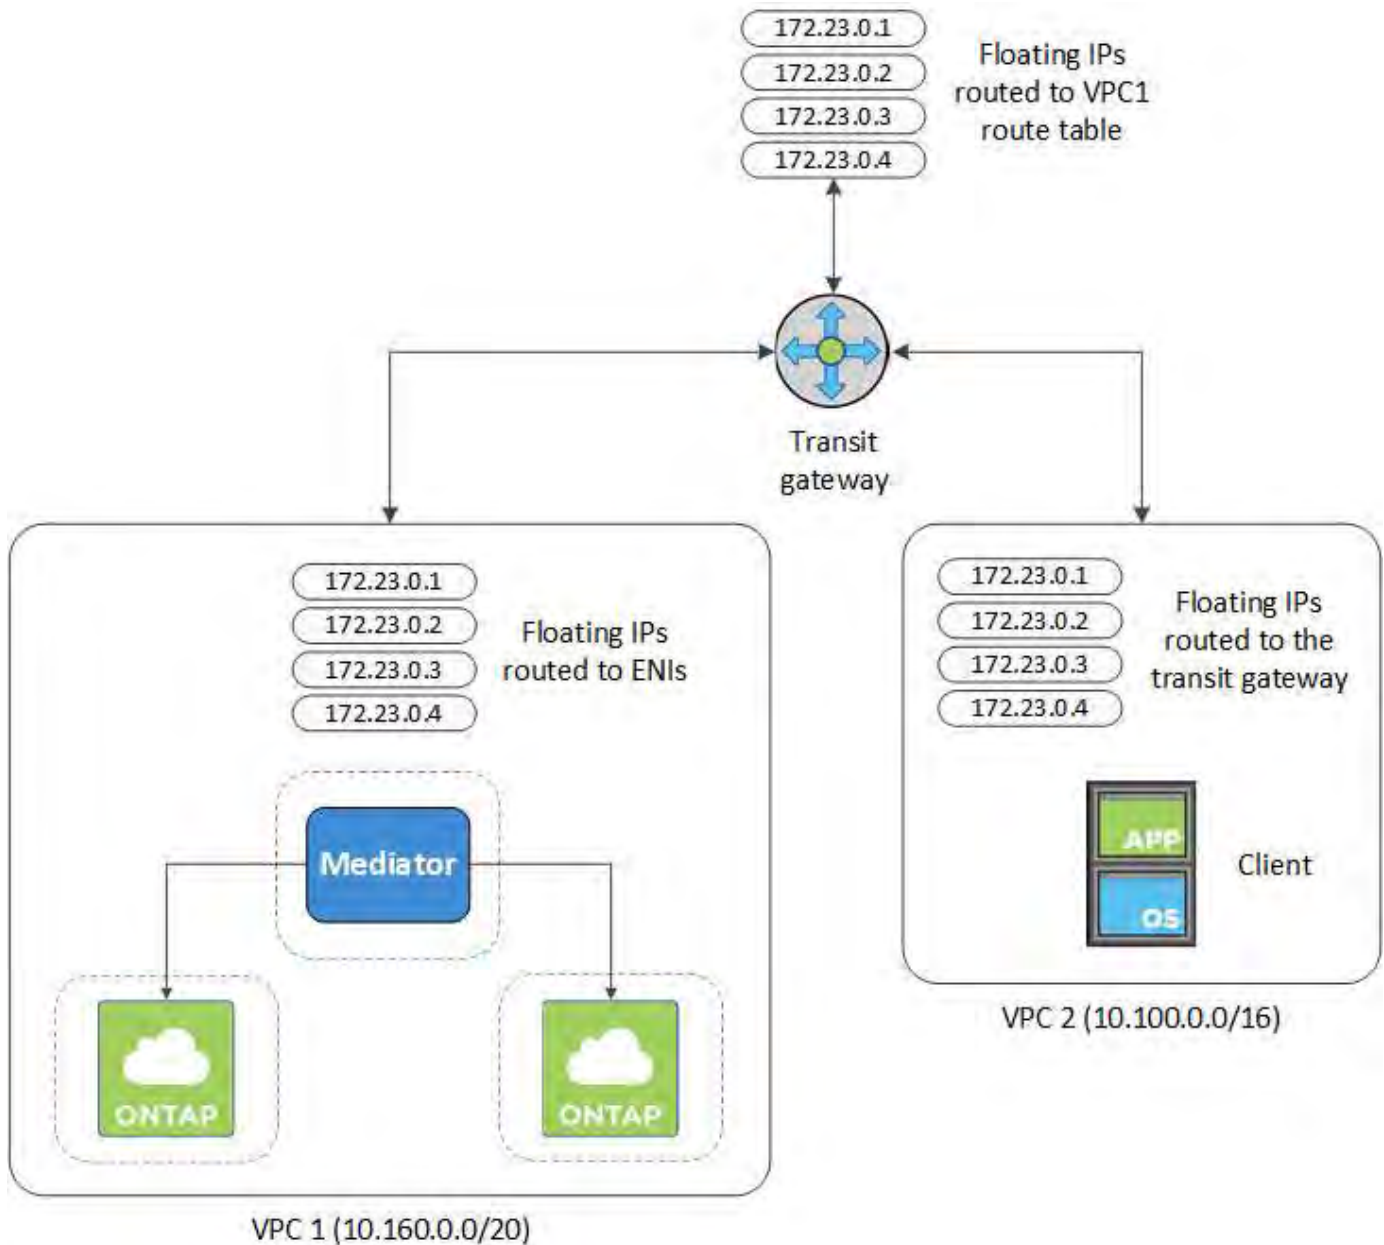
Set up an AWS transit gateway to enable access to an HA pair's [floating IP addresses](#) from outside the VPC where the HA pair resides.

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

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

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

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



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

### Steps

1. [Create a transit gateway and attach the VPCs to the gateway.](#)
2. Associate the VPCs with the transit gateway route table.
  - a. In the **VPC** service, click **Transit Gateway Route Tables**.
  - b. Select the route table.
  - c. Click **Associations** and then select **Create association**.
  - d. Choose the attachments (the VPCs) to associate and then click **Create association**.
3. Create routes in the transit gateway's route table by specifying the HA pair's floating IP addresses.

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



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

### Auto failover

Cluster Management : 172.23.0.1

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

### Access

SVM Management : 172.23.0.4

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

Transit Gateway Route Table: tgw-rtb-0ea8ee291c7aeddd3

Details Associations Propagations **Routes** Tags

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

Create route Replace route Delete route

Filter by attributes or search by keyword

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

**Floating IP Addresses**

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

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

Route Table: rb-0569a1bd740ed03f

Summary Routes Subnet Associations Route Propagation Tags

Edit routes

View All routes

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

VPC1  
Floating IP Addresses

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

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

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

Summary Routes Subnet Associations Route Propagation Tags

Edit routes

View All routes

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

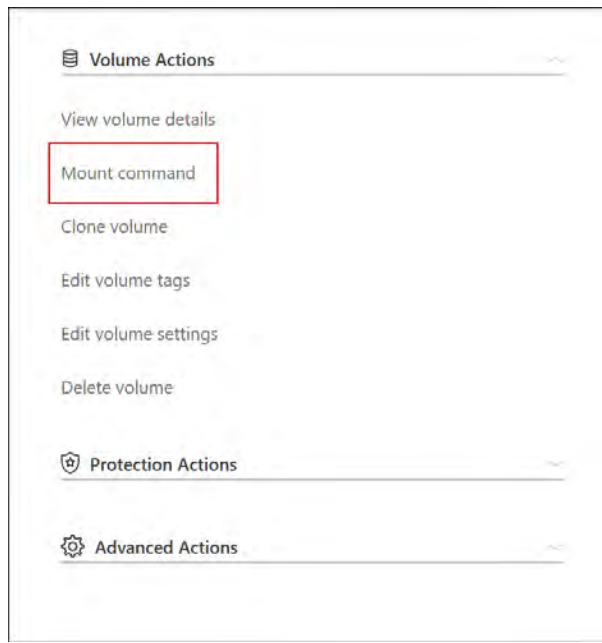
VPC2  
Floating IP Addresses

- Update the security groups settings to All traffic for the VPC.
  - Under Virtual Private Cloud, click **Subnets**.
  - Click the **Route table** tab, select the desired environment for one of the floating IP addresses for an HA pair.
  - Click **Security groups**.
  - Select **Edit Inbound Rules**.
  - Click **Add rule**.
    - Under Type, select **All traffic**, and then select the VPC IP address.
  - Click **Save Rules** to apply the changes.
- Mount volumes to clients using the floating IP address.

You can find the correct IP address in BlueXP through the **Mount Command** option under the Manage



## Volumes panel in BlueXP.



8. If you're mounting an NFS volume, configure the export policy to match the subnet of the client VPC.

[Learn how to edit a volume.](#)

### Related links

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

### Deploy Cloud Volumes ONTAP HA pairs in an AWS shared subnet

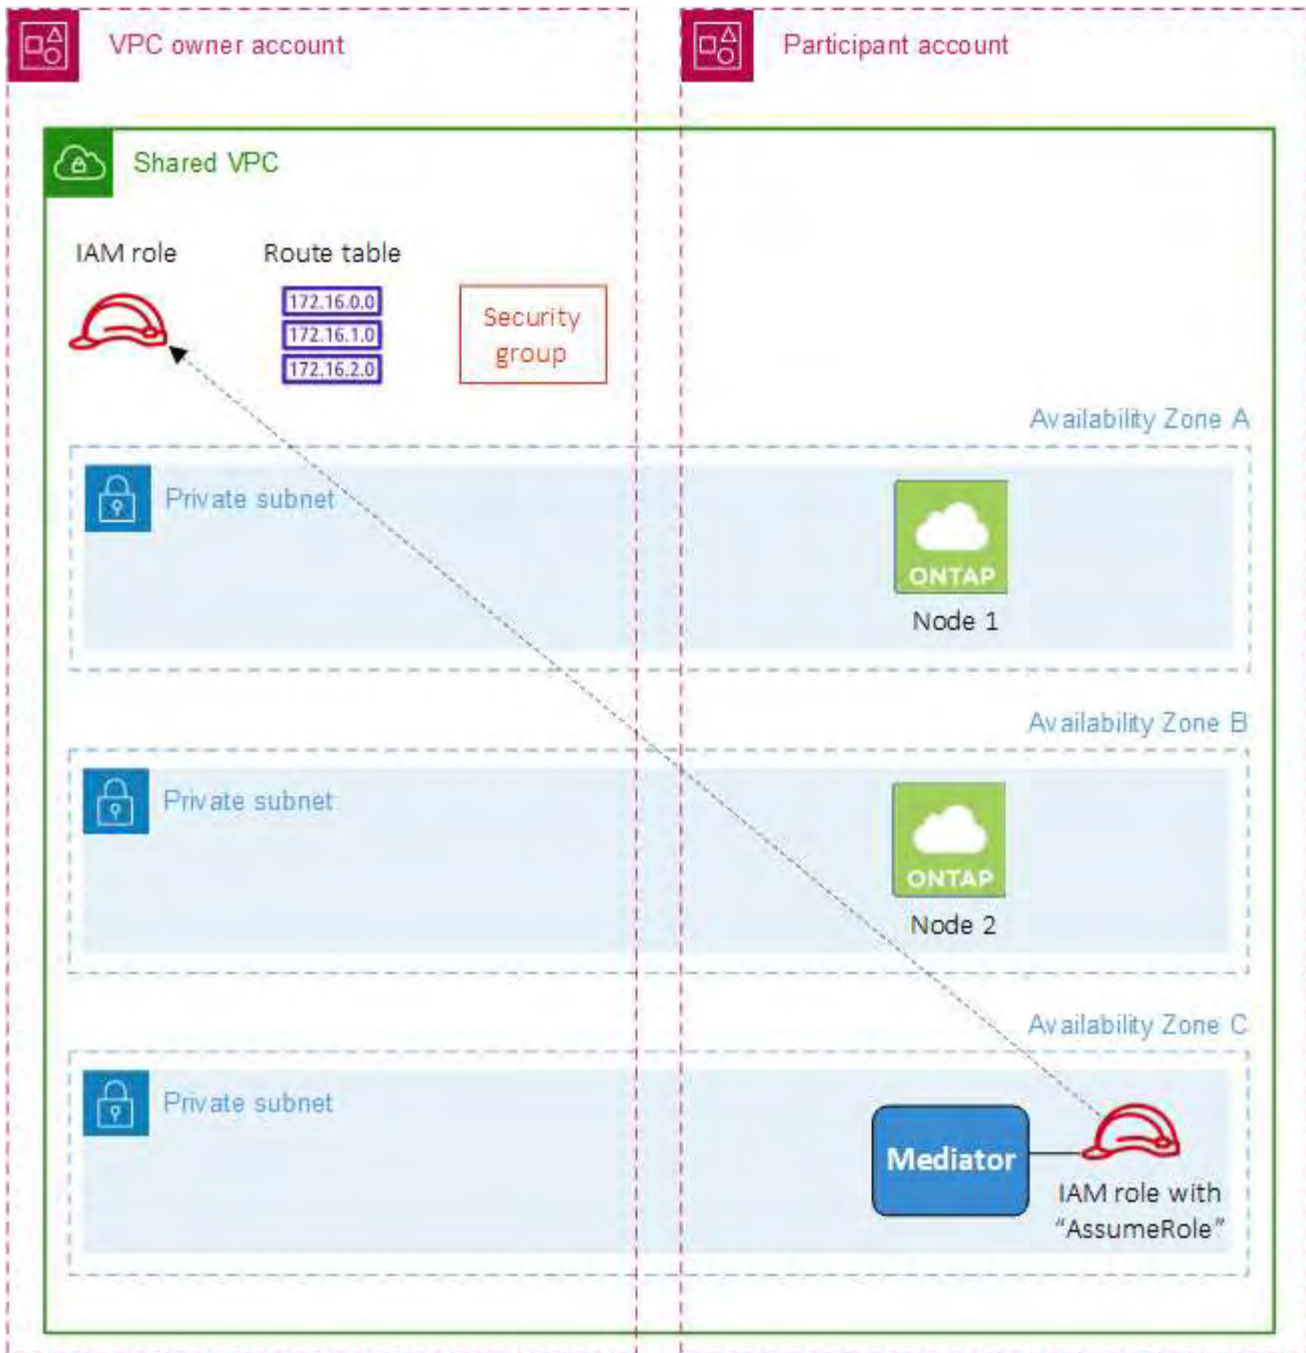
Starting with the 9.11.1 release, Cloud Volumes ONTAP HA pairs are supported in AWS with VPC sharing. VPC sharing enables your organization to share subnets with other AWS accounts. To use this configuration, you must set up your AWS environment and then deploy the HA pair using the API.

With [VPC sharing](#), a Cloud Volumes ONTAP HA configuration is spread across two accounts:

- The VPC owner account, which owns the networking (the VPC, subnets, route tables, and Cloud Volumes ONTAP security group)
- The participant account, where the EC2 instances are deployed in shared subnets (this includes the two HA nodes and the mediator)

In the case of a Cloud Volumes ONTAP HA configuration that is deployed across multiple Availability Zones, the HA mediator needs specific permissions to write to the route tables in the VPC owner account. You need to provide those permissions by setting up an IAM role that the mediator can assume.

The following image shows the components involved this deployment:



As described in the steps below, you'll need to share the subnets with the participant account, and then create the IAM role and security group in the VPC owner account.

When you create the Cloud Volumes ONTAP working environment, BlueXP automatically creates and attaches an IAM role to the mediator. This role assumes the IAM role that you created in the VPC owner account in order to make changes to the route tables associated with the HA pair.

### Steps

1. Share the subnets in the VPC owner account with the participant account.

This step is required to deploy the HA pair in shared subnets.

[AWS documentation: Share a subnet](#)

2. In the VPC owner account, create a security group for Cloud Volumes ONTAP.

[Refer to the security group rules for Cloud Volumes ONTAP](#). Note that you don't need to create a security group for the HA mediator. BlueXP does that for you.

3. In the VPC owner account, create an IAM role that includes the following permissions:

```
  "Action": [
    "ec2:AssignPrivateIpAddresses",
    "ec2:CreateRoute",
    "ec2>DeleteRoute",
    "ec2:DescribeNetworkInterfaces",
    "ec2:DescribeRouteTables",
    "ec2:DescribeVpcs",
    "ec2:ReplaceRoute",
    "ec2:UnassignPrivateIpAddresses"
```

4. Use the BlueXP API to create a new Cloud Volumes ONTAP working environment.

Note that you must specify the following fields:

- "securityGroupId"

The "securityGroupId" field should specify the security group that you created in the VPC owner account (see step 2 above).

- "assumeRoleArn" in the "haParams" object

The "assumeRoleArn" field should include the ARN of the IAM role that you created in the VPC owner account (see step 3 above).

For example:

```
"haParams": {
  "assumeRoleArn":
  "arn:aws:iam::642991768967:role/mediator_role_assume_fromdev"
}
```

[Learn about the Cloud Volumes ONTAP API](#)

### Configure placement group creation for Cloud Volumes ONTAP HA pairs in AWS single AZs

Cloud Volumes ONTAP high-availability (HA) deployments in AWS single availability Zone (AZ) can fail and roll back if the creation of the placement group fails. Creation of the placement group also fails and the deployment rolls back if the Cloud Volumes ONTAP node and mediator instance are not available. To avoid this, you can modify the configuration to allow the deployment to finish even if the placement group creation fails.

On bypassing the rollback process, the Cloud Volumes ONTAP deployment process completes successfully, and notifies you that the placement group creation is incomplete.

### Steps

1. Use SSH to connect to the Connector host and log in.
2. Navigate to `/opt/application/netapp/cloudmanager/docker_occm/data`.
3. Edit `app.conf` by changing the value of the `rollback-on-placement-group-failure` parameter to `false`. The default value of this parameter is `true`.

```
{
  "occm" : {
    "aws" : {
      "rollback-on-placement-group-failure" : false
    }
  }
}
```

4. Save the file and log off the Connector. You don't need to restart the Connector.

### AWS security group inbound and outbound rules for Cloud Volumes ONTAP

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

#### Rules for Cloud Volumes ONTAP

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

#### Inbound rules

When you create a working environment and choose a predefined security group, you can choose to allow traffic within one of the following:

- **Selected VPC only:** the source for inbound traffic is the subnet range of the VPC for the Cloud Volumes ONTAP system and the subnet range of the VPC where the Connector resides. This is the recommended option.
- **All VPCs:** the source for inbound traffic is the 0.0.0.0/0 IP range.

Protocol	Port	Purpose
All ICMP	All	Pinging the instance
HTTP	80	HTTP access to the ONTAP System Manager web console using the IP address of the cluster management LIF
HTTPS	443	Connectivity with the Connector and HTTPS access to the ONTAP System Manager web console using the IP address of the cluster management LIF

Protocol	Port	Purpose
SSH	22	SSH access to the IP address of the cluster management LIF or a node management LIF
TCP	111	Remote procedure call for NFS
TCP	139	NetBIOS service session for CIFS
TCP	161-162	Simple network management protocol
TCP	445	Microsoft SMB/CIFS over TCP with NetBIOS framing
TCP	635	NFS mount
TCP	749	Kerberos
TCP	2049	NFS server daemon
TCP	3260	iSCSI access through the iSCSI data LIF
TCP	4045	NFS lock daemon
TCP	4046	Network status monitor for NFS
TCP	10000	Backup using NDMP
TCP	11104	Management of intercluster communication sessions for SnapMirror
TCP	11105	SnapMirror data transfer using intercluster LIFs
UDP	111	Remote procedure call for NFS
UDP	161-162	Simple network management protocol
UDP	635	NFS mount
UDP	2049	NFS server daemon
UDP	4045	NFS lock daemon
UDP	4046	Network status monitor for NFS
UDP	4049	NFS rquotad protocol

## Outbound rules

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

### Basic outbound rules

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

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

## Advanced outbound rules

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



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

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

Service	Protocol	Port	Source	Destination	Purpose
AutoSupport	HTTPS	443	Node management LIF	support.netapp.com	AutoSupport (HTTPS is the default)
	HTTP	80	Node management LIF	support.netapp.com	AutoSupport (only if the transport protocol is changed from HTTPS to HTTP)
	TCP	3128	Node management LIF	Connector	Sending AutoSupport messages through a proxy server on the Connector, if an outbound internet connection isn't available
Backup to S3	TCP	5010	Intercluster LIF	Backup endpoint or restore endpoint	Back up and restore operations for the Backup to S3 feature
Cluster	All traffic	All traffic	All LIFs on one node	All LIFs on the other node	Intercluster communications (Cloud Volumes ONTAP HA only)
	TCP	3000	Node management LIF	HA mediator	ZAPI calls (Cloud Volumes ONTAP HA only)
	ICMP	1	Node management LIF	HA mediator	Keep alive (Cloud Volumes ONTAP HA only)
Configuration backups	HTTP	80	Node management LIF	http://<connector-IP-address>/occm/offbo xconfig	Send configuration backups to the Connector. <a href="#">ONTAP documentation</a>
DHCP	UDP	68	Node management LIF	DHCP	DHCP client for first-time setup
DHCPS	UDP	67	Node management LIF	DHCP	DHCP server
DNS	UDP	53	Node management LIF and data LIF (NFS, CIFS)	DNS	DNS
NDMP	TCP	1860 0–18 699	Node management LIF	Destination servers	NDMP copy
SMTP	TCP	25	Node management LIF	Mail server	SMTP alerts, can be used for AutoSupport
SNMP	TCP	161	Node management LIF	Monitor server	Monitoring by SNMP traps
	UDP	161	Node management LIF	Monitor server	Monitoring by SNMP traps
	TCP	162	Node management LIF	Monitor server	Monitoring by SNMP traps
	UDP	162	Node management LIF	Monitor server	Monitoring by SNMP traps



Service	Protocol	Port	Source	Destination	Purpose
SnapMirror	TCP	11104	Intercluster LIF	ONTAP intercluster LIFs	Management of intercluster communication sessions for SnapMirror
	TCP	11105	Intercluster LIF	ONTAP intercluster LIFs	SnapMirror data transfer
Syslog	UDP	514	Node management LIF	Syslog server	Syslog forward messages

### Rules for the HA mediator external security group

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

#### Inbound rules

The predefined security group for the HA mediator includes the following inbound rule.

Protocol	Port	Source	Purpose
TCP	3000	CIDR of the Connector	RESTful API access from the Connector

#### Outbound rules

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

#### Basic outbound rules

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

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

#### Advanced outbound rules

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

Protocol	Port	Destination	Purpose
HTTP	80	IP address of the Connector on AWS EC2 instance	Download upgrades for the mediator
HTTPS	443	ec2.amazonaws.com	Assist with storage failover
UDP	53	ec2.amazonaws.com	Assist with storage failover



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

### Rules for the HA configuration internal security group

The predefined internal security group for a Cloud Volumes ONTAP HA configuration includes the following rules. This security group enables communication between the HA nodes and between the mediator and the nodes.

BlueXP always creates this security group. You do not have the option to use your own.

### Inbound rules

The predefined security group includes the following inbound rules.

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

### Outbound rules

The predefined security group includes the following outbound rules.

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

### Rules for the Connector

[View security group rules for the Connector](#)

## Set up Cloud Volumes ONTAP to use a customer-managed key in AWS

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

### Steps

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

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

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

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

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

[AWS Documentation: Editing Keys](#)

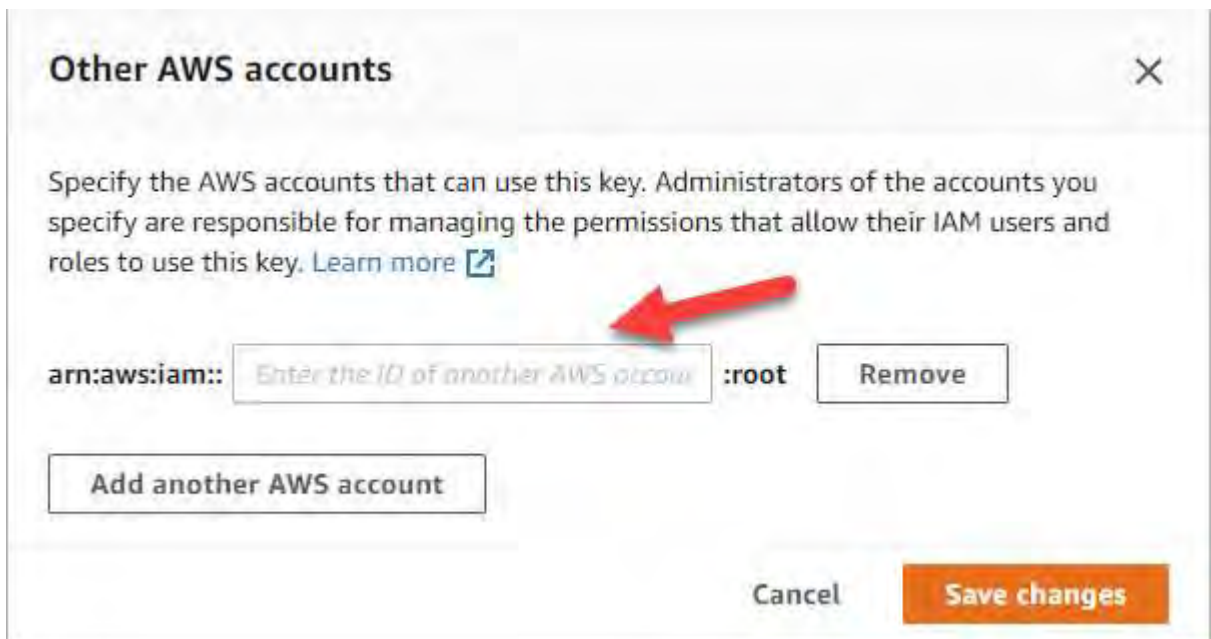
3. If the CMK is in a different AWS account, complete the following steps:

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

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

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

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



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

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

```

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

```

For additional details about this process, refer to the [AWS Documentation: Allowing users in other accounts to use a KMS key](#).

4. If you are using a customer-managed CMK, modify the key policy for the CMK by adding the Cloud Volumes ONTAP IAM role as a *key user*.

This step is required if you enabled data tiering on Cloud Volumes ONTAP and want to encrypt the data

stored in the S3 bucket.

You'll need to perform this step *after* you deploy Cloud Volumes ONTAP because the IAM role is created when you create a working environment. (Of course, you do have the option to use an existing Cloud Volumes ONTAP IAM role, so it's possible to perform this step before.)

[AWS Documentation: Editing Keys](#)

## Set up AWS IAM roles for Cloud Volumes ONTAP nodes

IAM roles with the required permissions must be attached to each Cloud Volumes ONTAP node. The same is true for the HA mediator. It's easiest to let BlueXP create the IAM roles for you, but you can use your own roles.

This task is optional. When you create a Cloud Volumes ONTAP working environment, the default option is to let BlueXP create the IAM roles for you. If your business's security policies require you to create the IAM roles yourself, then follow the steps below.



Providing your own IAM role is required in AWS Secret Cloud. [Learn how to deploy Cloud Volumes ONTAP in C2S.](#)

### Steps

1. Go to the AWS IAM console.
2. Create IAM policies that include the following permissions:
  - Base policy for Cloud Volumes ONTAP nodes

## Standard regions

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": "s3:ListAllMyBuckets",
    "Resource": "arn:aws:s3:::*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:ListBucket",
      "s3:GetBucketLocation"
    ],
    "Resource": "arn:aws:s3:::fabric-pool-*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:GetObject",
      "s3:PutObject",
      "s3>DeleteObject"
    ],
    "Resource": "arn:aws:s3:::fabric-pool-*",
    "Effect": "Allow"
  }
]
}
```

## GovCloud (US) regions

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": "s3:ListAllMyBuckets",
    "Resource": "arn:aws-us-gov:s3:::*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:ListBucket",
      "s3:GetBucketLocation"
    ],
    "Resource": "arn:aws-us-gov:s3:::fabric-pool-*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:GetObject",
      "s3:PutObject",
      "s3>DeleteObject"
    ],
    "Resource": "arn:aws-us-gov:s3:::fabric-pool-*",
    "Effect": "Allow"
  }]
}
```

### Top Secret regions

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": "s3:ListAllMyBuckets",
    "Resource": "arn:aws-iso:s3:::*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:ListBucket",
      "s3:GetBucketLocation"
    ],
    "Resource": "arn:aws-iso:s3:::fabric-pool-*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:GetObject",
      "s3:PutObject",
      "s3>DeleteObject"
    ],
    "Resource": "arn:aws-iso:s3:::fabric-pool-*",
    "Effect": "Allow"
  }]
}
```

## Secret regions



```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": "s3:ListAllMyBuckets",
    "Resource": "arn:aws-iso-b:s3:::*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:ListBucket",
      "s3:GetBucketLocation"
    ],
    "Resource": "arn:aws-iso-b:s3:::fabric-pool-*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:GetObject",
      "s3:PutObject",
      "s3>DeleteObject"
    ],
    "Resource": "arn:aws-iso-b:s3:::fabric-pool-*",
    "Effect": "Allow"
  }]
}

```

- Backup policy for Cloud Volumes ONTAP nodes

If you plan to use BlueXP backup and recovery with your Cloud Volumes ONTAP systems, the IAM role for the nodes must include the second policy shown below.

## Standard regions

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "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"
    }
  ]
}
```

## GovCloud (US) regions

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": "arn:aws-us-gov: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-us-gov:s3:::netapp-backup*/**",
      "Effect": "Allow"
    }
  ]
}

```

### Top Secret regions

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": "arn:aws-iso: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-iso:s3:::netapp-backup*/**",
      "Effect": "Allow"
    }
  ]
}

```

## Secret regions

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": "arn:aws-iso-b: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-iso-b:s3:::netapp-backup*/**",
      "Effect": "Allow"
    }
  ]
}

```

- HA mediator

```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:AssignPrivateIpAddresses",
      "ec2:CreateRoute",
      "ec2>DeleteRoute",
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeRouteTables",
      "ec2:DescribeVpcs",
      "ec2:ReplaceRoute",
      "ec2:UnassignPrivateIpAddresses",
      "sts:AssumeRole",
      "ec2:DescribeSubnets"
    ],
    "Resource": "*"
  }
]
}

```

3. Create an IAM role and attach the policies that you created to the role.

### Result

You now have IAM roles that you can select when you create a new Cloud Volumes ONTAP working environment.

### More information

- [AWS documentation: Creating IAM policies](#)
- [AWS documentation: Creating IAM roles](#)

## Set up licensing for Cloud Volumes ONTAP in AWS

After you decide which licensing option you want to use with Cloud Volumes ONTAP, a few steps are required before you can choose that licensing option when creating a new working environment.

### Freemium

Select the Freemium offering to use Cloud Volumes ONTAP free of charge with up to 500 GiB of provisioned capacity. [Learn more about the Freemium offering.](#)

### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the

prompts to subscribe to the pay-as-you-go offering in the AWS Marketplace.

You won't be charged through the marketplace subscription unless you exceed 500 GiB of provisioned capacity, at which time the system is automatically converted to the [Essentials package](#).

**Edit Credentials & Add Subscription**

Select a subscription option and click **Continue**. The AWS Marketplace enables you to view pricing details and then subscribe.

- Pay-Per-TiB - Annual Contract**  
Pay for Cloud Volumes ONTAP with an annual, upfront payment.
- Pay-as-you-go**  
Pay for Cloud Volumes ONTAP at an hourly rate.

**The next steps:**

- 1 AWS Marketplace**  
Subscribe and then click **Set Up Your Account** to configure your account.
- 2 Cloud Manager**  
Save your subscription and associate the Marketplace subscription with your AWS credentials.

**Continue** **Cancel**

b. After you return to BlueXP, select **Freemium** when you reach the charging methods page.

**Select Charging Method**

- Professional** **By capacity** ▾
- Essential** **By capacity** ▾
- Freemium (Up to 500 GiB)** **By capacity** ▾
- Per Node** **By node** ▾

[View step-by-step instructions to launch Cloud Volumes ONTAP in AWS.](#)

## Capacity-based license

Capacity-based licensing enables you to pay for Cloud Volumes ONTAP per TiB of capacity. Capacity-based licensing is available in the form of a *package*: the Essentials package or the Professional package.

The Essentials and Professional packages are available with the following consumption models:

- A license (bring your own license (BYOL)) purchased from NetApp
- An hourly, pay-as-you-go (PAYGO) subscription from the AWS Marketplace
- An annual contract from the AWS Marketplace

[Learn more about capacity-based licensing.](#)

The following sections describe how to get started with each of these consumption models.

### BYOL

Pay upfront by purchasing a license (BYOL) from NetApp to deploy Cloud Volumes ONTAP systems in any cloud provider.

NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAP](#).

### Steps

1. [Contact NetApp Sales to obtain a license](#)
2. [Add your NetApp Support Site account to BlueXP](#)

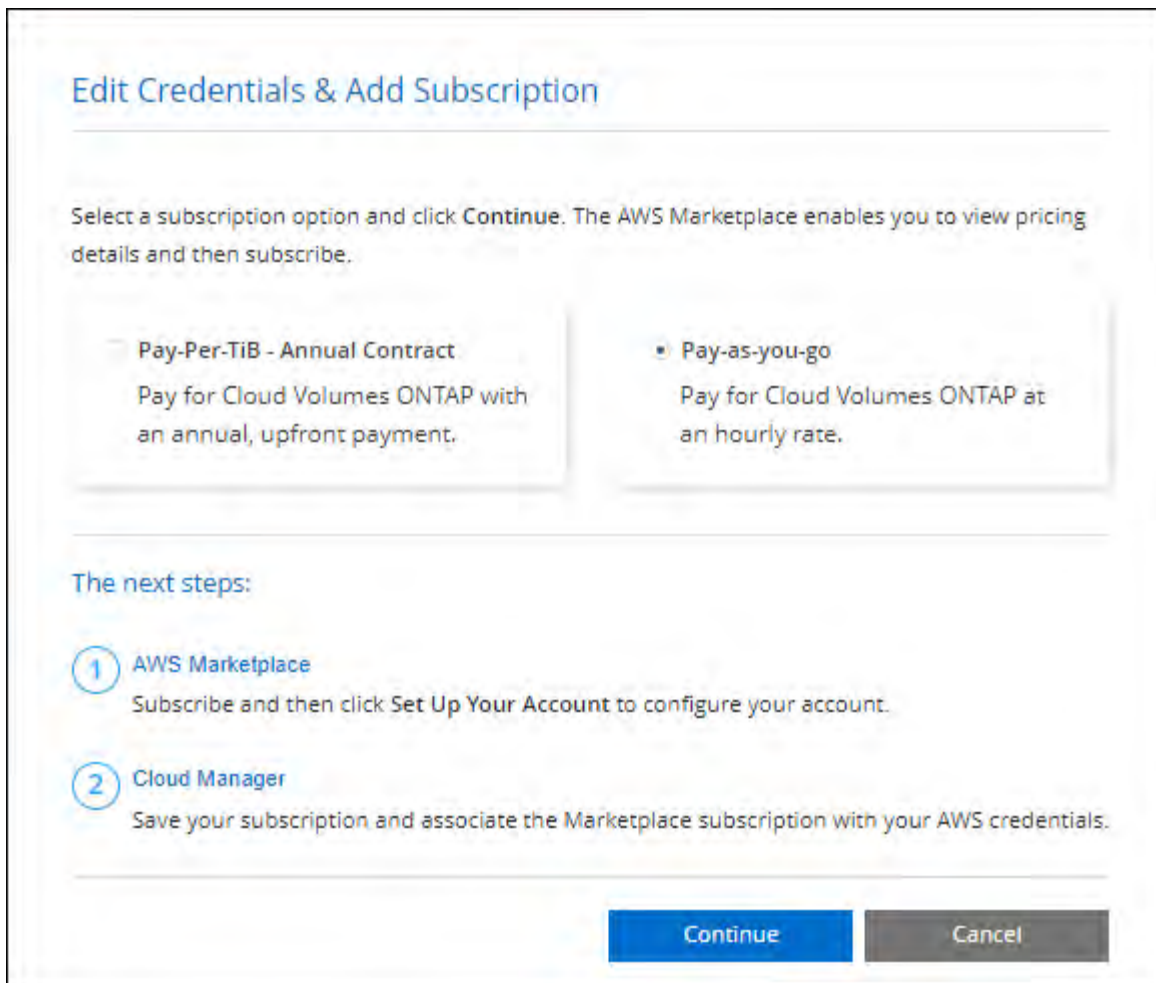
BlueXP automatically queries NetApp's licensing service to obtain details about the licenses associated with your NetApp Support Site account. If there are no errors, BlueXP automatically adds the licenses to the digital wallet.

Your license must be available from the BlueXP digital wallet before you can use it with Cloud Volumes ONTAP. If needed, you can [manually add the license to the BlueXP digital wallet](#).

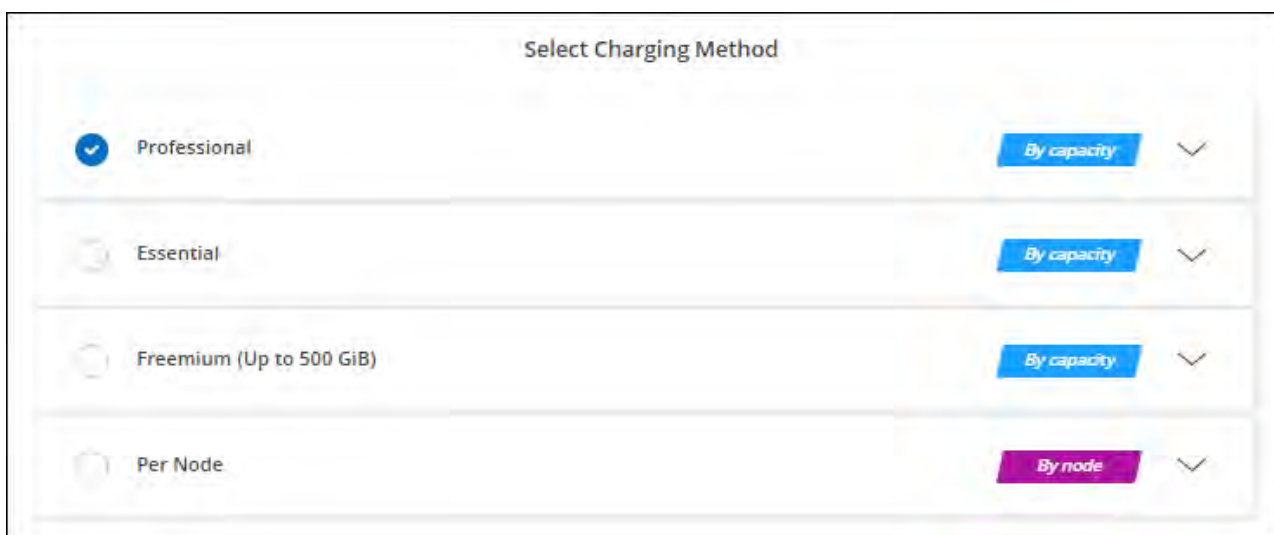
3. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the pay-as-you-go offering in the AWS Marketplace.

The license that you purchased from NetApp is always charged first, but you'll be charged from the hourly rate in the marketplace if you exceed your licensed capacity or if the term of your license expires.





- b. After you return to BlueXP, select a capacity-based package when you reach the charging methods page.



[View step-by-step instructions to launch Cloud Volumes ONTAP in AWS.](#)

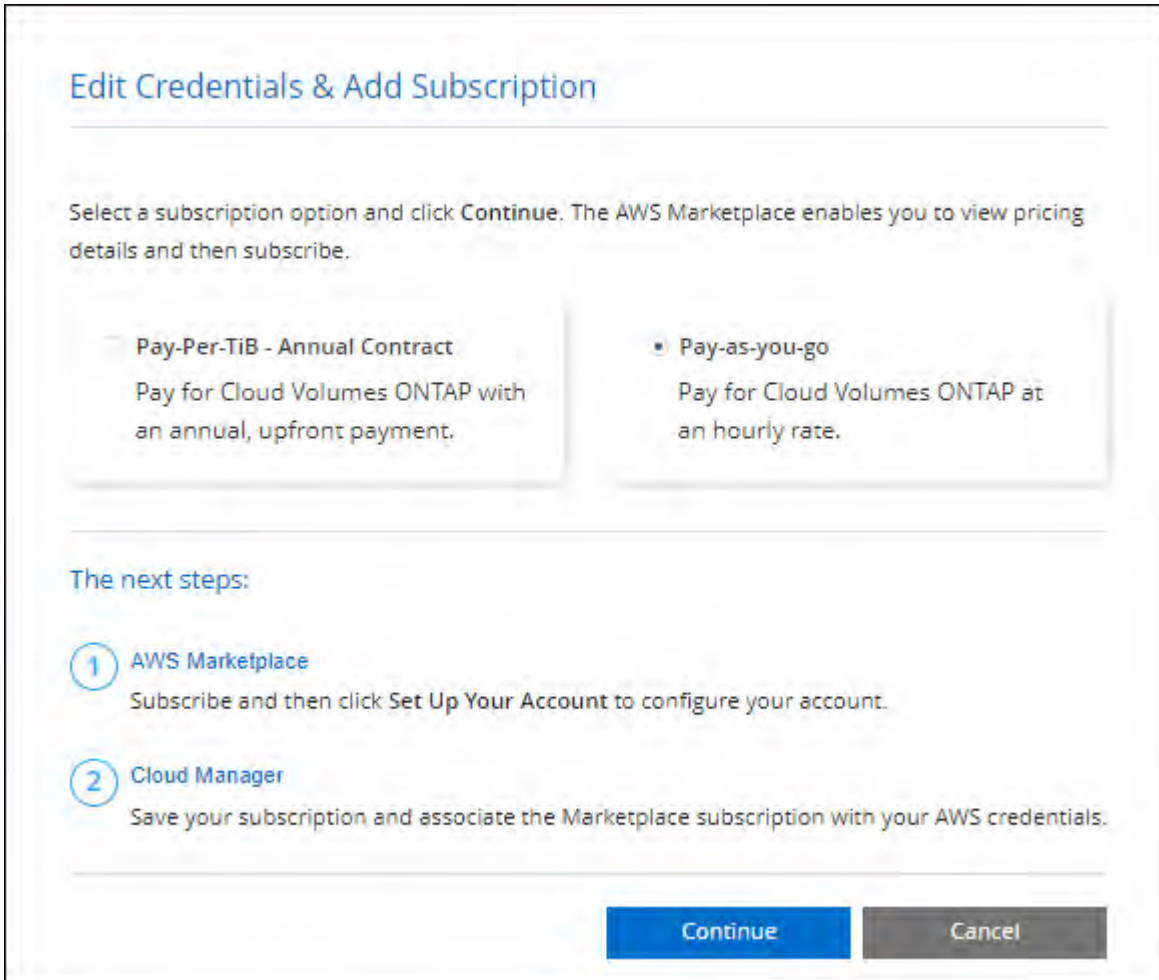
### **PAYGO subscription**

Pay hourly by subscribing to the offer from your cloud provider's marketplace.

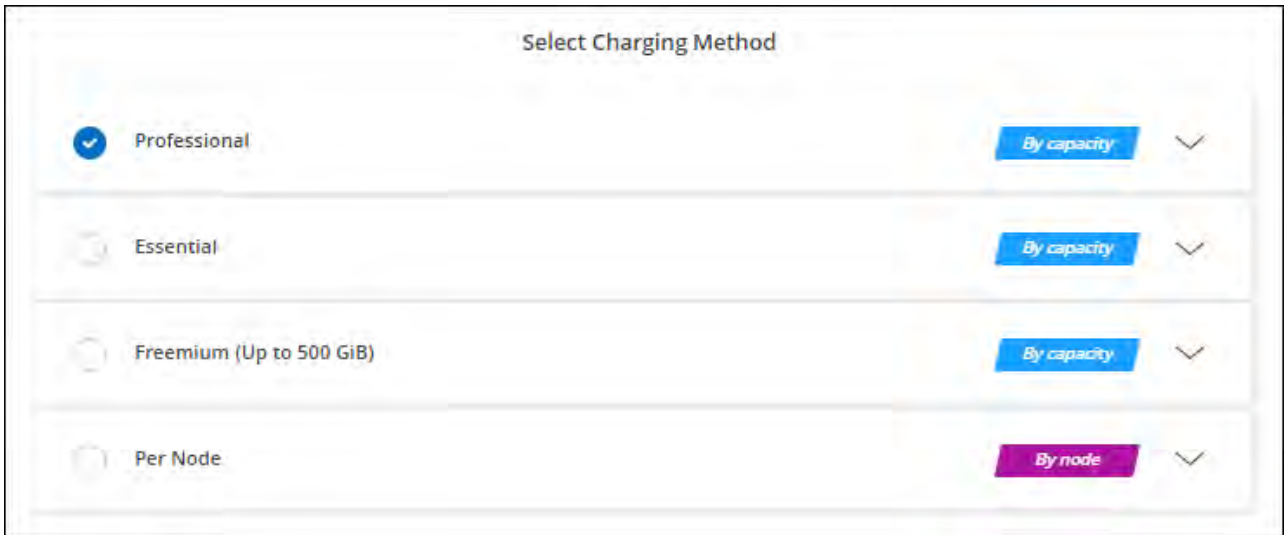
When you create a Cloud Volumes ONTAP working environment, BlueXP prompts you to subscribe to the agreement that's available in the AWS Marketplace. That subscription is then associated with the working environment for charging. You can use that same subscription for additional working environments.

## Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the pay-as-you-go offering in the AWS Marketplace.



- b. After you return to BlueXP, select a capacity-based package when you reach the charging methods page.



[View step-by-step instructions to launch Cloud Volumes ONTAP in AWS.](#)



You can manage the AWS Marketplace subscriptions associated with your AWS accounts from the Settings > Credentials page. [Learn how to manage your AWS accounts and subscriptions](#)

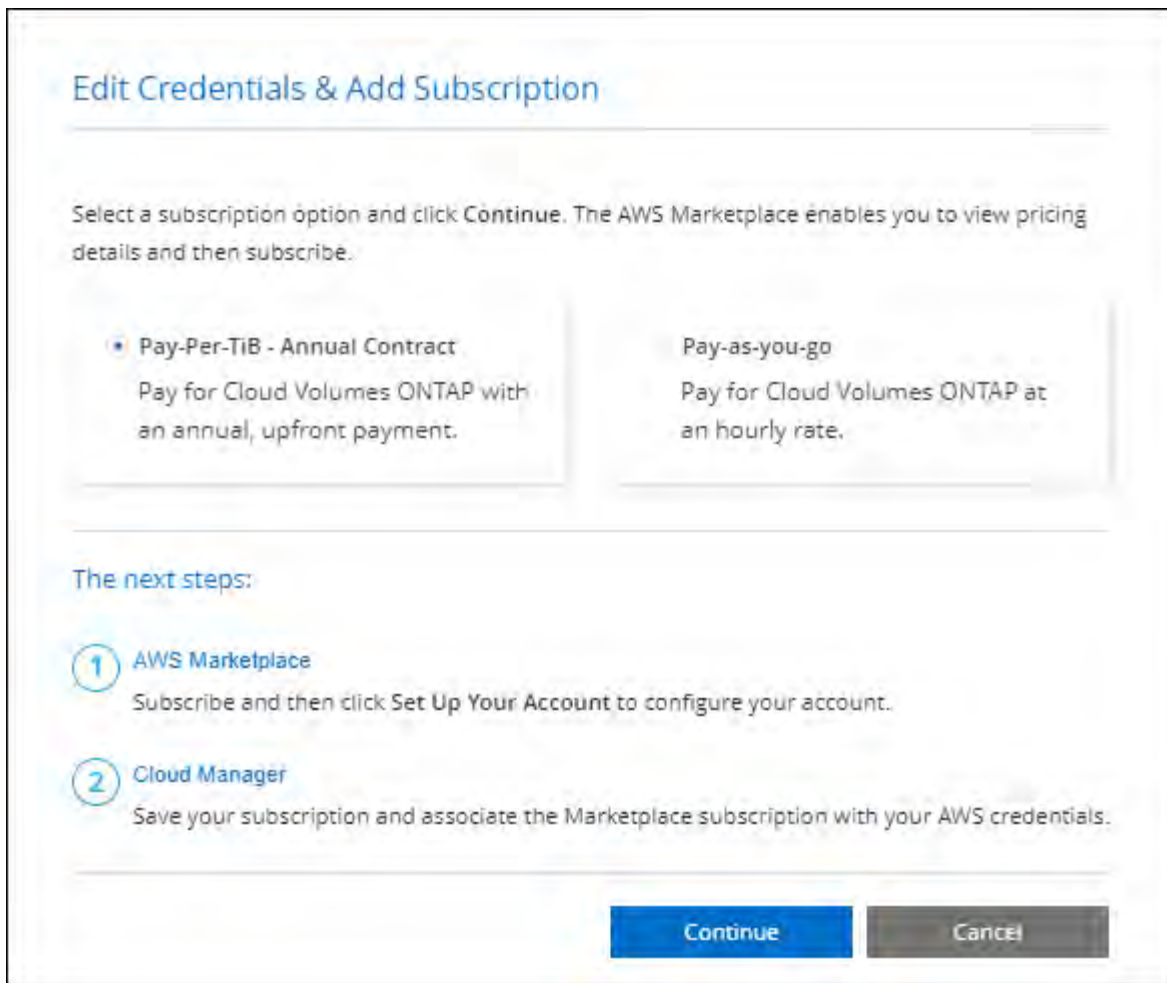
### Annual contract

Pay annually by purchasing an annual contract from your cloud provider's marketplace.

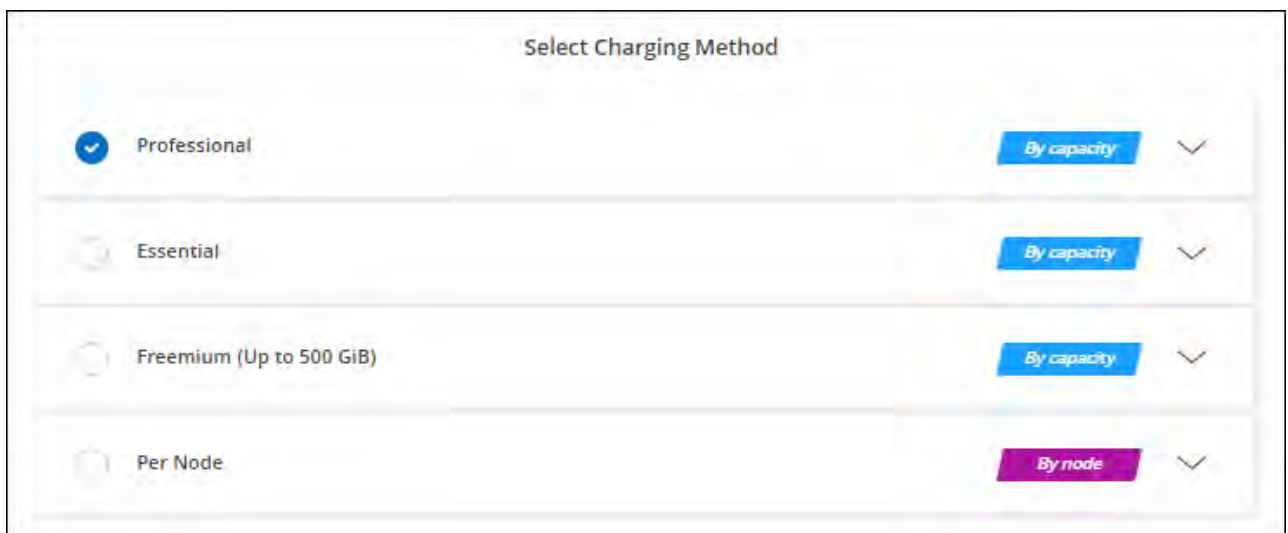
Similar to an hourly subscription, BlueXP prompts you to subscribe to the annual contract that's available in the AWS Marketplace.

### Steps

1. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the annual contract in the AWS Marketplace.



- b. After you return to BlueXP, select a capacity-based package when you reach the charging methods page.



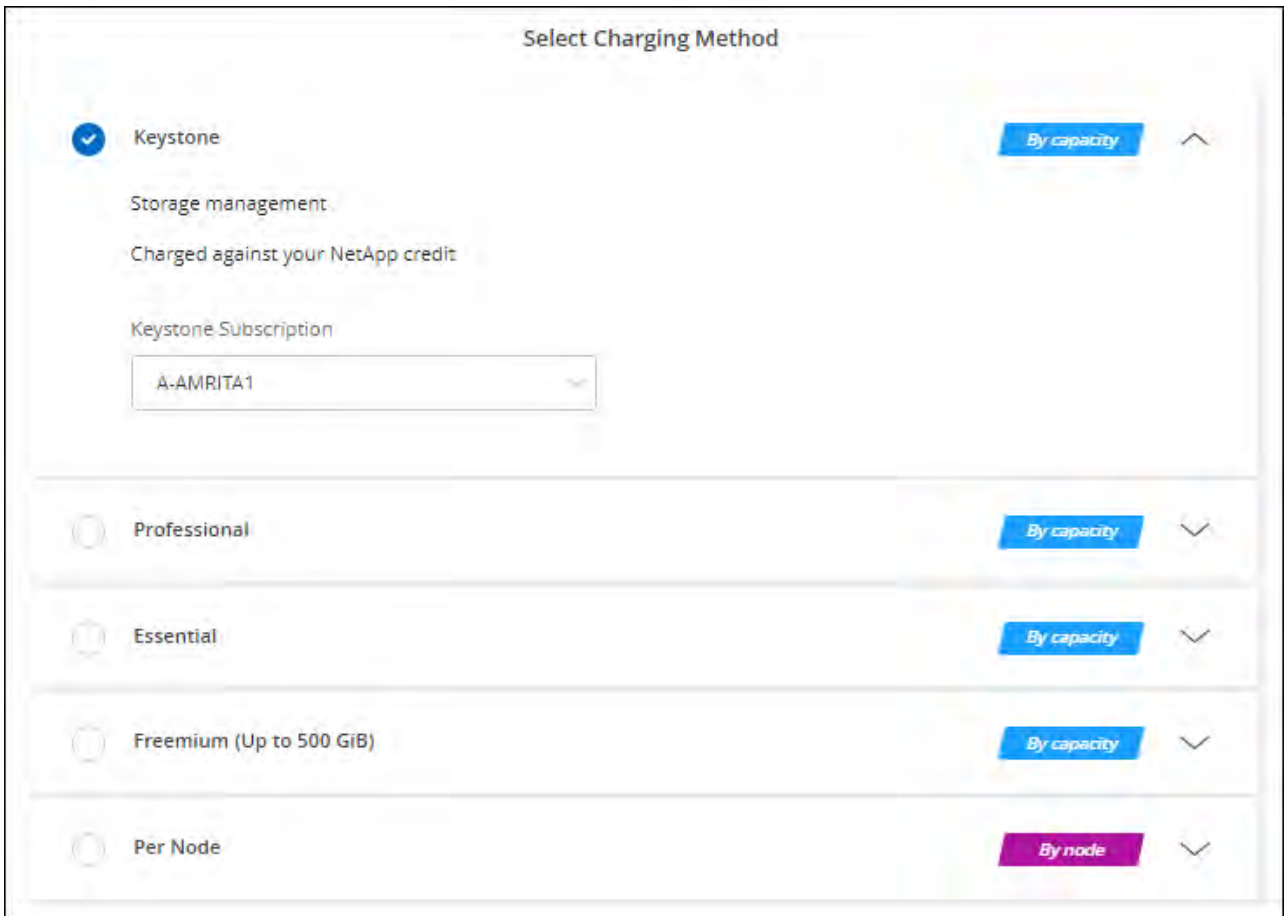
[View step-by-step instructions to launch Cloud Volumes ONTAP in AWS.](#)

## Keystone Subscription

A Keystone Subscription is a pay-as-you-grow subscription-based service. [Learn more about NetApp Keystone Subscriptions.](#)

### Steps

1. If you don't have a subscription yet, [contact NetApp](#)
2. [Contact NetApp](#) to authorize your BlueXP user account with one or more Keystone Subscriptions.
3. After NetApp authorizes your account, [link your subscriptions for use with Cloud Volumes ONTAP](#).
4. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. Select the Keystone Subscription charging method when prompted to choose a charging method.



[View step-by-step instructions to launch Cloud Volumes ONTAP in AWS.](#)

## Launch Cloud Volumes ONTAP in AWS

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

### Before you begin

You need the following to create a working environment.

- A Connector that's up and running.

- You should have a [Connector that is associated with your project or workspace](#).
- [You should be prepared to leave the Connector running at all times](#).
- An understanding of the configuration that you want to use.

You should have prepared by choosing a configuration and by obtaining AWS networking information from your administrator. For details, refer to [Planning your Cloud Volumes ONTAP configuration](#).

- An understanding of what's required to set up licensing for Cloud Volumes ONTAP.

[Learn how to set up licensing](#).

- DNS and Active Directory for CIFS configurations.

For details, refer to [Networking requirements for Cloud Volumes ONTAP in AWS](#).

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

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

### About this task

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

### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the prompts.
3. **Choose a Location:** Select **Amazon Web Services** and **Cloud Volumes ONTAP Single Node**.
4. If you're prompted, [create a Connector](#).
5. **Details and Credentials:** Optionally change the AWS credentials and subscription, enter a working environment name, add tags if needed, and then enter a password.

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

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

Field	Description
Add tags	<p>AWS tags are metadata for your AWS resources. BlueXP adds the tags to the Cloud Volumes ONTAP instance and each AWS resource associated with the instance.</p> <p>You can add up to four tags from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four tags when creating a working environment.</p> <p>For information about tags, refer to <a href="#">AWS Documentation: Tagging your Amazon EC2 Resources</a>.</p>
User name and password	<p>These are the credentials for the Cloud Volumes ONTAP cluster administrator account. You can use these credentials to connect to Cloud Volumes ONTAP through ONTAP System Manager or the ONTAP CLI. Keep the default <i>admin</i> user name or change it to a custom user name.</p>
Edit Credentials	<p>Choose the AWS credentials associated with the account where you want to deploy this system. You can also associate the AWS marketplace subscription to use with this Cloud Volumes ONTAP system.</p> <p>Click <b>Add Subscription</b> to associate the selected credentials with a new AWS marketplace subscription. The subscription can be for an annual contract or to pay for Cloud Volumes ONTAP at an hourly rate.</p> <p><a href="#">Learn how to add additional AWS credentials to BlueXP.</a></p>

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

### [Subscribe to BlueXP from the AWS marketplace](#)

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



#### Cloud Manager (for Cloud Volumes ONTAP)

---

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

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

*Subscribe*

You are already subscribed to this product

---

**Pricing Details**

Software Fees

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

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





If you would like to utilize WORM and data tiering, you must disable BlueXP backup and recovery and deploy a Cloud Volumes ONTAP working environment with version 9.8 or above.

7. **Location & Connectivity:** Enter the network information that you recorded in the [AWS worksheet](#).

The following table describes fields for which you might need guidance:

Field	Description
VPC	If you have an AWS Outpost, you can deploy a single node Cloud Volumes ONTAP system in that Outpost by selecting the Outpost VPC. The experience is the same as any other VPC that resides in AWS.
Generated security group	If you let BlueXP generate the security group for you, you need to choose how you'll allow traffic: <ul style="list-style-type: none"><li>• If you choose <b>Selected VPC only</b>, the source for inbound traffic is the subnet range of the selected VPC and the subnet range of the VPC where the Connector resides. This is the recommended option.</li><li>• If you choose <b>All VPCs</b>, the source for inbound traffic is the 0.0.0.0/0 IP range.</li></ul>
Use existing security group	If you use an existing firewall policy, ensure that it includes the required rules. <a href="#">Learn about firewall rules for Cloud Volumes ONTAP</a> .

8. **Data Encryption:** Choose no data encryption or AWS-managed encryption.

For AWS-managed encryption, you can choose a different Customer Master Key (CMK) from your account or another AWS account.



You can't change the AWS data encryption method after you create a Cloud Volumes ONTAP system.

[Learn how to set up the AWS KMS for Cloud Volumes ONTAP](#).

[Learn more about supported encryption technologies](#).

9. **Charging Methods and NSS Account:** Specify which charging option would you like to use with this system, and then specify a NetApp Support Site account.

- [Learn about licensing options for Cloud Volumes ONTAP](#).
- [Learn how to set up licensing](#).

10. **Cloud Volumes ONTAP Configuration** (annual AWS marketplace contract only): Review the default configuration and click **Continue** or click **Change Configuration** to select your own configuration.

If you keep the default configuration, then you only need to specify a volume and then review and approve the configuration.

11. **Preconfigured Packages:** Select one of the packages to quickly launch Cloud Volumes ONTAP, or click **Change Configuration** to select your own configuration.

If you choose one of the packages, then you only need to specify a volume and then review and approve



the configuration.

12. **IAM Role:** It's best to keep the default option to let BlueXP create the role for you.

If you prefer to use your own policy, it must meet [policy requirements for Cloud Volumes ONTAP nodes](#).

13. **Licensing:** Change the Cloud Volumes ONTAP version as needed and select an instance type and the instance tenancy.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then BlueXP updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.13.1 and 9.13.1 P4 is available. The update does not occur from one release to another—for example, from 9.13 to 9.14.

14. **Underlying Storage Resources:** Choose a disk type, configure the underlying storage, and choose whether to keep data tiering enabled.

Note the following:

- The disk type is for the initial volume (and aggregate). You can choose a different disk type for subsequent volumes (and aggregates).
- If you choose a gp3 or io1 disk, BlueXP uses the Elastic Volumes feature in AWS to automatically increase the underlying storage disk capacity as needed. You can choose the initial capacity based on your storage needs and revise it after Cloud Volumes ONTAP is deployed. [Learn more about support for Elastic Volumes in AWS](#).
- If you choose a gp2 or st1 disk, you can select a disk size for all disks in the initial aggregate and for any additional aggregates that BlueXP creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.
- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates.

[Learn how data tiering works](#).

15. **Write Speed & WORM:**

- a. Choose **Normal** or **High** write speed, if desired.

[Learn more about write speed](#).

- b. Activate write once, read many (WORM) storage, if desired.

WORM can't be enabled if data tiering was enabled for Cloud Volumes ONTAP versions 9.7 and below. Reverting or downgrading to Cloud Volumes ONTAP 9.8 is blocked after enabling WORM and tiering.

[Learn more about WORM storage](#).

- c. If you activate WORM storage, select the retention period.

16. **Create Volume:** Enter details for the new volume or click **Skip**.

[Learn about supported client protocols and versions](#).

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

need guidance:

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, BlueXP enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, BlueXP automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

### Volume Details, Protection & Protocol

Details & Protection	Protocol
<p>Volume Name: <input style="width: 200px;" type="text" value="vol"/>      Size (GB): <input style="width: 80px;" type="text" value="250"/> <span style="float: right; font-size: 0.8em;">?</span></p> <p>Snapshot Policy: <input style="width: 300px;" type="text" value="default"/> <span style="float: right; font-size: 0.8em;">▼</span></p> <p><span style="font-size: 0.8em;">?</span> Default Policy</p>	<p style="text-align: center;"> <span style="margin: 0 10px;">NFS</span> <span style="border-bottom: 2px solid #0070C0; display: inline-block; width: 100px; text-align: center; margin: 0 10px;">CIFS</span> <span style="margin: 0 10px;">iSCSI</span> </p> <p>Share name: <input style="width: 150px;" type="text" value="vol_share"/>      Permissions: <input style="width: 150px;" type="text" value="Full Control"/> <span style="float: right; font-size: 0.8em;">▼</span></p> <p>Users / Groups: <input style="width: 300px;" type="text" value="engineering"/></p> <p style="font-size: 0.8em; color: #0070C0;">Valid users and groups separated by a semicolon</p>

17. **CIFS Setup:** If you chose the CIFS protocol, set up a CIFS server.

Field	Description
DNS Primary and Secondary IP Address	The IP addresses of the DNS servers that provide name resolution for the CIFS server. The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers. If you configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=Computers,OU=corp</b> in this field.
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. Refer to the <a href="#">BlueXP automation docs</a> for details.  Note that you can configure an NTP server only when creating a CIFS server. It's not configurable after you create the CIFS server.

18. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and edit the volume tiering policy, if needed.

For more information, refer to [Understanding volume usage profiles](#), [Data tiering overview](#), and [KB: What Inline Storage Efficiency features are supported with CVO?](#)

19. **Review & Approve:** Review and confirm your selections.
  - a. Review details about the configuration.
  - b. Click **More information** to review details about support and the AWS resources that BlueXP will purchase.
  - c. Select the **I understand...** check boxes.
  - d. Click **Go**.

## Result

BlueXP launches the Cloud Volumes ONTAP instance. You can track the progress in the timeline.

If you have any issues launching the Cloud Volumes ONTAP instance, review the failure message. You can also select the working environment and click Re-create environment.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

## After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.
- If you want to apply quotas to volumes, use ONTAP System Manager or the ONTAP CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

## Launch a Cloud Volumes ONTAP HA pair in AWS

If you want to launch a Cloud Volumes ONTAP HA pair in AWS, you need to create an HA working environment in BlueXP.

## Limitation

At this time, HA pairs are not supported with AWS Outposts.

## About this task

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

## Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the prompts.
3. **Choose a Location:** Select **Amazon Web Services** and **Cloud Volumes ONTAP HA**.

Some AWS Local Zones are available.

Before you can use AWS Local Zones, you must enable Local Zones and create a subnet in the Local Zone in your AWS account. Follow the **Opt in to an AWS Local Zone** and **Extend your Amazon VPC to the Local Zone** steps in the [AWS tutorial "Get Started Deploying Low Latency Applications with AWS Local Zones"](#).

If you are running a Connector version 3.9.36 or below, you need to add the following permission to the AWS Connector role in the AWS EC2 console: DescribeAvailabilityZones.

4. **Details and Credentials:** Optionally change the AWS credentials and subscription, enter a working environment name, add tags if needed, and then enter a password.

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

Field	Description
Working Environment Name	BlueXP uses the working environment name to name both the Cloud Volumes ONTAP system and the Amazon EC2 instance. It also uses the name as the prefix for the predefined security group, if you select that option.
Add tags	<p>AWS tags are metadata for your AWS resources. BlueXP adds the tags to the Cloud Volumes ONTAP instance and each AWS resource associated with the instance.</p> <p>You can add up to four tags from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four tags when creating a working environment.</p> <p>For information about tags, refer to <a href="#">AWS Documentation: Tagging your Amazon EC2 Resources</a>.</p>
User name and password	These are the credentials for the Cloud Volumes ONTAP cluster administrator account. You can use these credentials to connect to Cloud Volumes ONTAP through ONTAP System Manager or the ONTAP CLI. Keep the default <i>admin</i> user name or change it to a custom user name.
Edit Credentials	<p>Choose the AWS credentials and marketplace subscription to use with this Cloud Volumes ONTAP system.</p> <p>Click <b>Add Subscription</b> to associate the selected credentials with a new AWS marketplace subscription. The subscription can be for an annual contract or to pay for Cloud Volumes ONTAP at an hourly rate.</p> <p>If you purchased a license directly from NetApp (bring your own license (BYOL)), then an AWS subscription isn't required. NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to <a href="#">Restricted availability of BYOL licensing for Cloud Volumes ONTAP</a>.</p> <p><a href="#">Learn how to add additional AWS credentials to BlueXP.</a></p>

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

#### [Subscribe to BlueXP from the AWS marketplace](#)



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

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

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



If you would like to utilize WORM and data tiering, you must disable BlueXP backup and recovery and deploy a Cloud Volumes ONTAP working environment with version 9.8 or above.

6. **HA Deployment Models:** Choose an HA configuration.

For an overview of the deployment models, refer to [Cloud Volumes ONTAP HA for AWS](#).

7. **Location and Connectivity** (single AZ) or **Region & VPC** (multiple AZs): Enter the network information that you recorded in the AWS worksheet.

The following table describes fields for which you might need guidance:

Field	Description
Generated security group	If you let BlueXP generate the security group for you, you need to choose how you'll allow traffic: <ul style="list-style-type: none"><li>• If you choose <b>Selected VPC only</b>, the source for inbound traffic is the subnet range of the selected VPC and the subnet range of the VPC where the Connector resides. This is the recommended option.</li><li>• If you choose <b>All VPCs</b>, the source for inbound traffic is the 0.0.0.0/0 IP range.</li></ul>
Use existing security group	If you use an existing firewall policy, ensure that it includes the required rules. <a href="#">Learn about firewall rules for Cloud Volumes ONTAP</a> .

8. **Connectivity and SSH Authentication:** Choose connection methods for the HA pair and the mediator.

9. **Floating IPs:** If you chose multiple AZs, specify the floating IP addresses.

The IP addresses must be outside of the CIDR block for all VPCs in the region. For additional details, refer to [AWS networking requirements for Cloud Volumes ONTAP HA in multiple AZs](#).

10. **Route Tables:** If you chose multiple AZs, select the route tables that should include routes to the floating IP addresses.

If you have more than one route table, it is very important to select the correct route tables. Otherwise, some clients might not have access to the Cloud Volumes ONTAP HA pair. For more information about route tables, refer to the [AWS Documentation: Route Tables](#).

11. **Data Encryption:** Choose no data encryption or AWS-managed encryption.

For AWS-managed encryption, you can choose a different Customer Master Key (CMK) from your account or another AWS account.



You can't change the AWS data encryption method after you create a Cloud Volumes ONTAP system.

[Learn how to set up the AWS KMS for Cloud Volumes ONTAP.](#)

[Learn more about supported encryption technologies.](#)

12. **Charging Methods and NSS Account:** Specify which charging option would you like to use with this system, and then specify a NetApp Support Site account.
  - [Learn about licensing options for Cloud Volumes ONTAP.](#)
  - [Learn how to set up licensing.](#)

13. **Cloud Volumes ONTAP Configuration** (annual AWS Marketplace contract only): Review the default configuration and click **Continue** or click **Change Configuration** to select your own configuration.

If you keep the default configuration, then you only need to specify a volume and then review and approve the configuration.

14. **Preconfigured Packages** (hourly or BYOL only): Select one of the packages to quickly launch Cloud Volumes ONTAP, or click **Change Configuration** to select your own configuration.

If you choose one of the packages, then you only need to specify a volume and then review and approve the configuration.

15. **IAM Role:** It's best to keep the default option to let BlueXP create the role for you.

If you prefer to use your own policy, it must meet [policy requirements for Cloud Volumes ONTAP nodes and the HA mediator](#).

16. **Licensing:** Change the Cloud Volumes ONTAP version as needed and select an instance type and the instance tenancy.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then BlueXP updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.13.1 and 9.13.1 P4 is available. The update does not occur from one release to another—for example, from 9.13 to 9.14.

17. **Underlying Storage Resources:** Choose a disk type, configure the underlying storage, and choose whether to keep data tiering enabled.

Note the following:

- The disk type is for the initial volume (and aggregate). You can choose a different disk type for subsequent volumes (and aggregates).
- If you choose a gp3 or io1 disk, BlueXP uses the Elastic Volumes feature in AWS to automatically increase the underlying storage disk capacity as needed. You can choose the initial capacity based on your storage needs and revise it after Cloud Volumes ONTAP is deployed. [Learn more about support for Elastic Volumes in AWS.](#)
- If you choose a gp2 or st1 disk, you can select a disk size for all disks in the initial aggregate and for any additional aggregates that BlueXP creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.
- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates.

[Learn how data tiering works.](#)

18. **Write Speed & WORM:**

- a. Choose **Normal** or **High** write speed, if desired.

[Learn more about write speed.](#)

- b. Activate write once, read many (WORM) storage, if desired.

WORM can't be enabled if data tiering was enabled for Cloud Volumes ONTAP versions 9.7 and below. Reverting or downgrading to Cloud Volumes ONTAP 9.8 is blocked after enabling WORM and tiering.

[Learn more about WORM storage.](#)

- c. If you activate WORM storage, select the retention period.

19. **Create Volume:** Enter details for the new volume or click **Skip**.

[Learn about supported client protocols and versions.](#)

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

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, BlueXP enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.



Field	Description
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, BlueXP automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

### Volume Details, Protection & Protocol

#### Details & Protection

Volume Name:  Size (GB):

Snapshot Policy:

Default Policy

#### Protocol

NFS     CIFS     iSCSI

Share name:  Permissions:

Users / Groups:

Valid users and groups separated by a semicolon

20. **CIFS Setup:** If you selected the CIFS protocol, set up a CIFS server.

Field	Description
DNS Primary and Secondary IP Address	<p>The IP addresses of the DNS servers that provide name resolution for the CIFS server.</p> <p>The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.</p>
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.

Field	Description
Organizational Unit	The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers. If you configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=Computers,OU=corp</b> in this field.
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. Refer to the <a href="#">BlueXP automation docs</a> for details.  Note that you can configure an NTP server only when creating a CIFS server. It's not configurable after you create the CIFS server.

21. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and edit the volume tiering policy, if needed.

For more information, refer to [Choose a volume usage profile](#) and [Data tiering overview](#).

22. **Review & Approve:** Review and confirm your selections.

- a. Review details about the configuration.
- b. Click **More information** to review details about support and the AWS resources that BlueXP will purchase.
- c. Select the **I understand...** check boxes.
- d. Click **Go**.

### Result

BlueXP launches the Cloud Volumes ONTAP HA pair. You can track the progress in the timeline.

If you experience any issues launching the HA pair, review the failure message. You can also select the working environment and click Re-create environment.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

### After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.
- If you want to apply quotas to volumes, use ONTAP System Manager or the ONTAP CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

## Deploy Cloud Volumes ONTAP in AWS Secret Cloud or AWS Top Secret Cloud

Similar to a standard AWS region, you can use BlueXP in [AWS Secret Cloud](#) and in [AWS Top Secret Cloud](#) to deploy Cloud Volumes ONTAP, which provides enterprise-class features for your cloud storage. AWS Secret Cloud and Top Secret Cloud are closed regions specific to the U.S. Intelligence Community; the instructions on this page only

apply to AWS Secret Cloud and Top Secret Cloud region users.

### Before you begin

Before you get started, review the supported versions in AWS Secret Cloud and Top Secret Cloud, and learn about private mode in BlueXP.

- Review the following supported versions in AWS Secret Cloud and Top Secret Cloud:
  - Cloud Volumes ONTAP 9.12.1 P2
  - Version 3.9.32 of the Connector

The Connector is software that's required to deploy and manage Cloud Volumes ONTAP in AWS. You'll log in to BlueXP from the software that gets installed on the Connector instance. The SaaS website for BlueXP isn't supported in AWS Secret Cloud and Top Secret Cloud.

- Learn about private mode

In AWS Secret Cloud and Top Secret Cloud, BlueXP operates in *private mode*. In private mode, there is no connectivity to the BlueXP SaaS layer. Users access BlueXP locally from the web-based console that's available from the Connector, not from the SaaS layer.

To learn more about how private mode works, refer to [BlueXP private deployment mode](#).

### Step 1: Set up your networking

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

#### Steps

1. Choose the VPC and subnets in which you want to launch the Connector instance and Cloud Volumes ONTAP instances.
2. Ensure that your VPC and subnets will support connectivity between the Connector and Cloud Volumes ONTAP.
3. Set up a VPC endpoint to the S3 service.

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

### Step 2: Set up permissions

Set up IAM policies and roles that provide the Connector and Cloud Volumes ONTAP with the permissions that they need to perform actions in the AWS Secret Cloud or Top Secret Cloud.

You need an IAM policy and IAM role for each of the following:

- The Connector instance
- Cloud Volumes ONTAP instances
- For HA pairs, the Cloud Volumes ONTAP HA mediator instance (if you want to deploy HA pairs)

#### Steps

1. Go to the AWS IAM console and click **Policies**.
2. Create a policy for the Connector instance.



You create these policies to support the S3 buckets in your AWS environment. While creating the buckets later, ensure that the bucket names are prefixed with `fabric-pool-`. This requirement applies to both the AWS Secret Cloud and Top Secret Cloud regions.

## Secret regions

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeInstances",
      "ec2:DescribeInstanceStatus",
      "ec2:RunInstances",
      "ec2:ModifyInstanceAttribute",
      "ec2:DescribeRouteTables",
      "ec2:DescribeImages",
      "ec2:CreateTags",
      "ec2:CreateVolume",
      "ec2:DescribeVolumes",
      "ec2:ModifyVolumeAttribute",
      "ec2>DeleteVolume",
      "ec2:CreateSecurityGroup",
      "ec2>DeleteSecurityGroup",
      "ec2:DescribeSecurityGroups",
      "ec2:RevokeSecurityGroupEgress",
      "ec2:RevokeSecurityGroupIngress",
      "ec2:AuthorizeSecurityGroupEgress",
      "ec2:AuthorizeSecurityGroupIngress",
      "ec2:CreateNetworkInterface",
      "ec2:DescribeNetworkInterfaces",
      "ec2>DeleteNetworkInterface",
      "ec2:ModifyNetworkInterfaceAttribute",
      "ec2:DescribeSubnets",
      "ec2:DescribeVpcs",
      "ec2:DescribeDhcpOptions",
      "ec2:CreateSnapshot",
      "ec2>DeleteSnapshot",
      "ec2:DescribeSnapshots",
      "ec2:GetConsoleOutput",
      "ec2:DescribeKeyPairs",
      "ec2:DescribeRegions",
      "ec2>DeleteTags",
      "ec2:DescribeTags",
      "cloudformation:CreateStack",
      "cloudformation>DeleteStack",
      "cloudformation:DescribeStacks",
      "cloudformation:DescribeStackEvents",
      "cloudformation:ValidateTemplate",
      "iam:PassRole",
```

```

        "iam:CreateRole",
        "iam:DeleteRole",
        "iam:PutRolePolicy",
        "iam:ListInstanceProfiles",
        "iam:CreateInstanceProfile",
        "iam:DeleteRolePolicy",
        "iam:AddRoleToInstanceProfile",
        "iam:RemoveRoleFromInstanceProfile",
        "iam:DeleteInstanceProfile",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:GetBucketTagging",
        "s3:GetBucketLocation",
        "s3:ListAllMyBuckets",
        "kms:List*",
        "kms:Describe*",
        "ec2:AssociateIamInstanceProfile",
        "ec2:DescribeIamInstanceProfileAssociations",
        "ec2:DisassociateIamInstanceProfile",
        "ec2:DescribeInstanceAttribute",
        "ec2:CreatePlacementGroup",
        "ec2>DeletePlacementGroup"
    ],
    "Resource": "*"
},
{
    "Sid": "fabricPoolPolicy",
    "Effect": "Allow",
    "Action": [
        "s3:DeleteBucket",
        "s3:GetLifecycleConfiguration",
        "s3:PutLifecycleConfiguration",
        "s3:PutBucketTagging",
        "s3:ListBucketVersions"
    ],
    "Resource": [
        "arn:aws-iso-b:s3:::fabric-pool*"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances",
        "ec2:AttachVolume",

```

```

        "ec2:DetachVolume"
    ],
    "Condition": {
        "StringLike": {
            "ec2:ResourceTag/WorkingEnvironment": "*"
        }
    },
    "Resource": [
        "arn:aws-iso-b:ec2:*:*:instance/*"
    ]
},
{
    "Effect": "Allow",
    "Action": [
        "ec2:AttachVolume",
        "ec2:DetachVolume"
    ],
    "Resource": [
        "arn:aws-iso-b:ec2:*:*:volume/*"
    ]
}
]
}

```

### Top Secret regions

```

{
    "Version": "2012-10-17",
    "Statement": [{
        "Effect": "Allow",
        "Action": [
            "ec2:DescribeInstances",
            "ec2:DescribeInstanceStatus",
            "ec2:RunInstances",
            "ec2:ModifyInstanceAttribute",
            "ec2:DescribeRouteTables",
            "ec2:DescribeImages",
            "ec2:CreateTags",
            "ec2:CreateVolume",
            "ec2:DescribeVolumes",
            "ec2:ModifyVolumeAttribute",
            "ec2>DeleteVolume",
            "ec2:CreateSecurityGroup",
            "ec2>DeleteSecurityGroup",
            "ec2:DescribeSecurityGroups",
            "ec2:RevokeSecurityGroupEgress",

```

```
"ec2:RevokeSecurityGroupIngress",
"ec2:AuthorizeSecurityGroupEgress",
"ec2:AuthorizeSecurityGroupIngress",
"ec2:CreateNetworkInterface",
"ec2:DescribeNetworkInterfaces",
"ec2>DeleteNetworkInterface",
"ec2:ModifyNetworkInterfaceAttribute",
"ec2:DescribeSubnets",
"ec2:DescribeVpcs",
"ec2:DescribeDhcpOptions",
"ec2:CreateSnapshot",
"ec2>DeleteSnapshot",
"ec2:DescribeSnapshots",
"ec2:GetConsoleOutput",
"ec2:DescribeKeyPairs",
"ec2:DescribeRegions",
"ec2>DeleteTags",
"ec2:DescribeTags",
"cloudformation:CreateStack",
"cloudformation>DeleteStack",
"cloudformation:DescribeStacks",
"cloudformation:DescribeStackEvents",
"cloudformation:ValidateTemplate",
"iam:PassRole",
"iam:CreateRole",
"iam>DeleteRole",
"iam:PutRolePolicy",
"iam:ListInstanceProfiles",
"iam:CreateInstanceProfile",
"iam>DeleteRolePolicy",
"iam:AddRoleToInstanceProfile",
"iam:RemoveRoleFromInstanceProfile",
"iam>DeleteInstanceProfile",
"s3:GetObject",
"s3:ListBucket",
"s3:GetBucketTagging",
"s3:GetBucketLocation",
"s3:ListAllMyBuckets",
"kms:List*",
"kms:Describe*",
"ec2:AssociateIamInstanceProfile",
"ec2:DescribeIamInstanceProfileAssociations",
"ec2:DisassociateIamInstanceProfile",
"ec2:DescribeInstanceAttribute",
"ec2:CreatePlacementGroup",
"ec2>DeletePlacementGroup"
```



```

    ],
    "Resource": "*"
  },
  {
    "Sid": "fabricPoolPolicy",
    "Effect": "Allow",
    "Action": [
      "s3:DeleteBucket",
      "s3:GetLifecycleConfiguration",
      "s3:PutLifecycleConfiguration",
      "s3:PutBucketTagging",
      "s3:ListBucketVersions"
    ],
    "Resource": [
      "arn:aws-iso:s3:::fabric-pool*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "ec2:StartInstances",
      "ec2:StopInstances",
      "ec2:TerminateInstances",
      "ec2:AttachVolume",
      "ec2:DetachVolume"
    ],
    "Condition": {
      "StringLike": {
        "ec2:ResourceTag/WorkingEnvironment": "*"
      }
    },
    "Resource": [
      "arn:aws-iso:ec2:*:*:instance/*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": [
      "ec2:AttachVolume",
      "ec2:DetachVolume"
    ],
    "Resource": [
      "arn:aws-iso:ec2:*:*:volume/*"
    ]
  }
]

```

```
}
```

3. Create a policy for Cloud Volumes ONTAP.

## Secret regions

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": "s3:ListAllMyBuckets",
    "Resource": "arn:aws-iso-b:s3:::*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:ListBucket",
      "s3:GetBucketLocation"
    ],
    "Resource": "arn:aws-iso-b:s3:::fabric-pool-*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:GetObject",
      "s3:PutObject",
      "s3>DeleteObject"
    ],
    "Resource": "arn:aws-iso-b:s3:::fabric-pool-*",
    "Effect": "Allow"
  }
]
```

## Top Secret regions

```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Action": "s3:ListAllMyBuckets",
    "Resource": "arn:aws-iso:s3:::*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:ListBucket",
      "s3:GetBucketLocation"
    ],
    "Resource": "arn:aws-iso:s3:::fabric-pool-*",
    "Effect": "Allow"
  }, {
    "Action": [
      "s3:GetObject",
      "s3:PutObject",
      "s3>DeleteObject"
    ],
    "Resource": "arn:aws-iso:s3:::fabric-pool-*",
    "Effect": "Allow"
  }]
}

```

For HA pairs, if you plan to deploy a Cloud Volumes ONTAP HA pair, create a policy for the HA mediator.

```

{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "ec2:AssignPrivateIpAddresses",
      "ec2:CreateRoute",
      "ec2>DeleteRoute",
      "ec2:DescribeNetworkInterfaces",
      "ec2:DescribeRouteTables",
      "ec2:DescribeVpcs",
      "ec2:ReplaceRoute",
      "ec2:UnassignPrivateIpAddresses"
    ],
    "Resource": "*"
  }
]
}

```

4. Create IAM roles with the role type Amazon EC2 and attach the policies that you created in the previous steps.

**Create the role:**

Similar to the policies, you should have one IAM role for the Connector and one for the Cloud Volumes ONTAP nodes.

For HA pairs: Similar to the policies, you should have one IAM role for the Connector, one for the Cloud Volumes ONTAP nodes, and one for the HA mediator (if you want to deploy HA pairs).

**Select the role:**

You must select the Connector IAM role when you launch the Connector instance. You can select the IAM roles for Cloud Volumes ONTAP when you create a Cloud Volumes ONTAP working environment from BlueXP.

For HA pairs, you can select the IAM roles for Cloud Volumes ONTAP and the HA mediator when you create a Cloud Volumes ONTAP working environment from BlueXP.

**Step 3: Set up the AWS KMS**

If you want to use Amazon encryption with Cloud Volumes ONTAP, ensure that requirements are met for the AWS Key Management Service (KMS).

**Steps**

1. Ensure that an active Customer Master Key (CMK) exists in your account or in another AWS account.

The CMK can be an AWS-managed CMK or a customer-managed CMK.

2. If the CMK is in an AWS account separate from the account where you plan to deploy Cloud Volumes ONTAP, then you need to obtain the ARN of that key.

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

3. Add the IAM role for the Connector instance to the list of key users for a CMK.

This gives BlueXP permissions to use the CMK with Cloud Volumes ONTAP.

#### Step 4: Install the Connector and set up BlueXP

Before you can start using BlueXP to deploy Cloud Volumes ONTAP in AWS, you must install and set up the BlueXP Connector. The Connector enables BlueXP to manage resources and processes within your public cloud environment (this includes Cloud Volumes ONTAP).

##### Steps

1. Obtain a root certificate signed by a certificate authority (CA) in the Privacy Enhanced Mail (PEM) Base-64 encoded X.509 format. Consult your organization's policies and procedures for obtaining the certificate.



For AWS Secret Cloud regions, you should upload the `NSS Root CA 2` certificate, and for Top Secret Cloud, the `Amazon Root CA 4` certificate. Ensure that you upload only these certificates and not the entire chain. The file for the certificate chain is large, and the upload can fail. If you have additional certificates, you can upload them later, as described in the next step.

You'll need to upload the certificate during the setup process. BlueXP uses the trusted certificate when sending requests to AWS over HTTPS.

2. Launch the Connector instance:
  - a. Go to the AWS Intelligence Community Marketplace page for BlueXP.
  - b. On the Custom Launch tab, choose the option to launch the instance from the EC2 console.
  - c. Follow the prompts to configure the instance.

Note the following as you configure the instance:

- We recommend `t3.xlarge`.
- You must choose the IAM role that you created when you set up permissions.
- You should keep the default storage options.
- The required connection methods for the Connector are as follows: SSH, HTTP, and HTTPS.

3. Set up BlueXP from a host that has a connection to the Connector instance:
  - a. Open a web browser and enter `https://ipaddress` where `ipaddress` is the IP address of the Linux host where you installed the Connector.
  - b. Specify a proxy server for connectivity to AWS services.
  - c. Upload the certificate that you obtained in step 1.
  - d. Select **Set Up New BlueXP** and follow the prompts to set up the system.
    - **System Details:** Enter a name for the Connector and your company name.
    - **Create Admin User:** Create the admin user for the system.

This user account runs locally on the system. There's no connection to the `auth0` service available through BlueXP.

- **Review:** Review the details, accept the license agreement, and then select **Set Up**.

- e. To complete installation of the CA-signed certificate, restart the Connector instance from the EC2 console.
4. After the Connector restarts, log in using the administrator user account that you created in the Setup wizard.

### Step 5: (optional) Install a private mode certificate

This step is optional for AWS Secret Cloud and Top Secret Cloud regions, and is required only if you have additional certificates apart from the root certificates that you installed in the previous step.

#### Steps

1. List existing installed certificates.

- a. To collect the occm container docker id (identified name "ds-occm-1"), run the following command:

```
docker ps
```

- b. To get inside occm container, run the following command:

```
docker exec -it <docker-id> /bin/sh
```

- c. To collect the password from "TRUST\_STORE\_PASSWORD" environment variable, run the following command:

```
env
```

- d. To list all installed certificates in truststore, run the following command and use the password collected in the previous step:

```
keytool -list -v -keystore occm.truststore
```

2. Add a certificate.

- a. To collect occm container docker id (identified name "ds-occm-1"), run the following command:

```
docker ps
```

- b. To get inside occm container, run the following command:

```
docker exec -it <docker-id> /bin/sh
```

Save the new certificate file inside.

- c. To collect the password from "TRUST\_STORE\_PASSWORD" environment variable, run the following

command:

```
env
```

- d. To add the certificate to the truststore, run the following command and use the password from the previous step:

```
keytool -import -alias <alias-name> -file <certificate-file-name>  
-keystore occm.truststore
```

- e. To check that the certificate installed, run the following command:

```
keytool -list -v -keystore occm.truststore -alias <alias-name>
```

- f. To exit occm container, run the following command:

```
exit
```

- g. To reset occm container, run the following command:

```
docker restart <docker-id>
```

## Step 6: Add a license to the BlueXP digital wallet

If you purchased a license from NetApp, you need to add it to the BlueXP digital wallet so that you can select the license when you create a new Cloud Volumes ONTAP system. The digital wallet identifies these licenses as unassigned.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. On the **Cloud Volumes ONTAP** tab, select **Node Based Licenses** from the drop-down.
3. Click **Unassigned**.
4. Click **Add Unassigned Licenses**.
5. Enter the serial number of the license or upload the license file.
6. If you don't have the license file yet, you'll need to manually upload the license file from netapp.com.
  - a. Go to the [NetApp License File Generator](#) and log in using your NetApp Support Site credentials.
  - b. Enter your password, choose your product, enter the serial number, confirm that you have read and accepted the privacy policy, and then click **Submit**.
  - c. Choose whether you want to receive the serialnumber.NLF JSON file through email or direct download.
7. Click **Add License**.



## Result

BlueXP adds the license to the digital wallet. The license will be identified as unassigned until you associate it with a new Cloud Volumes ONTAP system. After that happens, the license moves to the BYOL tab in the digital wallet.

## Step 7: Launch Cloud Volumes ONTAP from BlueXP

You can launch Cloud Volumes ONTAP instances in AWS Secret Cloud and Top Secret Cloud by creating new working environments in BlueXP.

### Before you begin

For HA pairs, a key pair is required to enable key-based SSH authentication to the HA mediator.

### Steps

1. On the Working Environments page, click **Add Working Environment**.
2. Under **Create**, select Cloud Volumes ONTAP.

For HA: Under **Create**, select Cloud Volumes ONTAP or Cloud Volumes ONTAP HA.

3. Complete the steps in the wizard to launch the Cloud Volumes ONTAP system.



While making selections through the wizard, do not select **Data Sense & Compliance** and **Backup to Cloud** under **Services**. Under **Preconfigured Packages**, select **Change Configuration** only, and ensure that you haven't selected any other option. Preconfigured packages aren't supported in AWS Secret Cloud and Top Secret Cloud regions, and if selected, your deployment will fail.

### Notes for deploying Cloud Volumes ONTAP HA in multiple Availability Zones

Note the following as you complete the wizard for HA pairs.

- You should configure a transit gateway when you deploy Cloud Volumes ONTAP HA in multiple Availability Zones (AZs). For instructions, refer to [Set up an AWS transit gateway](#).
- Deploy the configuration as the following because only two AZs were available in the AWS Top Secret Cloud at the time of publication:
  - Node 1: Availability Zone A
  - Node 2: Availability Zone B
  - Mediator: Availability Zone A or B

### Notes for deploying Cloud Volumes ONTAP in both single and HA nodes

Note the following as you complete the wizard:

- You should leave the default option to use a generated security group.

The predefined security group includes the rules that Cloud Volumes ONTAP needs to operate successfully. If you have a requirement to use your own, you can refer to the security group section below.

- You must choose the IAM role that you created when preparing your AWS environment.
- The underlying AWS disk type is for the initial Cloud Volumes ONTAP volume.

You can choose a different disk type for subsequent volumes.

- The performance of AWS disks is tied to disk size.

You should choose the disk size that gives you the sustained performance that you need. Refer to AWS documentation for more details about EBS performance.

- The disk size is the default size for all disks on the system.



If you need a different size later, you can use the Advanced allocation option to create an aggregate that uses disks of a specific size.

## Result

BlueXP launches the Cloud Volumes ONTAP instance. You can track the progress in the timeline.

## Step 8: Install security certificates for data tiering

You need to manually install security certificates for enabling data tiering in AWS Secret Cloud and Top Secret Cloud regions.

### Before you begin

1. Create S3 buckets.



Ensure that the bucket names are prefixed with `fabric-pool-`. For example `fabric-pool-testbucket`.

2. Keep the root certificates that you installed in `step 4` handy.

### Steps

1. Copy the text from the root certificates that you installed in `step 4`.
2. Securely connect to the Cloud Volumes ONTAP system by using the CLI.
3. Install the root certificates. You might need to press the `ENTER` key multiple times:

```
security certificate install -type server-ca -cert-name <certificate-name>
```

4. When prompted, enter the entire copied text, including and from `----- BEGIN CERTIFICATE -----` to `----- END CERTIFICATE -----`.
5. Keep a copy of the CA-signed digital certificate for future reference.
6. Retain the CA name and certificate serial number.
7. Configure the object store for AWS Secret Cloud and Top Secret Cloud regions: `set -privilege advanced -confirmations off`
8. Run this command to configure the object store.



All Amazon Resource Names (ARNs) should be suffixed with `-iso-b`, such as `arn:aws-iso-b`. For example, if a resource requires an ARN with a region, for Top Secret Cloud, use the naming convention as `us-iso-b` for the `-server` flag. For AWS Secret Cloud, use `us-iso-b-1`.

```
storage aggregate object-store config create -object-store-name
<S3Bucket> -provider-type AWS_S3 -auth-type EC2-IAM -server <s3.us-iso-
b-1.server_name> -container-name <fabric-pool-testbucket> -is-ssl
-enabled true -port 443
```

9. Verify that the object store was created successfully: `storage aggregate object-store show -instance`
10. Attach the object store to the aggregate. This should be repeated for every new aggregate: `storage aggregate object-store attach -aggregate <aggr1> -object-store-name <S3Bucket>`

## Get started in Microsoft Azure

### Learn about Cloud Volumes ONTAP deployment options in Azure

NetApp provides two options for deploying Cloud Volumes ONTAP on Azure. Cloud Volumes ONTAP traditionally relies on BlueXP for deployment and orchestration. Beginning with Cloud Volumes ONTAP 9.16.1, you can take advantage of Azure marketplace direct deployment, a streamlined process that provides access to a limited, but still powerful set of Cloud Volumes ONTAP features and options.

When you deploy Cloud Volumes ONTAP directly from the Azure marketplace, you're not required to set up the BlueXP Connector or meet other security and onboarding criteria required for deploying Cloud Volumes ONTAP through BlueXP. From the Azure marketplace, you can quickly deploy Cloud Volumes ONTAP in a few clicks and explore its core features and capabilities in your environment.

On completing the deployment in the Azure marketplace, you can discover these systems in BlueXP. After discovery, you can manage them as working environments and take advantage of all the BlueXP capabilities. Refer to [Discover the deployed systems in BlueXP](#).

Here is the feature comparison between the two options. Note that the features of a standalone instance deployed through the Azure marketplace change when it is discovered in BlueXP.

	Azure marketplace	BlueXP
<b>Onboarding</b>	Shorter and easier, minimal preparation required for direct deployment	Longer onboarding process, including the installation of the BlueXP Connector
<b>Supported virtual machine (VM) types</b>	Eds_v5 and Ls_v3 instance types	Full range of VM types. <a href="#">Supported configurations in Azure</a>
<b>License</b>	Free license	Any capacity-based license. <a href="#">Cloud Volumes ONTAP licensing</a>
<b>NetApp support</b>	Not included	Available, based on the license type
<b>Capacity</b>	Up to 500 GiB	Expandable by configuration
<b>Deployment model</b>	High-availability (HA) mode deployment in single availability zone (AZ)	All supported configurations, including single node and HA modes, single and multiple AZ deployments

	Azure marketplace	BlueXP
<b>Supported disk type</b>	Premium SSD v2 Managed Disks	Wider support. <a href="#">Default configuration for Cloud Volumes ONTAP</a>
<b>Write speed (fast write mode)</b>	Not supported	Supported, based on your configuration. <a href="#">Learn about write speeds in Cloud Volumes ONTAP.</a>
<b>Orchestration capabilities</b>	Not available	Available through BlueXP, based on the license type
<b>Number of supported storage VMs</b>	One per deployment	Multiple storage VMs, based on your configuration. <a href="#">Supported number of storage VMs</a>
<b>Changing the instance type</b>	Not supported	Supported
<b>FabricPool tiering</b>	Not supported	Supported

**Related links**

- Azure marketplace direct deployment: [Deploy Cloud Volumes ONTAP from the Azure marketplace](#)
- BlueXP deployment: [Quick start for Cloud Volumes ONTAP in Azure](#)
- [BlueXP documentation](#)

**Get started in BlueXP**

**Quick start for Cloud Volumes ONTAP in Azure**

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

**1**

**Create a Connector**

If you don't have a [Connector](#) yet, you need to create one. [Learn how to create a Connector in Azure](#)

Note that if you want to deploy Cloud Volumes ONTAP in a subnet where no internet access is available, then you need to manually install the Connector and access the BlueXP user interface that's running on that Connector. [Learn how to manually install the Connector in a location without internet access](#)

**2**

**Plan your configuration**

BlueXP offers preconfigured packages that match your workload requirements, or you can create your own configuration. If you choose your own configuration, you should understand the options available to you. For information, refer to [Plan your Cloud Volumes ONTAP configuration in Azure.](#)

**3**

**Set up your networking**

- Ensure that your VNet and subnets will support connectivity between the Connector and Cloud Volumes ONTAP.

b. Enable outbound internet access from the target VPC for NetApp AutoSupport.

This step isn't required if you're deploying Cloud Volumes ONTAP in a location where no internet access is available.

[Learn more about networking requirements.](#)



#### **Launch Cloud Volumes ONTAP using BlueXP**

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

#### **Related links**

- [Creating a Connector from BlueXP](#)
- [Creating a Connector from the Azure Marketplace](#)
- [Installing the Connector software on a Linux host](#)
- [What BlueXP does with permissions](#)

#### **Plan your Cloud Volumes ONTAP configuration in Azure**

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

#### **Choose a Cloud Volumes ONTAP license**

Several licensing options are available for Cloud Volumes ONTAP. Each option enables you to choose a consumption model that meets your needs.

- [Learn about licensing options for Cloud Volumes ONTAP](#)
- [Learn how to set up licensing](#)

#### **Choose a supported region**

Cloud Volumes ONTAP is supported in most Microsoft Azure regions. [View the full list of supported regions.](#)

#### **Choose a supported VM type**

Cloud Volumes ONTAP supports several VM types, depending on the license type that you choose.

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

#### **Understand storage limits**

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

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

## Size your system in Azure

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

### Virtual machine type

Look at the supported virtual machine types in the [Cloud Volumes ONTAP Release Notes](#) and then review details about each supported VM type. Be aware that each VM type supports a specific number of data disks.

- [Azure documentation: General purpose virtual machine sizes](#)
- [Azure documentation: Memory optimized virtual machine sizes](#)

### Azure disk type with single node systems

When you create volumes for Cloud Volumes ONTAP, you need to choose the underlying cloud storage that Cloud Volumes ONTAP uses as a disk.

Single node systems can use these types of Azure Managed Disks:

- *Premium SSD Managed Disks* provide high performance for I/O-intensive workloads at a higher cost.
- *Premium SSD v2 Managed Disks* provide higher performance with lower latency at a lower cost, compared to Premium SSD Managed Disks.
- *Standard SSD Managed Disks* provide consistent performance for workloads that require low IOPS.
- *Standard HDD Managed Disks* are a good choice if you don't need high IOPS and want to reduce your costs.

For additional details about the use cases for these disks, refer to [Microsoft Azure Documentation: What disk types are available in Azure?](#)

### Azure disk type with HA pairs

HA systems use Premium SSD Shared Managed Disks which both provide high performance for I/O-intensive workloads at a higher cost. HA deployments created before the 9.12.1 release use Premium page blobs.

### Azure disk size

When you launch Cloud Volumes ONTAP instances, you must choose the default disk size for aggregates. BlueXP uses this disk size for the initial aggregate, and for any additional aggregates that it creates when you use the simple provisioning option. You can create aggregates that use a disk size different from the default by [using the advanced allocation option](#).



All disks in an aggregate must be the same size.

When choosing a disk size, you should take several factors into consideration. The disk size impacts how much you pay for storage, the size of volumes that you can create in an aggregate, the total capacity available to Cloud Volumes ONTAP, and storage performance.

The performance of Azure Premium Storage is tied to the disk size. Larger disks provide higher IOPS and throughput. For example, choosing 1 TiB disks can provide better performance than 500 GiB disks, at a higher cost.

There are no performance differences between disk sizes for Standard Storage. You should choose disk size based on the capacity that you need.

Refer to Azure for IOPS and throughput by disk size:

- [Microsoft Azure: Managed Disks pricing](#)
- [Microsoft Azure: Page Blobs pricing](#)

### View default system disks

In addition to the storage for user data, BlueXP also purchases cloud storage for Cloud Volumes ONTAP system data (boot data, root data, core data, and NVRAM). For planning purposes, it might help for you to review these details before you deploy Cloud Volumes ONTAP.

[View the default disks for Cloud Volumes ONTAP system data in Azure.](#)



The Connector also requires a system disk. [View details about the Connector's default configuration.](#)

### Collect networking information

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

Azure information	Your value
Region	
Virtual network (VNet)	
Subnet	
Network security group (if using your own)	

### Choose a write speed

BlueXP enables you to choose a write speed setting for Cloud Volumes ONTAP. Before you choose a write speed, you should understand the differences between the normal and high settings and risks and recommendations when using high write speed. [Learn more about write speed.](#)

### Choose a volume usage profile

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

NetApp storage efficiency features provide the following benefits:

#### Thin provisioning

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

#### Deduplication

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

## Compression

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

## Set up Azure networking for Cloud Volumes ONTAP

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

### Requirements for Cloud Volumes ONTAP

The following networking requirements must be met in Azure.

#### Outbound internet access

Cloud Volumes ONTAP systems require outbound internet access for accessing external endpoints for various functions. Cloud Volumes ONTAP can't operate properly if these endpoints are blocked in environments with strict security requirements.

The BlueXP Connector also contacts several endpoints for day-to-day operations, as well as the BlueXP web-based console. For information about the BlueXP endpoints, refer to [View endpoints contacted from the Connector](#) and [Prepare networking for using the BlueXP console](#).

#### Cloud Volumes ONTAP endpoints

Cloud Volumes ONTAP uses these endpoints to communicate with various services.

Endpoints	Applicable for	Purpose	BlueXP deployment modes	Impact if unavailable
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a>	Authentication	Used for BlueXP authentication.	Standard and restricted modes.	User authentication fails and the following services remain unavailable: <ul style="list-style-type: none"><li>• Cloud Volumes ONTAP services</li><li>• ONTAP services</li><li>• Protocols and proxy services</li></ul>
<a href="https://vault.azure.net">https://vault.azure.net</a>	Key Vault	Used to retrieve client secret keys from the Azure Key Vault when using customer-managed keys (CMK).	Standard, restricted, and private modes.	Cloud Volumes ONTAP services are unavailable.
<a href="https://cloudmanager.cloud.netapp.com/tenancy">https://cloudmanager.cloud.netapp.com/tenancy</a>	Tenancy	Used to retrieve the Cloud Volumes ONTAP resources from BlueXP tenancy to authorize resources and users.	Standard and restricted modes.	Cloud Volumes ONTAP resources and the users are not authorized.



Endpoints	Applicable for	Purpose	BlueXP deployment modes	Impact if unavailable
<a href="https://support.netapp.com/ods/asupmessage">https://support.netapp.com/ods/asupmessage</a> <a href="https://support.netapp.com/asupprod/post/1.0/postAsup">https://support.netapp.com/asupprod/post/1.0/postAsup</a>	AutoSupport	Used to send AutoSupport telemetry data to NetApp support.	Standard and restricted modes.	AutoSupport information remains undelivered.
<a href="https://management.azure.com">https://management.azure.com</a> <a href="https://login.microsoftonline.com">https://login.microsoftonline.com</a> <a href="https://blob.core.windows.net">https://blob.core.windows.net</a> <a href="https://core.windows.net">https://core.windows.net</a>	Public regions	Communication with Azure services.	Standard, restricted, and private modes.	Cloud Volumes ONTAP cannot communicate with Azure service to perform specific BlueXP operations in Azure.
<a href="https://management.chinacloudapi.cn">https://management.chinacloudapi.cn</a> <a href="https://login.chinacloudapi.cn">https://login.chinacloudapi.cn</a> <a href="https://blob.core.chinacloudapi.cn">https://blob.core.chinacloudapi.cn</a> <a href="https://core.chinacloudapi.cn">https://core.chinacloudapi.cn</a>	China Region	Communication with Azure services.	Standard, restricted, and private modes.	Cloud Volumes ONTAP cannot communicate with Azure service to perform specific BlueXP operations in Azure.
<a href="https://management.microsoftazure.de">https://management.microsoftazure.de</a> <a href="https://login.microsoftonline.de">https://login.microsoftonline.de</a> <a href="https://blob.core.cloudapi.de">https://blob.core.cloudapi.de</a> <a href="https://core.cloudapi.de">https://core.cloudapi.de</a>	Germany Region	Communication with Azure services.	Standard, restricted, and private modes.	Cloud Volumes ONTAP cannot communicate with Azure service to perform specific BlueXP operations in Azure.
<a href="https://management.usgovcloudapi.net">https://management.usgovcloudapi.net</a> <a href="https://login.microsoftonline.us">https://login.microsoftonline.us</a> <a href="https://blob.core.usgovcloudapi.net">https://blob.core.usgovcloudapi.net</a> <a href="https://core.usgovcloudapi.net">https://core.usgovcloudapi.net</a>	Government regions	Communication with Azure services.	Standard, restricted, and private modes.	Cloud Volumes ONTAP cannot communicate with Azure service to perform specific BlueXP operations in Azure.
<a href="https://management.azure.microsoft.scloud">https://management.azure.microsoft.scloud</a> <a href="https://login.microsoftonline.microsoft.scloud">https://login.microsoftonline.microsoft.scloud</a> <a href="https://blob.core.microsoft.scloud">https://blob.core.microsoft.scloud</a> <a href="https://core.microsoft.scloud">https://core.microsoft.scloud</a>	Government DoD regions	Communication with Azure services.	Standard, restricted, and private modes.	Cloud Volumes ONTAP cannot communicate with Azure service to perform specific BlueXP operations in Azure.

## Network configurations to support Connector proxy

You can use the proxy servers configured for the BlueXP Connector to enable outbound internet access from Cloud Volumes ONTAP. BlueXP supports two types of proxies:

- **Explicit proxy:** The outbound traffic from Cloud Volumes ONTAP uses the HTTP address of the proxy server specified during the Connector proxy configuration. The Connector administrator might also have

configured user credentials and root CA certificates for additional authentication. If a root CA certificate is available for the explicit proxy, make sure to obtain and upload the same certificate to your Cloud Volumes ONTAP working environment using the [ONTAP CLI: security certificate install](#) command.

- **Transparent proxy:** The network is configured to automatically route outbound traffic from Cloud Volumes ONTAP through the Connector proxy. When setting up a transparent proxy, the Connector administrator needs to provide only a root CA certificate for connectivity from Cloud Volumes ONTAP, not the HTTP address of the proxy server. Make sure that you obtain and upload the same root CA certificate to your Cloud Volumes ONTAP working environment using the [ONTAP CLI: security certificate install](#) command.

For information about configuring proxy servers for the BlueXP Connector, refer to the [Configure a Connector to use a proxy server](#).

## IP addresses

BlueXP automatically allocates the required number of private IP addresses to Cloud Volumes ONTAP in Azure. You need to make sure that your networking has enough private IP addresses available.

The number of LIFs that BlueXP allocates for Cloud Volumes ONTAP depends on whether you deploy a single node system or an HA pair. A LIF is an IP address associated with a physical port. An SVM management LIF is required for management tools like SnapCenter.



An iSCSI LIF provides client access over the iSCSI protocol and is used by the system for other important networking workflows. These LIFs are required and should not be deleted.

### IP addresses for a single node system

BlueXP allocates 5 or 6 IP addresses to a single node system:

- Cluster management IP
- Node management IP
- Intercluster IP for SnapMirror
- NFS/CIFS IP
- iSCSI IP



The iSCSI IP provides client access over the iSCSI protocol. It is also used by the system for other important networking workflows. This LIF is required and should not be deleted.

- SVM management (optional - not configured by default)

### IP addresses for HA pairs

BlueXP allocates IP addresses to 4 NICs (per node) during deployment.

Note that BlueXP creates an SVM management LIF on HA pairs, but not on single node systems in Azure.

### NIC0

- Node management IP
- Intercluster IP
- iSCSI IP



The iSCSI IP provides client access over the iSCSI protocol. It is also used by the system for other important networking workflows. This LIF is required and should not be deleted.

### NIC1

- Cluster network IP

### NIC2

- Cluster Interconnect IP (HA IC)

### NIC3

- Pageblob NIC IP (disk access)



NIC3 is only applicable to HA deployments that use page blob storage.

The above IP addresses do not migrate on failover events.

Additionally, 4 frontend IPs (FIPs) are configured to migrate on failover events. These frontend IPs live in the load balancer.

- Cluster management IP
- NodeA data IP (NFS/CIFS)
- NodeB data IP (NFS/CIFS)
- SVM management IP

### Secure connections to Azure services

By default, BlueXP enables an Azure Private Link for connections between Cloud Volumes ONTAP and Azure page blob storage accounts.

In most cases, there's nothing that you need to do—BlueXP manages the Azure Private Link for you. But if you use Azure Private DNS, then you'll need to edit a configuration file. You should also be aware of a requirement for the Connector location in Azure.

You can also disable the Private Link connection, if required by your business needs. If you disable the link, BlueXP configures Cloud Volumes ONTAP to use a service endpoint instead.

[Learn more about using Azure Private Links or service endpoints with Cloud Volumes ONTAP.](#)

### Connections to other ONTAP systems

To replicate data between a Cloud Volumes ONTAP system in Azure and ONTAP systems in other networks, you must have a VPN connection between the Azure VNet and the other network—for example, your corporate network.

For instructions, refer to [Microsoft Azure Documentation: Create a Site-to-Site connection in the Azure portal.](#)

### Port for the HA interconnect

A Cloud Volumes ONTAP HA pair includes an HA interconnect, which allows each node to continually check

whether its partner is functioning and to mirror log data for the other's nonvolatile memory. The HA interconnect uses TCP port 10006 for communication.

By default, communication between the HA interconnect LIFs is open and there are no security group rules for this port. But if you create a firewall between the HA interconnect LIFs, then you need to ensure that TCP traffic is open for port 10006 so that the HA pair can operate properly.

### Only one HA pair in an Azure resource group

You must use a *dedicated* resource group for each Cloud Volumes ONTAP HA pair that you deploy in Azure. Only one HA pair is supported in a resource group.

BlueXP experiences connection issues if you try to deploy a second Cloud Volumes ONTAP HA pair in an Azure resource group.

### Security group rules

BlueXP creates Azure security groups that include the inbound and outbound rules for Cloud Volumes ONTAP to operate successfully. [View security group rules for the Connector](#).

The Azure security groups for Cloud Volumes ONTAP require the appropriate ports to be open for internal communication between the nodes. [Learn about ONTAP internal ports](#).

We do not recommend modifying the predefined security groups or using custom security groups. However, if you must, note that the deployment process requires the Cloud Volumes ONTAP system to have full access within its own subnet. After the deployment is complete, if you decide to modify the network security group, ensure to keep the cluster ports and HA network ports open. This ensures seamless communication within the Cloud Volumes ONTAP cluster (any-to-any communication between the nodes).

### Inbound rules for single node systems

When you create a working environment and choose a predefined security group, you can choose to allow traffic within one of the following:

- **Selected VNet only:** The source for inbound traffic is the subnet range of the VNet for the Cloud Volumes ONTAP system and the subnet range of the VNet where the Connector resides. This is the recommended option.
- **All VNets:** The source for inbound traffic is the 0.0.0.0/0 IP range.
- **Disabled:** This option restricts the public network access to your storage account, and disables data tiering for Cloud Volumes ONTAP systems. This is a recommended option if your private IP addresses should not be exposed even within the same VNet due to security regulations and policies.

Priority and name	Port and protocol	Source and destination	Description
1000 inbound_ssh	22 TCP	Any to Any	SSH access to the IP address of the cluster management LIF or a node management LIF
1001 inbound_http	80 TCP	Any to Any	HTTP access to the ONTAP System Manager web console using the IP address of the cluster management LIF
1002 inbound_111_tcp	111 TCP	Any to Any	Remote procedure call for NFS

Priority and name	Port and protocol	Source and destination	Description
1003 inbound_111_udp	111 UDP	Any to Any	Remote procedure call for NFS
1004 inbound_139	139 TCP	Any to Any	NetBIOS service session for CIFS
1005 inbound_161-162_tcp	161-162 TCP	Any to Any	Simple network management protocol
1006 inbound_161-162_udp	161-162 UDP	Any to Any	Simple network management protocol
1007 inbound_443	443 TCP	Any to Any	Connectivity with the Connector and HTTPS access to the ONTAP System Manager web console using the IP address of the cluster management LIF
1008 inbound_445	445 TCP	Any to Any	Microsoft SMB/CIFS over TCP with NetBIOS framing
1009 inbound_635_tcp	635 TCP	Any to Any	NFS mount
1010 inbound_635_udp	635 UDP	Any to Any	NFS mount
1011 inbound_749	749 TCP	Any to Any	Kerberos
1012 inbound_2049_tcp	2049 TCP	Any to Any	NFS server daemon
1013 inbound_2049_udp	2049 UDP	Any to Any	NFS server daemon
1014 inbound_3260	3260 TCP	Any to Any	iSCSI access through the iSCSI data LIF
1015 inbound_4045-4046_tcp	4045-4046 TCP	Any to Any	NFS lock daemon and network status monitor
1016 inbound_4045-4046_udp	4045-4046 UDP	Any to Any	NFS lock daemon and network status monitor
1017 inbound_10000	10000 TCP	Any to Any	Backup using NDMP
1018 inbound_11104-11105	11104-11105 TCP	Any to Any	SnapMirror data transfer

Priority and name	Port and protocol	Source and destination	Description
3000 inbound_deny_all_tcp	Any port TCP	Any to Any	Block all other TCP inbound traffic
3001 inbound_deny_all_udp	Any port UDP	Any to Any	Block all other UDP inbound traffic
65000 AllowVnetInBound	Any port Any protocol	VirtualNetwork to VirtualNetwork	Inbound traffic from within the VNet
65001 AllowAzureLoadBalancerInBound	Any port Any protocol	AzureLoadBalancer to Any	Data traffic from the Azure Standard Load Balancer
65500 DenyAllInBound	Any port Any protocol	Any to Any	Block all other inbound traffic

### Inbound rules for HA systems

When you create a working environment and choose a predefined security group, you can choose to allow traffic within one of the following:

- **Selected VNet only:** The source for inbound traffic is the subnet range of the VNet for the Cloud Volumes ONTAP system and the subnet range of the VNet where the Connector resides. This is the recommended option.
- **All VNets:** The source for inbound traffic is the 0.0.0.0/0 IP range.



HA systems have less inbound rules than single node systems because inbound data traffic goes through the Azure Standard Load Balancer. Because of this, traffic from the Load Balancer should be open, as shown in the "AllowAzureLoadBalancerInBound" rule.

- **Disabled:** This option restricts the public network access to your storage account, and disables data tiering for Cloud Volumes ONTAP systems. This is a recommended option if your private IP addresses should not be exposed even within the same VNet due to security regulations and policies.

Priority and name	Port and protocol	Source and destination	Description
100 inbound_443	443 Any protocol	Any to Any	Connectivity with the Connector and HTTPS access to the ONTAP System Manager web console using the IP address of the cluster management LIF
101 inbound_111_tcp	111 Any protocol	Any to Any	Remote procedure call for NFS
102 inbound_2049_tcp	2049 Any protocol	Any to Any	NFS server daemon

Priority and name	Port and protocol	Source and destination	Description
111 inbound_ssh	22 Any protocol	Any to Any	SSH access to the IP address of the cluster management LIF or a node management LIF
121 inbound_53	53 Any protocol	Any to Any	DNS and CIFS
65000 AllowVnetInBound	Any port Any protocol	VirtualNetwork to VirtualNetwork	Inbound traffic from within the VNet
65001 AllowAzureLoad BalancerInBound	Any port Any protocol	AzureLoadBalan cer to Any	Data traffic from the Azure Standard Load Balancer
65500 DenyAllInBound	Any port Any protocol	Any to Any	Block all other inbound traffic

## Outbound rules

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

### Basic outbound rules

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

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

### Advanced outbound rules

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



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

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



Service	Port	Protocol	Source	Destination	Purpose
AutoSupport	HTTPS	443	Node management LIF	support.netapp.com	AutoSupport (HTTPS is the default)
	HTTP	80	Node management LIF	support.netapp.com	AutoSupport (only if the transport protocol is changed from HTTPS to HTTP)
	TCP	3128	Node management LIF	Connector	Sending AutoSupport messages through a proxy server on the Connector, if an outbound internet connection isn't available
Configuration backups	HTTP	80	Node management LIF	http://<connector-IP-address>/occm/offbo xconfig	Send configuration backups to the Connector. <a href="#">ONTAP documentation</a> .
DHCP	68	UDP	Node management LIF	DHCP	DHCP client for first-time setup
DHCPS	67	UDP	Node management LIF	DHCP	DHCP server
DNS	53	UDP	Node management LIF and data LIF (NFS, CIFS)	DNS	DNS
NDMP	18600–18699	TCP	Node management LIF	Destination servers	NDMP copy
SMTP	25	TCP	Node management LIF	Mail server	SMTP alerts, can be used for AutoSupport
SNMP	161	TCP	Node management LIF	Monitor server	Monitoring by SNMP traps
	161	UDP	Node management LIF	Monitor server	Monitoring by SNMP traps
	162	TCP	Node management LIF	Monitor server	Monitoring by SNMP traps
	162	UDP	Node management LIF	Monitor server	Monitoring by SNMP traps
SnapMirror	11104	TCP	Intercluster LIF	ONTAP intercluster LIFs	Management of intercluster communication sessions for SnapMirror
	11105	TCP	Intercluster LIF	ONTAP intercluster LIFs	SnapMirror data transfer
Syslog	514	UDP	Node management LIF	Syslog server	Syslog forward messages

### Requirements for the Connector

If you haven't created a Connector yet, you should review networking requirements for the Connector as well.

- [View networking requirements for the Connector](#)
- [Security group rules in Azure](#)

### Related topics

- [Verify AutoSupport setup for Cloud Volumes ONTAP](#)
- [Learn about ONTAP internal ports.](#)

## Set up Cloud Volumes ONTAP to use a customer-managed key in Azure

Data is automatically encrypted on Cloud Volumes ONTAP in Azure using Azure Storage Service Encryption with a Microsoft-managed key. But you can use your own encryption key instead by following the steps on this page.

### Data encryption overview

Cloud Volumes ONTAP data is automatically encrypted in Azure using [Azure Storage Service Encryption](#). The default implementation uses a Microsoft-managed key. No setup is required.

If you want to use a customer-managed key with Cloud Volumes ONTAP, then you need to complete the following steps:

1. From Azure, create a key vault and then generate a key in that vault.
2. From BlueXP, use the API to create a Cloud Volumes ONTAP working environment that uses the key.

### How data is encrypted

BlueXP uses a disk encryption set, which enables management of encryption keys with managed disks not page blobs. Any new data disks also use the same disk encryption set. Lower versions will use Microsoft-managed key, instead of the customer-managed key.

After you create a Cloud Volumes ONTAP working environment that is configured to use a customer-managed key, Cloud Volumes ONTAP data is encrypted as follows.

Cloud Volumes ONTAP configuration	System disks used for key encryption	Data disks used for key encryption
Single node	<ul style="list-style-type: none"> <li>• Boot</li> <li>• Core</li> <li>• NVRAM</li> </ul>	<ul style="list-style-type: none"> <li>• Root</li> <li>• Data</li> </ul>
Azure HA single availability zone with page blobs	<ul style="list-style-type: none"> <li>• Boot</li> <li>• Core</li> <li>• NVRAM</li> </ul>	None
Azure HA single availability zone with shared managed disks	<ul style="list-style-type: none"> <li>• Boot</li> <li>• Core</li> <li>• NVRAM</li> </ul>	<ul style="list-style-type: none"> <li>• Root</li> <li>• Data</li> </ul>

Cloud Volumes ONTAP configuration	System disks used for key encryption	Data disks used for key encryption
Azure HA multiple availability zones with shared managed disks	<ul style="list-style-type: none"> <li>• Boot</li> <li>• Core</li> <li>• NVRAM</li> </ul>	<ul style="list-style-type: none"> <li>• Root</li> <li>• Data</li> </ul>

All Azure storage accounts for Cloud Volumes ONTAP are encrypted using a customer-managed key. If you want to encrypt your storage accounts during their creation, you must create and provide the ID of the resource in the Cloud Volumes ONTAP creation request. This applies for all type of deployments. If you do not provide it, the storage accounts still will be encrypted, but BlueXP will first create the storage accounts with Microsoft-managed key encryption and then will update the storage accounts to use the customer-managed key.

### Key rotation in Cloud Volumes ONTAP

When you configure your encryption keys, you must use the Azure portal to set up and enable automatic key rotation. Creating and enabling a new version of encryption keys ensures that Cloud Volumes ONTAP can automatically detect and use the latest key version for encryption, ensuring your data remains secure without the need for manual intervention.

For information about configuring your keys and setting up key rotation, refer to the following Microsoft Azure documentation topics:

- [Configure cryptographic key auto-rotation in Azure Key Vault](#)
- [Azure PowerShell - Enable customer-managed keys](#)



After configuring the keys, ensure that you have selected *Enable auto rotation*, so that Cloud Volumes ONTAP can use the new keys when the previous keys expire. If you don't enable this option on the Azure portal, Cloud Volumes ONTAP can't automatically detect the new keys, which might cause issues with storage provisioning.

### Create a user-assigned managed identity

You have the option to create a resource called a user-assigned managed identity. Doing so allows you to encrypt your storage accounts when you create a Cloud Volumes ONTAP working environment. We recommend creating this resource prior to creating a key vault and generating a key.

The resource has the following ID: `userassignedidentity`.

### Steps

1. In Azure, go to Azure services and select **Managed Identities**.
2. Click **Create**.
3. Provide the following details:
  - **Subscription:** Choose a subscription. We recommend choosing the same subscription as the Connector subscription.
  - **Resource group:** Use an existing resource group or create a new one.
  - **Region:** Optionally, select the same region as the Connector.
  - **Name:** Enter a name for the resource.

4. Optionally, add tags.
5. Click **Create**.

### Create a key vault and generate a key

The key vault must reside in the same Azure subscription and region in which you plan to create the Cloud Volumes ONTAP system.

If you [created a user-assigned managed identity](#), while creating the key vault, you should also create an access policy for the key vault.

### Steps

1. [Create a key vault in your Azure subscription](#).

Note the following requirements for the key vault:

- The key vault must reside in the same region as the Cloud Volumes ONTAP system.
- The following options should be enabled:
  - **Soft-delete** (this option is enabled by default, but must *not* be disabled)
  - **Purge protection**
  - **Azure Disk Encryption for volume encryption** (for single node systems, HA pairs in multiple zones, and HA single AZ deployments)



Usage of Azure customer-managed encryption keys is contingent upon having Azure Disk encryption enabled for the key vault.

- The following option should be enabled if you created a user-assigned managed identity:
    - **Vault access policy**
2. If you selected Vault access policy, click Create to create an access policy for the key vault. If not, skip to step 3.
    - a. Select the following permissions:
      - get
      - list
      - decrypt
      - encrypt
      - unwrap key
      - wrap key
      - verify
      - sign
    - b. Select the user-assigned managed identity (resource) as the principal.
    - c. Review and create the access policy.
  3. [Generate a key in the key vault](#).

Note the following requirements for the key:

- The key type must be **RSA**.
- The recommended RSA key size is **2048**, but other sizes are supported.

### Create a working environment that uses the encryption key

After you create the key vault and generate an encryption key, you can create a new Cloud Volumes ONTAP system that is configured to use the key. These steps are supported by using the BlueXP API.

### Required permissions

If you want to use a customer-managed key with a single node Cloud Volumes ONTAP system, ensure that the BlueXP Connector has the following permissions:

```
"Microsoft.Compute/diskEncryptionSets/read",
"Microsoft.Compute/diskEncryptionSets/write",
"Microsoft.Compute/diskEncryptionSets/delete"
"Microsoft.KeyVault/vaults/deploy/action",
"Microsoft.KeyVault/vaults/read",
"Microsoft.KeyVault/vaults/accessPolicies/write",
"Microsoft.ManagedIdentity/userAssignedIdentities/assign/action"
```

[View the latest list of permissions](#)

### Steps

1. Obtain the list of key vaults in your Azure subscription by using the following BlueXP API call.

For an HA pair: GET /azure/ha/metadata/vaults

For single node: GET /azure/vsa/metadata/vaults

Make note of the **name** and **resourceGroup**. You'll need to specify those values in the next step.

[Learn more about this API call.](#)

2. Obtain the list of keys within the vault by using the following BlueXP API call.

For an HA pair: GET /azure/ha/metadata/keys-vault

For single node: GET /azure/vsa/metadata/keys-vault

Make note of the **keyName**. You'll need to specify that value (along with the vault name) in the next step.

[Learn more about this API call.](#)

3. Create a Cloud Volumes ONTAP system by using the following BlueXP API call.

- a. For an HA pair:

POST /azure/ha/working-environments

The request body must include the following fields:

```
"azureEncryptionParameters": {
  "key": "keyName",
  "vaultName": "vaultName"
}
```



Include the "userAssignedIdentity": " userAssignedIdentityId" field if you created this resource to be used for storage account encryption.

[Learn more about this API call.](#)

b. For a single node system:

POST /azure/vsa/working-environments

The request body must include the following fields:

```
"azureEncryptionParameters": {
  "key": "keyName",
  "vaultName": "vaultName"
}
```



Include the "userAssignedIdentity": " userAssignedIdentityId" field if you created this resource to be used for storage account encryption.

[Learn more about this API call.](#)

## Result

You have a new Cloud Volumes ONTAP system that is configured to use your customer-managed key for data encryption.

## Set up licensing for Cloud Volumes ONTAP in Azure

After you decide which licensing option you want to use with Cloud Volumes ONTAP, a few steps are required before you can choose that licensing option when creating a new working environment.

### Freemium

Select the Freemium offering to use Cloud Volumes ONTAP free of charge with up to 500 GiB of provisioned capacity. [Learn more about the Freemium offering.](#)

### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the pay-as-you-go offering in the Azure Marketplace.

You won't be charged through the marketplace subscription unless you exceed 500 GiB of provisioned capacity, at which time the system is automatically converted to the [Essentials package](#).

**Edit Credentials & Add Subscription**

Associate Subscription to Credentials ⓘ

Credentials  
Managed Service Identity

Azure Subscription  
OCCM Dev (Default)

Marketplace Subscription  
4 marketplace subscription isn't associated with the selected Azure subscription.

+ Add Subscription

Apply Cancel

b. After you return to BlueXP, select **Freemium** when you reach the charging methods page.

**Select Charging Method**

<input type="radio"/>	Professional	By capacity
<input type="radio"/>	Essential	By capacity
<input checked="" type="radio"/>	Freemium (Up to 500 GiB)	By capacity
<input type="radio"/>	Per Node	By node

[View step-by-step instructions to launch Cloud Volumes ONTAP in Azure.](#)

### Capacity-based license

Capacity-based licensing enables you to pay for Cloud Volumes ONTAP per TiB of capacity. Capacity-based licensing is available in the form of a *package*: the Essentials package or the Professional package.

The Essentials and Professional packages are available with the following consumption models:

- A license (bring your own license (BYOL)) purchased from NetApp
- An hourly, pay-as-you-go (PAYGO) subscription from the Azure Marketplace
- An annual contract

[Learn more about capacity-based licensing.](#)

The following sections describe how to get started with each of these consumption models.

## BYOL

Pay upfront by purchasing a license (BYOL) from NetApp to deploy Cloud Volumes ONTAP systems in any cloud provider.



NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAP](#).

## Steps

1. [Contact NetApp Sales to obtain a license](#)
2. [Add your NetApp Support Site account to BlueXP](#)

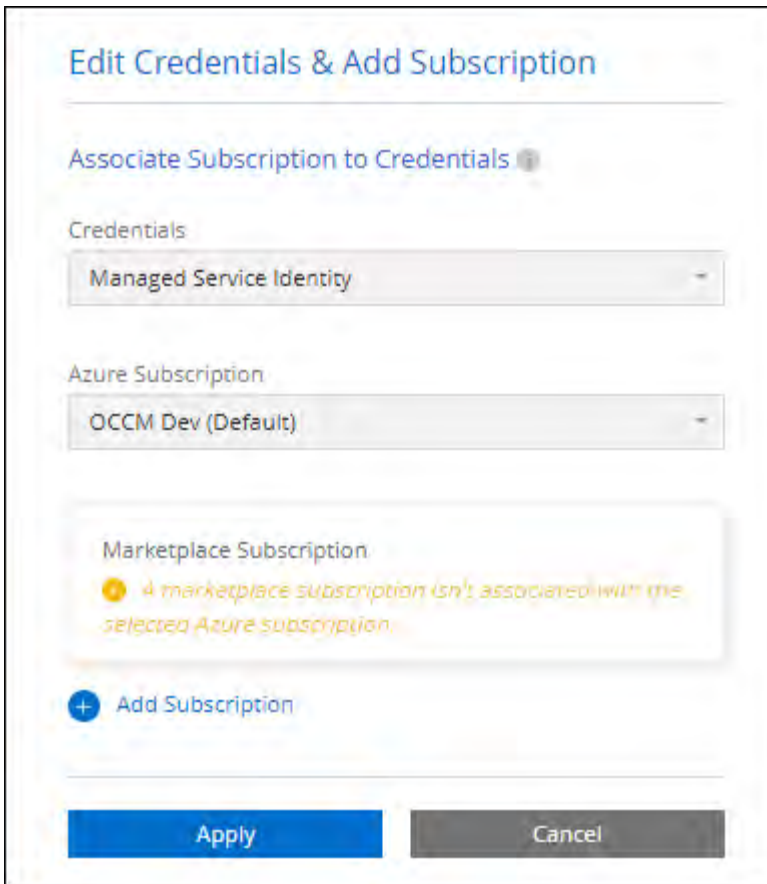
BlueXP automatically queries NetApp's licensing service to obtain details about the licenses associated with your NetApp Support Site account. If there are no errors, BlueXP automatically adds the licenses to the digital wallet.

Your license must be available from the BlueXP digital wallet before you can use it with Cloud Volumes ONTAP. If needed, you can [manually add the license to the BlueXP digital wallet](#).

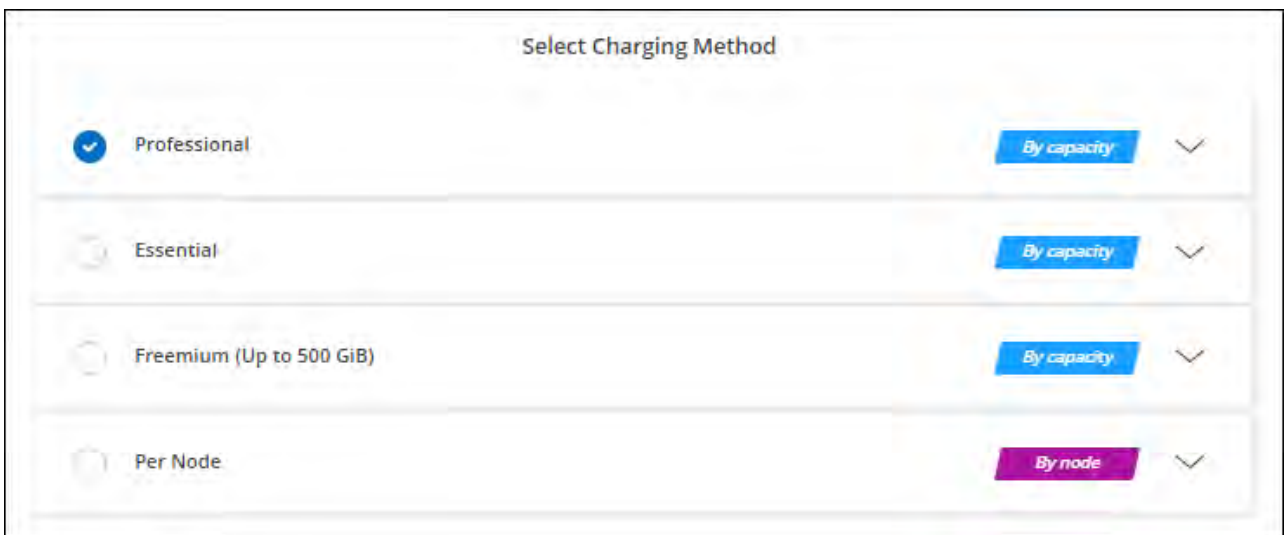
3. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the pay-as-you-go offering in the Azure Marketplace.

The license that you purchased from NetApp is always charged first, but you'll be charged from the hourly rate in the marketplace if you exceed your licensed capacity or if the term of your license expires.





- b. After you return to BlueXP, select a capacity-based package when you reach the charging methods page.



[View step-by-step instructions to launch Cloud Volumes ONTAP in Azure.](#)

### **PAYGO subscription**

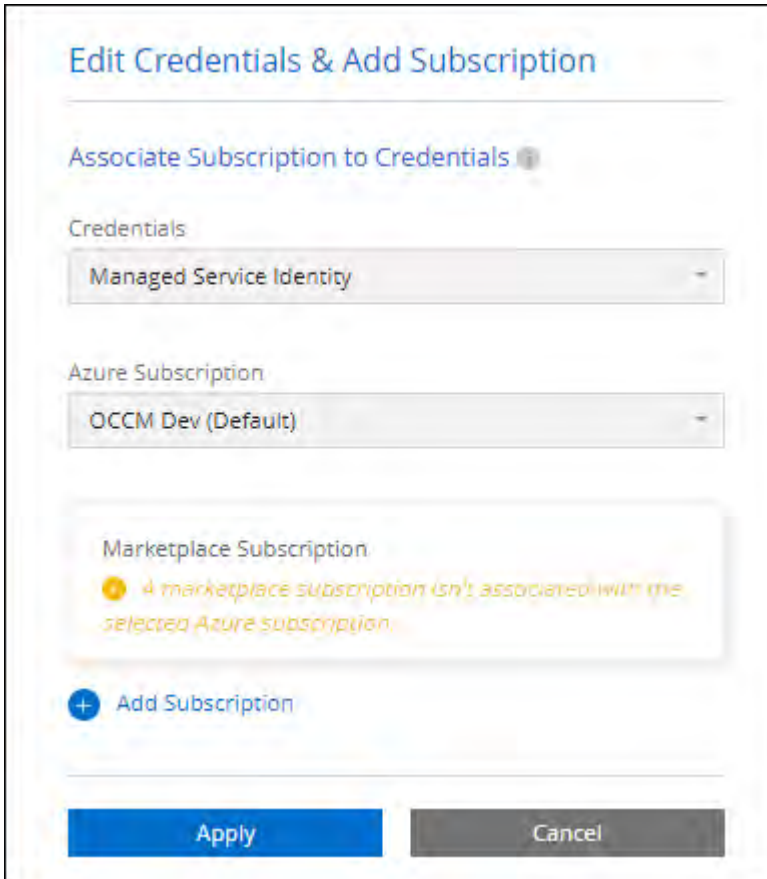
Pay hourly by subscribing to the offer from your cloud provider's marketplace.

When you create a Cloud Volumes ONTAP working environment, BlueXP prompts you to subscribe to the agreement that's available in the Azure Marketplace. That subscription is then associated with the working

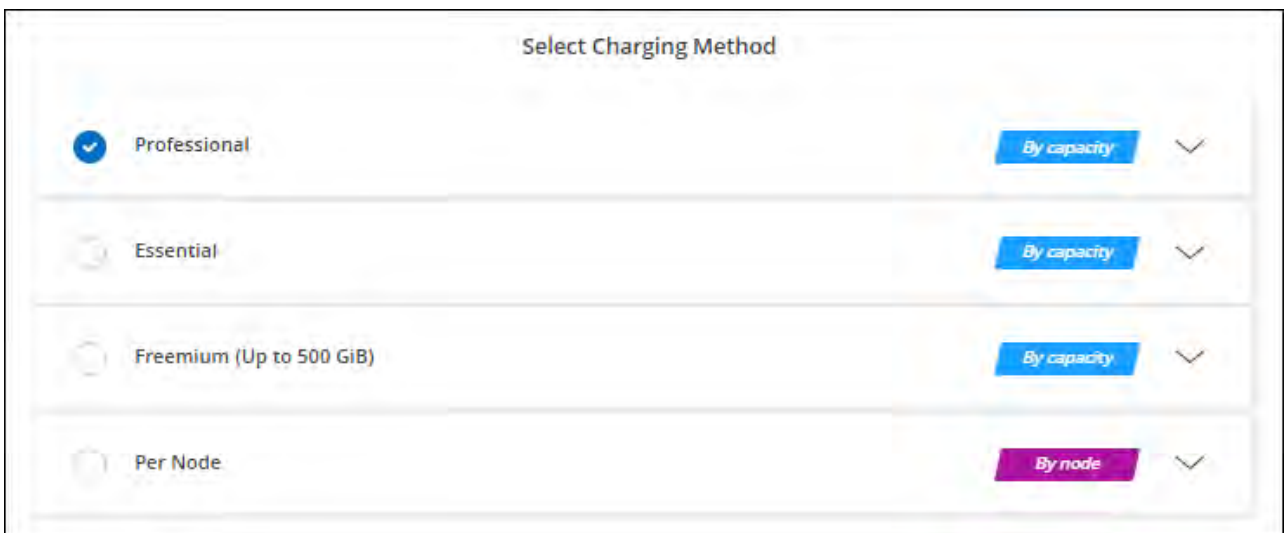
environment for charging. You can use that same subscription for additional working environments.

### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the pay-as-you-go offering in the Azure Marketplace.



- b. After you return to BlueXP, select a capacity-based package when you reach the charging methods page.



[View step-by-step instructions to launch Cloud Volumes ONTAP in Azure.](#)



You can manage the Azure Marketplace subscriptions associated with your Azure accounts from the Settings > Credentials page. [Learn how to manage your Azure accounts and subscriptions](#)

## Annual contract

Pay for Cloud Volumes ONTAP annually by purchasing an annual contract.

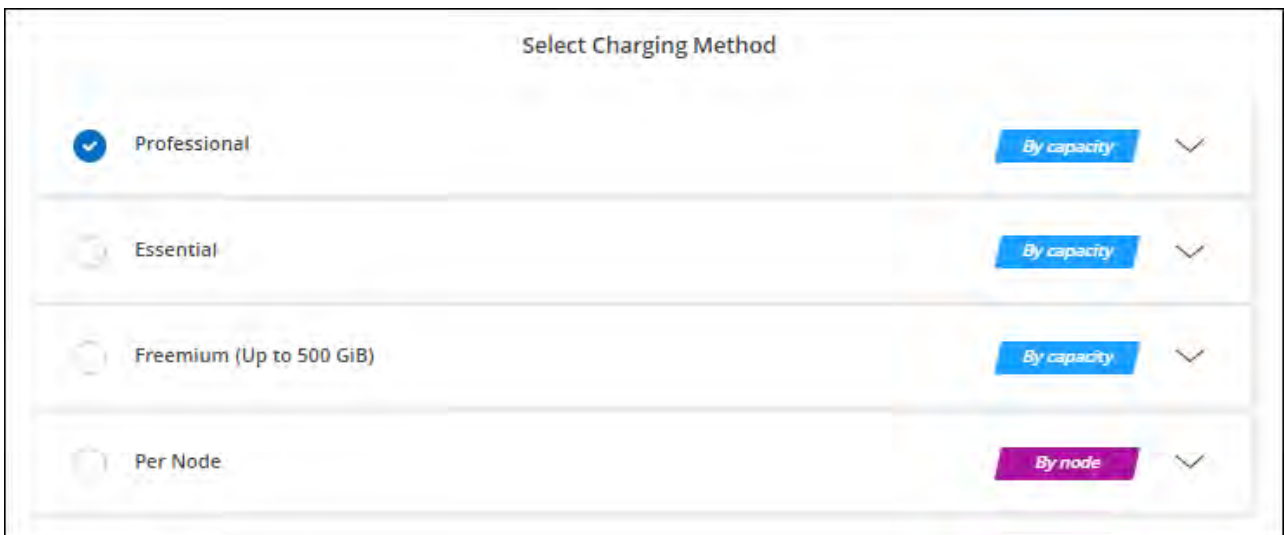
### Steps

1. Contact your NetApp sales representative to purchase an annual contract.

The contract is available as a *private* offer in the Azure Marketplace.

After NetApp shares the private offer with you, you can select the annual plan when you subscribe from the Azure Marketplace during working environment creation.

2. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription > Continue**.
  - b. In the Azure portal, select the annual plan that was shared with your Azure account and then click **Subscribe**.
  - c. After you return to BlueXP, select a capacity-based package when you reach the charging methods page.



[View step-by-step instructions to launch Cloud Volumes ONTAP in Azure.](#)

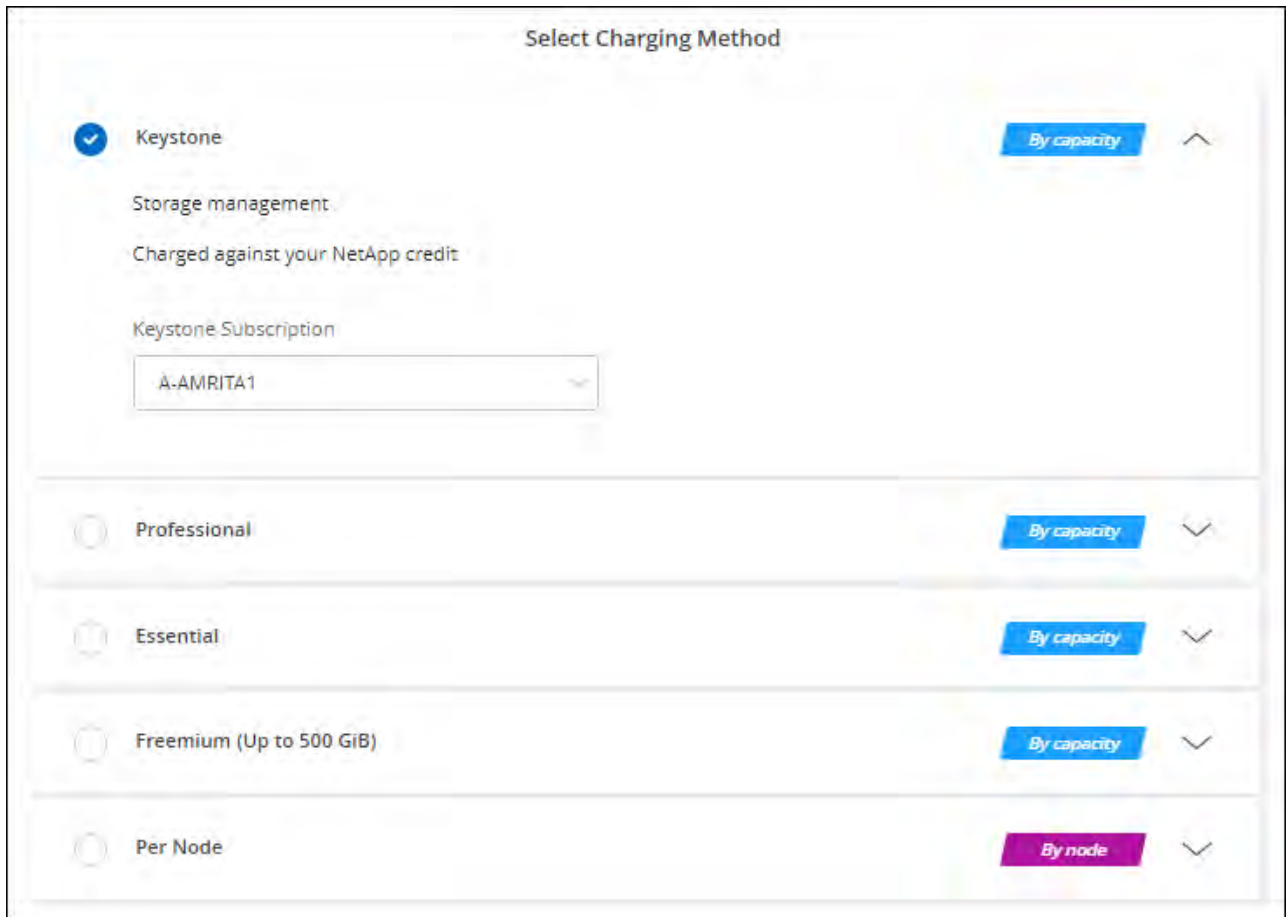
## Keystone Subscription

A Keystone Subscription is a pay-as-you-grow subscription-based service. [Learn more about NetApp Keystone Subscriptions.](#)

### Steps

1. If you don't have a subscription yet, [contact NetApp](#)
2. [Contact NetApp](#) to authorize your BlueXP user account with one or more Keystone Subscriptions.

3. After NetApp authorizes your account, [link your subscriptions for use with Cloud Volumes ONTAP](#).
4. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. Select the Keystone Subscription charging method when prompted to choose a charging method.



[View step-by-step instructions to launch Cloud Volumes ONTAP in Azure.](#)

### Enable high-availability mode in Azure

Microsoft Azure's high-availability (HA) mode should be enabled to reduce unplanned failover times and to enable NFSv4 support for Cloud Volumes ONTAP. In this mode, your Cloud Volumes ONTAP HA nodes can achieve a low (60 seconds) recovery time objective (RTO) during unplanned failovers on CIFS and NFSv4 clients.

Beginning with Cloud Volumes ONTAP 9.10.1, we reduced the unplanned failover time for Cloud Volumes ONTAP HA pairs running in Microsoft Azure and added support for NFSv4. To make these enhancements available to Cloud Volumes ONTAP, you need to enable the high-availability feature on your Azure subscription.

BlueXP will prompt you with these details in an Action Required message when the feature needs to be enabled on an Azure subscription.

Note the following:

- There are no problems with the high availability of your Cloud Volumes ONTAP HA pair. This Azure feature works in concert with ONTAP to reduce the client observed application outage time for NFS protocols that

result from unplanned failover events.

- Enabling this feature is non-disruptive to Cloud Volumes ONTAP HA pairs.
- Enabling this feature on your Azure subscription doesn't cause issues to other VMs.
- Cloud Volumes ONTAP uses an internal Azure Load Balancer during failovers of cluster and SVM management LIFs on CIFS and NFS clients.
- When the HA mode is enabled, BlueXP scans the system every 12 hours to update the internal Azure Load Balancer rules.

An Azure user who has "Owner" privileges can enable the feature from the Azure CLI.

### Steps

1. [Access the Azure Cloud Shell from the Azure Portal](#)
2. Register the high-availability mode feature:

```
az account set -s AZURE_SUBSCRIPTION_NAME_OR_ID
az feature register --name EnableHighAvailabilityMode --namespace
Microsoft.Network
az provider register -n Microsoft.Network
```

3. Optionally verify that the feature is now registered:

```
az feature show --name EnableHighAvailabilityMode --namespace
Microsoft.Network
```

The Azure CLI should return a result similar to the following:

```
{
  "id": "/subscriptions/xxxxxxxx-xxxx-xxxx-xxxx-
xxxxxxxxxxxx/providers/Microsoft.Features/providers/Microsoft.Network/fe
atures/EnableHighAvailabilityMode",
  "name": "Microsoft.Network/EnableHighAvailabilityMode",
  "properties": {
    "state": "Registered"
  },
  "type": "Microsoft.Features/providers/features"
}
```

### Enable VMOrchestratorZonalMultiFD for Cloud Volumes ONTAP in Azure

For deploying VM instances in locally-redundant storage (LRS) single availability zones (AZ), you should activate the Microsoft `Microsoft.Compute/VMOrchestratorZonalMultiFD` feature for your subscriptions. In a high-availability (HA) mode, this feature facilitates deploying nodes in

separate fault domains in the same availability zone.

Unless you activate this feature, zonal deployment doesn't occur, and the previous LRS non-zonal deployment becomes effective.

For information about VM deployment in single availability zone, refer to [High-availability pairs in Azure](#).

Perform these steps as a user with "Owner" privileges:

### Steps

1. Access Azure Cloud Shell from the Azure portal. For information, refer to [Microsoft Azure documentation: Get started with Azure Cloud Shell](#).
2. Register for the `Microsoft.Compute/VMOrchestratorZonalMultiFD` feature by running this command:

```
az account set -s <Azure_subscription_name_or_ID>
az feature register --name VMOrchestratorZonalMultiFD --namespace Microsoft.Compute
```

3. Verify the registration status and output sample:

```
az feature show -n VMOrchestratorZonalMultiFD --namespace Microsoft.Compute
{
  "id": "/subscriptions/
  </ID>/providers/Microsoft.Features/providers/Microsoft.Compute/features/VMOrchestratorZonalMultiF
  D",
  "name": "Microsoft.Compute/VMOrchestratorZonalMultiFD",
  "properties": {
    "state": "Registered"
  },
  "type": "Microsoft.Features/providers/features"
}
```

## Launch Cloud Volumes ONTAP in Azure

You can launch a single node system or an HA pair in Azure by creating a Cloud Volumes ONTAP working environment in BlueXP.

### Before you begin

You need the following to create a working environment.

- A Connector that's up and running.
  - You should have a [Connector that is associated with your project or workspace](#).
  - [You should be prepared to leave the Connector running at all times](#).
- An understanding of the configuration that you want to use.

You should have chose a configuration and obtained Azure networking information from your administrator. For information, refer to [Planning your Cloud Volumes ONTAP configuration](#).

- An understanding of what's required to set up licensing for Cloud Volumes ONTAP.

[Learn how to set up licensing.](#)

### About this task

When BlueXP creates a Cloud Volumes ONTAP system in Azure, it creates several Azure objects, such as a resource group, network interfaces, and storage accounts. You can review a summary of the resources at the end of the wizard.

#### Potential for Data Loss

The best practice is to use a new, dedicated resource group for each Cloud Volumes ONTAP system.



Deploying Cloud Volumes ONTAP in an existing, shared resource group is not recommended due to the risk of data loss. While BlueXP can remove Cloud Volumes ONTAP resources from a shared resource group in case of deployment failure or deletion, an Azure user might accidentally delete Cloud Volumes ONTAP resources from a shared resource group.

### Launch a single-node Cloud Volumes ONTAP system in Azure

If you want to launch a single-node Cloud Volumes ONTAP system in Azure, you need to create a single node working environment in BlueXP.

#### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the prompts.
3. **Choose a Location:** Select **Microsoft Azure** and **Cloud Volumes ONTAP Single Node**.
4. If you're prompted, [create a Connector](#).
5. **Details and Credentials:** Optionally change the Azure credentials and subscription, specify a cluster name, add tags if needed, and then specify credentials.

The following table describes fields for which you might need guidance:

Field	Description
Working Environment Name	BlueXP uses the working environment name to name both the Cloud Volumes ONTAP system and the Azure virtual machine. It also uses the name as the prefix for the predefined security group, if you select that option.
Resource Group Tags	<p>Tags are metadata for your Azure resources. When you enter tags in this field, BlueXP adds them to the resource group associated with the Cloud Volumes ONTAP system.</p> <p>You can add up to four tags from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four tags when creating a working environment.</p> <p>For information about tags, refer to the <a href="#">Microsoft Azure Documentation: Using tags to organize your Azure resources</a>.</p>

Field	Description
User name and password	These are the credentials for the Cloud Volumes ONTAP cluster administrator account. You can use these credentials to connect to Cloud Volumes ONTAP through ONTAP System Manager or the ONTAP CLI. Keep the default <i>admin</i> user name or change it to a custom user name.
Edit Credentials	You can choose different Azure credentials and a different Azure subscription to use with this Cloud Volumes ONTAP system. You need to associate an Azure Marketplace subscription with the selected Azure subscription in order to deploy a pay-as-you-go Cloud Volumes ONTAP system. <a href="#">Learn how to add credentials.</a>

The following video shows how to associate a Marketplace subscription to an Azure subscription:

[Subscribe to BlueXP from the Azure Marketplace](#)

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

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



If you would like to utilize WORM and data tiering, you must disable BlueXP backup and recovery and deploy a Cloud Volumes ONTAP working environment with version 9.8 or above.

7. **Location:** Select a region, availability zone, VNet, and subnet, and then select the checkbox to confirm network connectivity between the Connector and the target location.




For China regions, single node deployments are supported only in Cloud Volumes ONTAP 9.12.1 GA and 9.13.0 GA. You can upgrade these versions to later patches and releases of Cloud Volumes ONTAP. If you want to deploy later Cloud Volumes ONTAP versions in China regions, contact NetApp Support. Only licenses purchased directly from NetApp are supported in China regions, marketplace subscriptions are not available.

8. **Connectivity:** Choose a new or existing resource group and then choose whether to use the predefined security group or to use your own.

The following table describes fields for which you might need guidance:



Field	Description
Resource Group	<p>Create a new resource group for Cloud Volumes ONTAP or use an existing resource group. The best practice is to use a new, dedicated resource group for Cloud Volumes ONTAP. While it is possible to deploy Cloud Volumes ONTAP in an existing, shared resource group, it's not recommended due to the risk of data loss. See the warning above for more details.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>If the Azure account that you're using has the <a href="#">required permissions</a>, BlueXP removes Cloud Volumes ONTAP resources from a resource group, in case of deployment failure or deletion.</p> </div>
Generated security group	<p>If you let BlueXP generate the security group for you, you need to choose how you'll allow traffic:</p> <ul style="list-style-type: none"> <li>• If you choose <b>Selected VNet only</b>, the source for inbound traffic is the subnet range of the selected VNet and the subnet range of the VNet where the Connector resides. This is the recommended option.</li> <li>• If you choose <b>All VNets</b>, the source for inbound traffic is the 0.0.0.0/0 IP range.</li> </ul>
Use existing	<p>If you choose an existing security group, then it must meet Cloud Volumes ONTAP requirements. <a href="#">View the default security group</a>.</p>

9. **Charging Methods and NSS Account:** Specify which charging option would you like to use with this system, and then specify a NetApp Support Site account.
  - [Learn about licensing options for Cloud Volumes ONTAP](#).
  - [Learn how to set up licensing](#).
10. **Preconfigured Packages:** Select one of the packages to quickly deploy a Cloud Volumes ONTAP system, or click **Create my own configuration**.

If you choose one of the packages, you only need to specify a volume and then review and approve the configuration.

11. **Licensing:** Change the Cloud Volumes ONTAP version if required, and select a virtual machine type.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then BlueXP updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.13.1 and 9.13.1 P4 is available. The update does not occur from one release to another—for example, from 9.13 to 9.14.

12. **Subscribe from the Azure Marketplace:** You see this page if BlueXP could not enable programmatic deployments of Cloud Volumes ONTAP. Follow the steps listed on the screen. refer to [Programmatic deployment of Marketplace products](#) for more information.
13. **Underlying Storage Resources:** Choose settings for the initial aggregate: a disk type, a size for each disk, and whether data tiering to Blob storage should be enabled.

Note the following:

- If the public access to your storage account is disabled within the VNet, you cannot enable data tiering in your Cloud Volumes ONTAP system. For information, refer to [Security group rules](#).
- The disk type is for the initial volume. You can choose a different disk type for subsequent volumes.
- The disk size is for all disks in the initial aggregate and for any additional aggregates that BlueXP creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.

For help choosing a disk type and size, refer to [Sizing your system in Azure](#).

- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates.

[Learn more about data tiering](#).

#### 14. **Write Speed & WORM:**

- Choose **Normal** or **High** write speed, if desired.

[Learn more about write speed](#).

- Activate write once, read many (WORM) storage, if desired.

This option is only available for certain VM types. To find out which VM types are supported, refer to [Supported configurations by license for HA pairs](#).

WORM can't be enabled if data tiering was enabled for Cloud Volumes ONTAP versions 9.7 and below. Reverting or downgrading to Cloud Volumes ONTAP 9.8 is blocked after enabling WORM and tiering.

[Learn more about WORM storage](#).

- If you activate WORM storage, select the retention period.

#### 15. **Create Volume:** Enter details for the new volume or click **Skip**.

[Learn about supported client protocols and versions](#).

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

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, BlueXP enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.

Field	Description
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, BlueXP automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

### Volume Details, Protection & Protocol

#### Details & Protection

Volume Name:  Size (GB):

Snapshot Policy:

Default Policy

#### Protocol

NFS     CIFS     iSCSI

Share name:  Permissions:

Users / Groups:

Valid users and groups separated by a semicolon

16. **CIFS Setup:** If you chose the CIFS protocol, set up a CIFS server.

Field	Description
DNS Primary and Secondary IP Address	<p>The IP addresses of the DNS servers that provide name resolution for the CIFS server.</p> <p>The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.</p>

Field	Description
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	<p>The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.</p> <p>To configure Azure AD Domain Services as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=AADDCC Computers</b> or <b>OU=AADDCC Users</b> in this field.</p> <p><a href="#">Azure Documentation: Create an Organizational Unit (OU) in an Azure AD Domain Services managed domain</a></p>
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	<p>Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. Refer to the <a href="#">BlueXP automation docs</a> for details.</p> <p>Note that you can configure an NTP server only when creating a CIFS server. It's not configurable after you create the CIFS server.</p>

17. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and change the volume tiering policy, if needed.

For more information, refer to [Understanding volume usage profiles](#) and [Data tiering overview](#).

18. **Review & Approve:** Review and confirm your selections.
- Review details about the configuration.
  - Click **More information** to review details about support and the Azure resources that BlueXP will purchase.
  - Select the **I understand...** check boxes.
  - Click **Go**.

### Result

BlueXP deploys the Cloud Volumes ONTAP system. You can track the progress in the timeline.

If you experience any issues deploying the Cloud Volumes ONTAP system, review the failure message. You can also select the working environment and click **Re-create environment**.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

### After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.

- If you want to apply quotas to volumes, use ONTAP System Manager or the ONTAP CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

### Launch a Cloud Volumes ONTAP HA pair in Azure

If you want to launch a Cloud Volumes ONTAP HA pair in Azure, you need to create an HA working environment in BlueXP.

#### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the prompts.
3. If you're prompted, [create a Connector](#).
4. **Details and Credentials:** Optionally change the Azure credentials and subscription, specify a cluster name, add tags if needed, and then specify credentials.

The following table describes fields for which you might need guidance:

Field	Description
Working Environment Name	BlueXP uses the working environment name to name both the Cloud Volumes ONTAP system and the Azure virtual machine. It also uses the name as the prefix for the predefined security group, if you select that option.
Resource Group Tags	<p>Tags are metadata for your Azure resources. When you enter tags in this field, BlueXP adds them to the resource group associated with the Cloud Volumes ONTAP system.</p> <p>You can add up to four tags from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four tags when creating a working environment.</p> <p>For information about tags, refer to the <a href="#">Microsoft Azure Documentation: Using tags to organize your Azure resources</a>.</p>
User name and password	These are the credentials for the Cloud Volumes ONTAP cluster administrator account. You can use these credentials to connect to Cloud Volumes ONTAP through ONTAP System Manager or the ONTAP CLI. Keep the default <i>admin</i> user name or change it to a custom user name.
Edit Credentials	You can choose different Azure credentials and a different Azure subscription to use with this Cloud Volumes ONTAP system. You need to associate an Azure Marketplace subscription with the selected Azure subscription in order to deploy a pay-as-you-go Cloud Volumes ONTAP system. <a href="#">Learn how to add credentials</a> .

The following video shows how to associate a Marketplace subscription to an Azure subscription:

[Subscribe to BlueXP from the Azure Marketplace](#)

5. **Services:** Enable or disable the individual services based on whether you want to use them with Cloud Volumes ONTAP.
  - [Learn more about BlueXP classification](#)

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



If you would like to utilize WORM and data tiering, you must disable BlueXP backup and recovery and deploy a Cloud Volumes ONTAP working environment with version 9.8 or above.

## 6. HA Deployment Models:

### a. Select **Single Availability Zone** or **Multiple Availability Zone**.

- For single availability zones, select an Azure region, availability zone, VNet, and subnet.


Beginning with Cloud Volumes ONTAP 9.15.1, you can deploy virtual machine (VM) instances in HA mode in single availability zones (AZs) in Azure. You need to select a zone and a region that support this deployment. If the zone or the region does not support zonal deployment, then the previous non-zonal deployment mode for LRS is followed. For understanding the supported configurations for shared managed disks, refer to [HA single availability zone configuration with shared managed disks](#).

- For multiple availability zones, select a region, VNet, subnet, zone for node 1, and zone for node 2.

### b. Select the **I have verified network connectivity...** check box.

## 7. **Connectivity:** Choose a new or existing resource group and then choose whether to use the predefined security group or to use your own.

The following table describes fields for which you might need guidance:

Field	Description
Resource Group	<p>Create a new resource group for Cloud Volumes ONTAP or use an existing resource group. The best practice is to use a new, dedicated resource group for Cloud Volumes ONTAP. While it is possible to deploy Cloud Volumes ONTAP in an existing, shared resource group, it's not recommended due to the risk of data loss. See the warning above for more details.</p> <p>You must use a dedicated resource group for each Cloud Volumes ONTAP HA pair that you deploy in Azure. Only one HA pair is supported in a resource group. BlueXP experiences connection issues if you try to deploy a second Cloud Volumes ONTAP HA pair in an Azure resource group.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>If the Azure account that you're using has the <a href="#">required permissions</a>, BlueXP removes Cloud Volumes ONTAP resources from a resource group, in case of deployment failure or deletion.</p> </div>
Generated security group	<p>If you let BlueXP generate the security group for you, you need to choose how you'll allow traffic:</p> <ul style="list-style-type: none"> <li>• If you choose <b>Selected VNet only</b>, the source for inbound traffic is the subnet range of the selected VNet and the subnet range of the VNet where the Connector resides. This is the recommended option.</li> <li>• If you choose <b>All VNets</b>, the source for inbound traffic is the 0.0.0.0/0 IP range.</li> </ul>

Field	Description
Use existing	If you choose an existing security group, then it must meet Cloud Volumes ONTAP requirements. <a href="#">View the default security group.</a>

8. **Charging Methods and NSS Account:** Specify which charging option would you like to use with this system, and then specify a NetApp Support Site account.
  - [Learn about licensing options for Cloud Volumes ONTAP.](#)
  - [Learn how to set up licensing.](#)
9. **Preconfigured Packages:** Select one of the packages to quickly deploy a Cloud Volumes ONTAP system, or click **Change configuration**.

If you choose one of the packages, you only need to specify a volume and then review and approve the configuration.

10. **Licensing:** Change the Cloud Volumes ONTAP version as needed and select a virtual machine type.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then BlueXP updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.13.1 and 9.13.1 P4 is available. The update does not occur from one release to another—for example, from 9.13 to 9.14.

11. **Subscribe from the Azure Marketplace:** Follow the steps if BlueXP could not enable programmatic deployments of Cloud Volumes ONTAP.
12. **Underlying Storage Resources:** Choose settings for the initial aggregate: a disk type, a size for each disk, and whether data tiering to Blob storage should be enabled.

Note the following:

- The disk size is for all disks in the initial aggregate and for any additional aggregates that BlueXP creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.

For help choosing a disk size, refer to [Size your system in Azure](#).

- If the public access to your storage account is disabled within the VNet, you cannot enable data tiering in your Cloud Volumes ONTAP system. For information, refer to [Security group rules](#).
- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates.

[Learn more about data tiering.](#)

- Starting with Cloud Volumes ONTAP 9.15.0P1, Azure page blobs are no longer supported for new high-availability pair deployments. If you currently use Azure page blobs in existing high-availability pair deployments, you can migrate to newer VM instance types in the Edsv4-series VMs and Edsv5-series VMs.

[Learn more about supported configurations in Azure.](#)

13. **Write Speed & WORM:**

- a. Choose **Normal** or **High** write speed, if desired.

[Learn more about write speed.](#)

- b. Activate write once, read many (WORM) storage, if desired.

This option is only available for certain VM types. To find out which VM types are supported, refer to [Supported configurations by license for HA pairs](#).

WORM can't be enabled if data tiering was enabled for Cloud Volumes ONTAP versions 9.7 and below. Reverting or downgrading to Cloud Volumes ONTAP 9.8 is blocked after enabling WORM and tiering.

[Learn more about WORM storage.](#)

- c. If you activate WORM storage, select the retention period.

14. **Secure Communication to Storage & WORM:** Choose whether to enable an HTTPS connection to Azure storage accounts, and activate write once, read many (WORM) storage, if desired.

The HTTPS connection is from a Cloud Volumes ONTAP 9.7 HA pair to Azure page blob storage accounts. Note that enabling this option can impact write performance. You can't change the setting after you create the working environment.

[Learn more about WORM storage.](#)

WORM can't be enabled if data tiering was enabled.

[Learn more about WORM storage.](#)

15. **Create Volume:** Enter details for the new volume or click **Skip**.

[Learn about supported client protocols and versions.](#)

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

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, BlueXP enters a value that provides access to all instances in the subnet.
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.



Field	Description
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, BlueXP automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

### Volume Details, Protection & Protocol

#### Details & Protection

Volume Name:  Size (GB):

Snapshot Policy:

Default Policy

#### Protocol

NFS     CIFS     iSCSI

Share name:  Permissions:

Users / Groups:

Valid users and groups separated by a semicolon

16. **CIFS Setup:** If you chose the CIFS protocol, set up a CIFS server.

Field	Description
DNS Primary and Secondary IP Address	<p>The IP addresses of the DNS servers that provide name resolution for the CIFS server.</p> <p>The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.</p>
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.

Field	Description
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	<p>The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.</p> <p>To configure Azure AD Domain Services as the AD server for Cloud Volumes ONTAP, you should enter <b>OU=AADDC Computers</b> or <b>OU=AADDC Users</b> in this field.</p> <p><a href="#">Azure Documentation: Create an Organizational Unit (OU) in an Azure AD Domain Services managed domain</a></p>
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	<p>Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. Refer to the <a href="#">BlueXP automation docs</a> for details.</p> <p>Note that you can configure an NTP server only when creating a CIFS server. It's not configurable after you create the CIFS server.</p>

17. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and change the volume tiering policy, if needed.

For more information, refer to [Choose a volume usage profile](#), [Data tiering overview](#), and [KB: What Inline Storage Efficiency features are supported with CVO?](#)

18. **Review & Approve:** Review and confirm your selections.
- Review details about the configuration.
  - Click **More information** to review details about support and the Azure resources that BlueXP will purchase.
  - Select the **I understand...** check boxes.
  - Click **Go**.

## Result

BlueXP deploys the Cloud Volumes ONTAP system. You can track the progress in the timeline.

If you experience any issues deploying the Cloud Volumes ONTAP system, review the failure message. You can also select the working environment and click **Re-create environment**.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

## After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.
- If you want to apply quotas to volumes, use ONTAP System Manager or the ONTAP CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

## Verify Azure platform image

### Azure marketplace image verification for Cloud Volumes ONTAP


Azure image verification complies with enhanced NetApp security requirements. Verifying an image file is a straightforward process. However, the Azure image signature verification requires specific considerations for the Azure VHD image file because it is altered in the Azure marketplace.



Azure image verification is supported on Cloud Volumes ONTAP 9.15.0 and later.

### Azure's alteration of published VHD files

The 1 MB (1048576 bytes) at the beginning and 512 bytes at the end of the VHD file is modified by Azure. NetApp signs the remaining VHD file.



In the example, the VHD file is of 10GB. The portion that NetApp signed is marked in green (10 GB - 1 MB - 512 bytes).

### Related links

- [Page Fault Blog: How to sign and verify using OpenSSL](#)
- [Use Azure Marketplace image to create VM image for your Azure Stack Edge Pro GPU | Microsoft Learn](#)
- [Export/Copy a managed disk to a storage account using the Azure CLI | Microsoft Learn](#)
- [Azure Cloud Shell Quickstart - Bash | Microsoft Learn](#)
- [How to install the Azure CLI | Microsoft Learn](#)
- [az storage blob copy | Microsoft Learn](#)
- [Sign in with Azure CLI — Login and Authentication | Microsoft Learn](#)

### Download the Azure image file for Cloud Volumes ONTAP

You can download the Azure image file from the [NetApp Support Site](#).

The *tar.gz* file contains the files required for image signature verification. Along with the *tar.gz* file, you should also download the *checksum* file for the image. The checksum file contains the `md5` and `sha256` checksums of the *tar.gz* file.

### Steps

1. Go to the [Cloud Volumes ONTAP product page on the NetApp Support Site](#) and download the required software version from the **Downloads** section.
2. On the Cloud Volumes ONTAP download page, click the downloadable file for the Azure image and download the *tar.gz* file.

# Cloud Volumes ONTAP 9.15.0P1

Date Posted : 17-May-2024

3. On Linux, run `md5sum AZURE-<version>_PKG.TAR.GZ`.

On macOS, run `sha256sum AZURE-<version>_PKG.TAR.GZ`.

4. Verify that the `md5sum` and `sha256sum` values match those in the downloaded Azure image.

5. On Linux and macOS, extract the `tar.gz` file using the `tar -xzf` command.

The extracted `tar.gz` file contains the digest (`.sig`) file, public key certificate (`.pem`) file, and chain certificate (`.pem`) file.

### Example output after extracting the `tar.gz` file:

```
$ ls cert/ -l
-rw-r----- 1 netapp netapp 384 May 13 13:00 9.15.0P1_azure_digest.sig
-rw-r----- 1 netapp netapp 2365 May 13 13:00 Certificate-
9.15.0P1_azure.pem
-rw-r----- 1 netapp netapp 8537 May 13 13:00 Certificate-Chain-
9.15.0P1_azure.pem
-rw-r----- 1 netapp netapp 8537 May 13 13:00 version_readme
```

### Export VHD images for Cloud Volumes ONTAP from the Azure marketplace

Once the VHD image is published to Azure cloud, it is no longer managed by NetApp. Instead, the published image is placed on the Azure marketplace. When the image is staged and published on the Azure marketplace, Azure modifies 1 MB at the beginning and 512 bytes at the end of the VHD. To verify the signature of the VHD file, you need to export the VHD image modified by Azure from the Azure marketplace.

### Before you begin

Ensure that the Azure CLI is installed on your system, or the Azure Cloud Shell is available through the Azure portal. For more information about how to install the Azure CLI, refer to [Azure documentation: How to install the Azure CLI](#).

## Steps

1. Map the Cloud Volumes ONTAP version on your system to the Azure marketplace image version using the contents of the `version_readme` file. The Cloud Volumes ONTAP version is represented by `buildname` and the Azure marketplace image version is represented by `version` in the version mappings.

In the following example, the Cloud Volumes ONTAP version 9.15.0P1 is mapped to the Azure marketplace image version 9150.01000024.05090105. This Azure marketplace image version is later used to set the image URN.

```
[
  "buildname": "9.15.0P1",
  "publisher": "netapp",
  "version": "9150.01000024.05090105"
]
```

2. Identify the region where you want to create the VMs. The region name is used as the value for the `locName` variable when setting the URN of the marketplace image. To list the available regions, run this command:

```
az account list-locations -o table
```

In this table, the region name appears in the `Name` field.

```
$ az account list-locations -o table
DisplayName          Name                RegionalDisplayName
-----
East US              eastus              (US) East US
East US 2            eastus2             (US) East US 2
South Central US    southcentralus     (US) South Central US
...
```

3. Review the SKU names for the corresponding Cloud Volumes ONTAP versions and VM deployment types in the table below. The SKU name is used as the value for the `skuName` variable when setting the URN of the marketplace image.

For example, all single node deployments with Cloud Volumes ONTAP 9.15.0 should use `ontap_cloud_byol` as the SKU name.

Cloud Volumes ONTAP version	VM deployment method	SKU name
9.16.1 and later	Through the marketplace direct deployment method	ontap_cloud_direct
9.16.1 and later	All deployments through BlueXP	ontap_cloud

9.15.1	All deployments through BlueXP	ontap_cloud
9.15.0	Single node	ontap_cloud_byol
9.15.0	High availability	ontap_cloud_byol_ha

4. After mapping the ONTAP version and Azure marketplace image, export the VHD file from the Azure marketplace using the Azure Cloud Shell or Azure CLI.

### Export VHD file using the Azure Cloud Shell on Linux

From the Azure Cloud Shell, export the marketplace image to the VHD file (for example, *9150.01000024.05090105.vhd*), and download it to your local Linux system. Perform these steps to get the VHD image from the Azure marketplace.

#### Steps

1. Set the URN and other parameters of the marketplace image. The URN format is `<publisher>:<offer>:<sku>:<version>`. Optionally, you can list NetApp marketplace images to confirm the correct image version.

```
PS /home/user1> $urn="netapp:netapp-ontap-
cloud:ontap_cloud_byol:9150.01000024.05090105"
PS /home/user1> $locName="eastus2"
PS /home/user1> $pubName="netapp"
PS /home/user1> $offerName="netapp-ontap-cloud"
PS /home/user1> $skuName="ontap_cloud_byol"
PS /home/user1> Get-AzVMImage -Location $locName -PublisherName $pubName
-Offer $offerName -Sku $skuName |select version
...
141.20231128
9.141.20240131
9.150.20240213
9150.01000024.05090105
...
```

2. Create a new managed disk from the marketplace image with the matching image version:

```
PS /home/user1> $diskName = "9150.01000024.05090105-managed-disk"
PS /home/user1> $diskRG = "fnf1"
PS /home/user1> az disk create -g $diskRG -n $diskName --image-reference
$urn
PS /home/user1> $sas = az disk grant-access --duration-in-seconds 3600
--access-level Read --name $diskName --resource-group $diskRG
PS /home/user1> $diskAccessSAS = ($sas | ConvertFrom-Json)[0].accessSas
```

3. Export the VHD file from the managed disk to Azure Storage. Create a container with the appropriate access level. In this example, we've used a container named `vm-images` with `Container` access level. Get the storage account access key from the Azure portal: **Storage Accounts > *examplesname* >**

## Access Key > key1 > key > Show > <copy>

```
PS /home/user1> $storageAccountName = "examplesaname"
PS /home/user1> $containerName = "vm-images"
PS /home/user1> $storageAccountKey = "<replace with the above access
key>"
PS /home/user1> $destBlobName = "9150.01000024.05090105.vhd"
PS /home/user1> $destContext = New-AzureStorageContext
-StorageAccountName $storageAccountName -StorageAccountKey
$storageAccountKey
PS /home/user1> Start-AzureStorageBlobCopy -AbsoluteUri $diskAccessSAS
-DestContainer $containerName -DestContext $destContext -DestBlob
$destBlobName
PS /home/user1> Get-AzureStorageBlobCopyState -Container $containerName
-Context $destContext -Blob $destBlobName
```

4. Download the generated image to your Linux system. Use the `wget` command to download the VHD file:

```
wget <URL of filename/Containers/vm-images/9150.01000024.05090105.vhd>
```

The URL follows a standard format. For automation, you can derive the URL string as shown below. Alternatively, you can use the Azure CLI `az` command to get the URL.

Example URL:

<https://examplesaname.blob.core.windows.net/vm-images/9150.01000024.05090105.vhd>

5. Clean up the managed disk

```
PS /home/user1> Revoke-AzDiskAccess -ResourceGroupName $diskRG -DiskName
$diskName
PS /home/user1> Remove-AzDisk -ResourceGroupName $diskRG -DiskName
$diskName
```

## Export VHD file using the Azure CLI on Linux

Export the marketplace image to a VHD file using the Azure CLI from a local Linux system.

### Steps

1. Log in to the Azure CLI and list marketplace images:

```
% az login --use-device-code
```

2. To sign in, use a web browser to open the page <https://microsoft.com/devicelogin> and enter the authentication code.

```
% az vm image list --all --publisher netapp --offer netapp-ontap-cloud
--sku ontap_cloud_byol
...
{
  "architecture": "x64",
  "offer": "netapp-ontap-cloud",
  "publisher": "netapp",
  "sku": "ontap_cloud_byol",
  "urn": "netapp:netapp-ontap-
cloud:ontap_cloud_byol:9150.01000024.05090105",
  "version": "9150.01000024.05090105"
},
...
```

3. Create a new managed disk from the marketplace image with the matching image version.

```
% export urn="netapp:netapp-ontap-
cloud:ontap_cloud_byol:9150.01000024.05090105"
% export diskName="9150.01000024.05090105-managed-disk"
% export diskRG="new_rg_your_rg"
% az disk create -g $diskRG -n $diskName --image-reference $urn
% az disk grant-access --duration-in-seconds 3600 --access-level Read
--name $diskName --resource-group $diskRG
{
  "accessSas": "https://md-
xxxxxx.blob.core.windows.net/xxxxxxx/abcd?sv=2018-03-
28&sr=b&si=xxxxxxx-xxxx-xxxx-xxxx-xxxxxxx&sigxxxxxxxxxxxxxxxxxxxxxxxxxxxx"
}
% export diskAccessSAS="https://md-
xxxxxx.blob.core.windows.net/xxxxxxx/abcd?sv=2018-03-
28&sr=b&si=xxxxxxx-xxxx-xx-xx-xx&sigxxxxxxxxxxxxxxxxxxxxxxxxxxxx"
```

To automate the process, the SAS needs to be extracted from the standard output. Refer to the appropriate documents for guidance.

4. Export the VHD file from the managed disk.

- a. Create a container with the appropriate access level. In this example, a container named `vm-images` with `Container` access level is used.
- b. Get the storage account access key from the Azure portal: **Storage Accounts > *examplesname* > Access Key > *key1* > *key* > Show > <copy>**

You can also use the `az` command for this step.



```

% export storageAccountName="examplesaname"
% export containerName="vm-images"
% export storageAccountKey="xxxxxxxxxxx"
% export destBlobName="9150.01000024.05090105.vhd"

% az storage blob copy start --source-uri $diskAccessSAS
--destination-container $containerName --account-name
$storageAccountName --account-key $storageAccountKey --destination
-blob $destBlobName

{
  "client_request_id": "xxxx-xxxx-xxxx-xxxx-xxxx",
  "copy_id": "xxxx-xxxx-xxxx-xxxx-xxxx",
  "copy_status": "pending",
  "date": "2022-11-02T22:02:38+00:00",
  "etag": "\"0xxxxxxxxxxxxxxxxxxxx\"",
  "last_modified": "2022-11-02T22:02:39+00:00",
  "request_id": "xxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "version": "2020-06-12",
  "version_id": null
}

```

##### 5. Check the status of the blob copy.

```

% az storage blob show --name $destBlobName --container-name
$containerName --account-name $storageAccountName

....
  "copy": {
    "completionTime": null,
    "destinationSnapshot": null,
    "id": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxx",
    "incrementalCopy": null,
    "progress": "10737418752/10737418752",
    "source": "https://md-
xxxxxx.blob.core.windows.net/xxxxx/abcd?sv=2018-03-28&sr=b&si=xxxxxxxx-
xxxx-xxxx-xxxx-xxxxxxxxxxxx",
    "status": "success",
    "statusDescription": null
  },
....

```

##### 6. Download the generated image to your Linux server.

```
wget <URL of file examplesname/Containers/vm-  
images/9150.01000024.05090105.vhd>
```

The URL follows a standard format. For automation, you can derive the URL string as shown below. Alternatively, you can use the Azure CLI `az` command to get the URL.

Example URL:

<https://examplesname.blob.core.windows.net/vm-images/9150.01000024.05090105.vhd>

## 7. Clean up the managed disk

```
az disk revoke-access --name $diskName --resource-group $diskRG  
az disk delete --name $diskName --resource-group $diskRG --yes
```

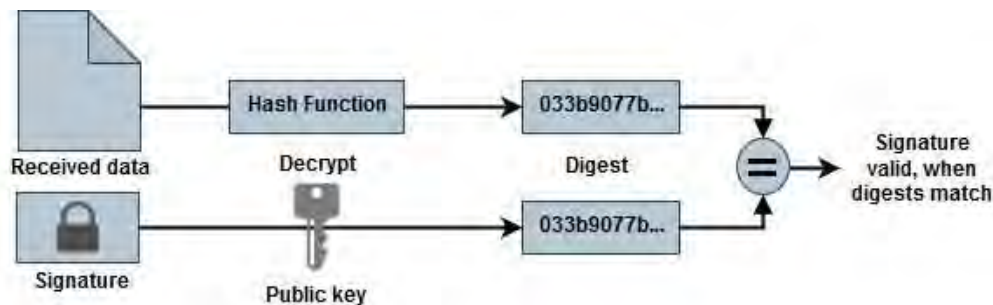
## Verify file signature

### Azure marketplace image signature verification for Cloud Volumes ONTAP

The Azure image verification process generates a digest file from the VHD file by stripping 1 MB at the beginning and 512 bytes at the end, then applying a hash function. To match the signing procedure, *sha256* is used for hashing.

### File signature verification workflow summary

The following is an overview of the file signature verification workflow process.



- Downloading the Azure image from the [NetApp Support Site](#) and extracting the digest (.sig) file, public key certificate (.pem) file, and chain certificate (.pem) file. Refer to [Download the Azure image digest file](#) for more information.
- Verification of the chain of trust.
- Extracting the public key (.pub) from the public key certificate (.pem).
- Decrypting the digest file by using the extracted public key.
- Comparing the result against a newly generated digest of a temporary file created from the image file after removing 1 MB at the beginning and 512 bytes at the end. This step is performed by using the OpenSSL command line tool. The OpenSSL CLI tool displays appropriate messaging on success or failure in matching the files.

```
openssl dgst -verify <public_key> -keyform <form> <hash_function>
-signature <digest_file> -binary <temporary_file>
```

## Verify Azure marketplace image signature for Cloud Volumes ONTAP on Linux

Verification of an exported VHD file signature on Linux includes validating the chain of trust, editing the file, and verifying the signature.

### Steps

1. Download the Azure image file from the [NetApp Support Site](#) and extract the digest (.sig) file, public key certificate (.pem) file, and chain certificate (.pem) file.

Refer to [Download the Azure image digest file](#) for more information.

2. Verify the chain of trust.

```
% openssl verify -CAfile Certificate-Chain-9.15.0P1_azure.pem
Certificate-9.15.0P1_azure.pem
Certificate-9.15.0P1_azure.pem: OK
```

3. Remove 1 MB (1,048,576 bytes) at the beginning and 512 bytes at the end of the VHD file. When using tail, the -c +K option generates bytes from the Kth byte of the file. Therefore, it passes 1048577 to tail -c.

```
% tail -c +1048577 ./9150.01000024.05090105.vhd > ./sign.tmp.tail
% head -c -512 ./sign.tmp.tail > sign.tmp
% rm ./sign.tmp.tail
```

4. Use OpenSSL to extract the public key from the certificate and verify the stripped file (sign.tmp) with the signature file and the public key.

The command prompt displays messages indicating success or failure based on the verification.

```
% openssl x509 -pubkey -noout -in ./Certificate-9.15.0P1_azure.pem >
./Code-Sign-Cert-Public-key.pub

% openssl dgst -verify Code-Sign-Cert-Public-key.pub -keyform PEM
-sha256 -signature digest.sig -binary ./sign.tmp
Verification OK

% openssl dgst -verify Code-Sign-Cert-Public-key.pub -keyform PEM
-sha256 -signature digest.sig -binary ./another_file_from_nowhere.tmp
Verification Failure
```

## 5. Clean up the workspace.

```
% rm ./9150.01000024.05090105.vhd ./sign.tmp
% rm *.sig *.pub *.pem
```

## Verify Azure marketplace image signature for Cloud Volumes ONTAP on macOS

Verification of an exported VHD file signature on Linux includes validating the chain of trust, editing the file, and verifying the signature.

### Steps

1. Download the Azure image file from the [NetApp Support Site](#) and extract the digest (.sig) file, public key certificate (.pem) file, and chain certificate (.pem) file.

Refer to [Download the Azure image digest file](#) for more information.

2. Verify the chain of trust.

```
% openssl verify -CAfile Certificate-Chain-9.15.0P1_azure.pem
Certificate-9.15.0P1_azure.pem
Certificate-9.15.0P1_azure.pem: OK
```

3. Remove 1MB (1,048,576 bytes) at the beginning and 512 bytes at the end of the VHD file. When using `tail`, the `-c +K` option generates bytes from the Kth byte of the file. Therefore, it passes 1048577 to `tail -c`. Note that on macOS, the `tail` command might take about ten minutes to complete.

```
% tail -c +1048577 ./9150.01000024.05090105.vhd > ./sign.tmp.tail
% head -c -512 ./sign.tmp.tail > sign.tmp
% rm ./sign.tmp.tail
```

4. Use OpenSSL to extract the public key from the certificate and verify the stripped file (sign.tmp) with the signature file and public key. The command prompt displays messages indicating success or failure based on the verification.

```
% openssl x509 -pubkey -noout -in ./Certificate-9.15.0P1_azure.pem >
./Code-Sign-Cert-Public-key.pub

% openssl dgst -verify Code-Sign-Cert-Public-key.pub -keyform PEM
-sha256 -signature digest.sig -binary ./sign.tmp
Verified OK

% openssl dgst -verify Code-Sign-Cert-Public-key.pub -keyform PEM
-sha256 -signature digest.sig -binary ./another_file_from_nowhere.tmp
Verification Failure
```

## 5. Clean up the workspace.

```
% rm ./9150.01000024.05090105.vhd ./sign.tmp
% rm *.sig *.pub *.pem
```

## Deploy Cloud Volumes ONTAP from the Azure marketplace

You can use Azure marketplace direct deployment to quickly and easily deploy Cloud Volumes ONTAP. From the Azure marketplace, you can quickly deploy Cloud Volumes ONTAP in a few clicks and explore its core features and capabilities in your environment.

For more information about this offering, refer to [Learn about Cloud Volumes ONTAP offerings on BlueXP and the marketplace](#).

### About this task

The Cloud Volumes ONTAP system deployed by using Azure marketplace direct deployment has these properties. Note that the features of a standalone instance deployed through the Azure marketplace change when it is discovered in BlueXP.

- The latest Cloud Volumes ONTAP version (9.16.1 or later).
- A free license for Cloud Volumes ONTAP that is limited to 500 GiB of provisioned capacity. This license includes no NetApp support and has no expiry date.
- Two nodes configured in a high availability (HA) mode in a single availability zone (AZ), provisioned with default serial numbers. The storage virtual machines (storage VMs) are deployed in a [flexible orchestration mode](#).
- An aggregate for the instance created by default.
- A Premium SSD v2 Managed Disk of 500 GiB provisioned capacity, and a root and a data disk.
- One data storage VM deployed, with NFS, CIFS, iSCSI, and NVMe/TCP data-services. You cannot add any additional data storage VMs.
- Licenses installed for NFS, CIFS (SMB), iSCSI, Autonomous Ransomware Protection (ARP), SnapLock, and SnapMirror.
- [ONTAP temperature-sensitive storage efficiency \(TSSE\)](#), volume encryption, and external key-management enabled by default.
- These features are not supported:
  - FabricPool tiering
  - Changing the storage VM type
  - Fast write mode

### Before you begin

- Ensure that you have a valid Azure marketplace subscription.
- Ensure you meet the networking requirements for an [HA deployment in a single AZ](#) in Azure. Refer to [Set up Azure networking for Cloud Volumes ONTAP](#).
- You need to be assigned one of these Azure roles to deploy Cloud Volumes ONTAP:
  - The `contributor` role with the default permissions. For more information, refer to the [Azure](#)

[documentation: Azure built-in roles.](#)

- A custom RBAC role with the following permissions. For more information, refer to the [Azure documentation: Azure custom roles.](#)

```
"permissions": [  
  {  
    "actions": [  
      "Microsoft.AAD/register/action",  
      "Microsoft.Resources/subscriptions/resourceGroups/write",  
      "Microsoft.Network/loadBalancers/write",  
      "Microsoft.ClassicCompute/virtualMachines/write",  
      "Microsoft.Compute/capacityReservationGroups/deploy/action",  
      "Microsoft.ClassicCompute/virtualMachines/networkInterfaces/associatedNetworkSecurityGroups/write",  
      "Microsoft.Network/networkInterfaces/write",  
      "Microsoft.Compute/virtualMachines/write",  
      "Microsoft.Compute/virtualMachines/extensions/write",  
      "Microsoft.Resources/deployments/validate/action",  
      "Microsoft.Resources/subscriptions/resourceGroups/read",  
      "Microsoft.Network/virtualNetworks/write",  
      "Microsoft.Network/virtualNetworks/read",  
      "Microsoft.Network/networkSecurityGroups/write",  
      "Microsoft.Network/networkSecurityGroups/read",  
      "Microsoft.Compute/disks/write",  
      "Microsoft.Compute/virtualMachineScaleSets/write",  
      "Microsoft.Resources/deployments/write",  
      "Microsoft.Network/virtualNetworks/subnets/read",  
      "Microsoft.Network/virtualNetworks/subnets/write"  
    ],  
    "notActions": [],  
    "dataActions": [],  
    "notDataActions": []  
  }  
]
```



If you have registered the resource provider "Microsoft.storage" to your subscription, then you don't need the `Microsoft.AAD/register/action` permission. For more information, refer to the [Azure documentation: Azure permissions for Storage.](#)

## Steps

1. From the Azure marketplace site, search for NetApp products.
2. Select **NetApp Cloud Volumes ONTAP direct**.
3. Click **Create** to launch the deployment wizard.
4. Select a plan. The **Plan** list typically displays the latest releases of Cloud Volumes ONTAP.
5. In the **Basics** tab, provide these details:
  - **Subscription:** Select a subscription. The deployment will be linked to the subscription number.
  - **Resource group:** Use an existing resource group or create a new one. Resource groups help in allocating all resources, such as disks and storage VMs, within a single group for a Cloud Volumes

ONTAP system.

- **Region:** Select a region that supports Azure HA deployment in a single AZ. You see only the available regions on the list.
- **Size:** Select an storage VM size for the supported Premium SSD v2 Managed Disk.
- **Zone:** Select a zone for the region you selected.
- **Admin Password:** Set a password. You use this admin password to log in to the system after the deployment.
- **Confirm Password:** Re-enter the same password for confirmation.
  - In the **Network** tab, add a virtual network and a subnet, or select them from the lists.



To comply with Microsoft Azure restrictions, you should create a new subnet when setting up a new virtual network. Likewise, if you choose an existing network, you should select an existing subnet.

- To select a predefined network security group, select **Yes**. Select **No** to assign a predefined Azure network security group with the necessary traffic rules. For more information, refer to [Security group rules for Azure](#).
- In the **Advanced** tab confirm whether the two Azure features necessary for this deployment have been set. Refer to [Enable an Azure feature for Cloud Volumes ONTAP single AZ deployments](#) and [Enable high-availability mode for Cloud Volumes ONTAP in Azure](#).
- You can define name and value pairs for the resources or resource groups in the **Tags** tab.
- In the **Review + create** tab, review the details and start the deployment.

### After you finish

Select the notification icon to view the progress of your deployment. After Cloud Volumes ONTAP is deployed, you can view the storage VM listed for operations.

Once accessible, use ONTAP System Manager or the ONTAP CLI to log in to the storage VM with the admin credentials that you set. Thereafter, you can create volumes, LUNs, or shares and start utilizing the storage capabilities of Cloud Volumes ONTAP.

### Troubleshoot deployment issues

Cloud Volumes ONTAP systems deployed directly through the Azure marketplace do not include support from NetApp. If any issues arise during deployment, you can independently troubleshoot and resolve them.

### Steps

1. On the Azure marketplace site, go to **Boot diagnostics > Serial log**.
2. Download and investigate the serial logs.
3. Consult the product documentation and knowledge base (KB) articles for troubleshooting.
  - [Azure marketplace documentation](#)
  - [NetApp documentation](#)
  - [NetApp KB articles](#)

### Discover the deployed systems in BlueXP

You can discover the Cloud Volumes ONTAP systems that you deployed using Azure marketplace direct

deployment and manage them as working environments in BlueXP. The BlueXP Connector discovers the systems, adds them as working environments, applies the necessary BlueXP licenses, and unlocks the full capabilities of BlueXP for these systems. The original HA configuration in a single AZ with PSSD v2 Managed Disks is retained, and the system is registered to the same Azure subscription and resource group as the original deployment.

### About this task

On discovering the Cloud Volumes ONTAP systems deployed using Azure marketplace direct deployment, the BlueXP Connector performs these tasks:

- Replaces the free licenses of the discovered systems as regular capacity-based [Freemium licenses](#).
- Retains the existing capabilities of the deployed systems, and adds the additional capabilities of BlueXP, such as data protection, data management, and security features.
- Replaces the installed licenses on the nodes with new ONTAP licenses for NFS, CIFS (SMB), iSCSI, ARP, SnapLock, and SnapMirror.
- Converts the generic node serial numbers to unique serial numbers.
- Assigns new system tags on the resources as required.
- Converts the dynamic IP addresses of the instance to static IP addresses.
- Enables the functionalities of [FabricPool tiering](#), [AutoSupport](#), and [write-once-read-many](#) (WORM) storage on the deployed systems. You can activate these features from the BlueXP console when you need them.
- Registers the instances to the NSS accounts used to discover them.
- Enables capacity management features in [automatic and manual modes](#) for the discovered systems.

### Before you begin

Ensure that the deployment is complete on the Azure marketplace. The BlueXP Connector can discover the systems only when the deployment is complete and are available for discovery.

### Steps

In BlueXP, you follow the standard procedure for discovering existing systems. Refer to [Add an existing Cloud Volumes ONTAP system to BlueXP](#).

### After you finish

After the discovery is complete, you can view the systems listed as working environments in BlueXP. You can perform various management tasks, such as [expanding the aggregate](#), [adding volumes](#), [provisioning additional storage VMs](#), and [changing the instance types](#).

### Related links

Refer to the ONTAP documentation for more information about creating storage:

- [Create volumes for NFS](#)
- [Create LUNs for iSCSI](#)
- [Create shares for CIFS](#)

## Get started in Google Cloud

### Quick start for Cloud Volumes ONTAP in Google Cloud

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



**1**

## Create a Connector

If you don't have a [Connector](#) yet, you need to create one. [Learn how to create a Connector in Google Cloud](#)

Note that if you want to deploy Cloud Volumes ONTAP in a subnet where no internet access is available, then you need to manually install the Connector and access the BlueXP user interface that's running on that Connector. [Learn how to manually install the Connector in a location without internet access](#)

**2**

## Plan your configuration

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

[Learn more about planning your configuration.](#)

**3**

## Set up your networking

- a. Ensure that your VPC and subnets will support connectivity between the Connector and Cloud Volumes ONTAP.
- b. If you plan to enable data tiering, [configure the Cloud Volumes ONTAP subnet for Private Google Access](#).
- c. If you're deploying an HA pair, ensure that you have four VPCs, each with their own subnet.
- d. If you're using a shared VPC, provide the *Compute Network User* role to the Connector service account.
- e. Enable outbound internet access from the target VPC for NetApp AutoSupport.

This step isn't required if you're deploying Cloud Volumes ONTAP in a location where no internet access is available.

[Learn more about networking requirements.](#)

**4**

## Set up a service account

Cloud Volumes ONTAP requires a Google Cloud service account for two purposes. The first is when you enable [data tiering](#) to tier cold data to low-cost object storage in Google Cloud. The second is when you enable the [BlueXP backup and recovery](#) to back up volumes to low-cost object storage.

You can set up one service account and use it for both purposes. The service account must have the **Storage Admin** role.

[Read step-by-step instructions.](#)

**5**

## Enable Google Cloud APIs

[Enable the following Google Cloud APIs in your project.](#) These APIs are required to deploy the Connector and Cloud Volumes ONTAP.

- Cloud Deployment Manager V2 API
- Cloud Logging API

- Cloud Resource Manager API
- Compute Engine API
- Identity and Access Management (IAM) API

## 6

### Launch Cloud Volumes ONTAP using BlueXP

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

#### Related links

- [Creating a Connector from BlueXP](#)
- [Installing the Connector software on a Linux host](#)
- [What BlueXP does with Google Cloud permissions](#)

## Plan your Cloud Volumes ONTAP configuration in Google Cloud

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

### Choose a Cloud Volumes ONTAP license

Several licensing options are available for Cloud Volumes ONTAP. Each option enables you to choose a consumption model that meets your needs.

- [Learn about licensing options for Cloud Volumes ONTAP](#)
- [Learn how to set up licensing](#)

### Choose a supported region

Cloud Volumes ONTAP is supported in most Google Cloud regions. [View the full list of supported regions.](#)

### Choose a supported machine type

Cloud Volumes ONTAP supports several machine types, depending on the license type that you choose.

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

### Understand storage limits

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

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

### Size your system in GCP

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

## Machine type

Look at the supported machine types in the [Cloud Volumes ONTAP Release Notes](#) and then review details from Google about each supported machine type. Match your workload requirements to the number of vCPUs and memory for the machine type. Note that each CPU core increases networking performance.

Refer to the following for more details:

- [Google Cloud documentation: N1 standard machine types](#)
- [Google Cloud documentation: Performance](#)

## GCP disk type

When you create volumes for Cloud Volumes ONTAP, you need to choose the underlying cloud storage that Cloud Volumes ONTAP uses for a disk. The disk type can be any of the following:

- *Zonal SSD persistent disks*: SSD persistent disks are best for workloads that require high rates of random IOPS.
- *Zonal Balanced persistent disks*: These SSDs balance performance and cost by providing lower IOPS per GB.
- *Zonal Standard persistent disks* : Standard persistent disks are economical and can handle sequential read/write operations.

For more details, refer to [Google Cloud documentation: Zonal Persistent disks \(Standard and SSD\)](#).

## GCP disk size

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

- All disks in an aggregate must be the same size.
- Determine the space that you need, while taking performance into consideration.
- The performance of persistent disks scales automatically with disk size and the number of vCPUs available to the system.

Refer to the following for more details:

- [Google Cloud documentation: Zonal Persistent disks \(Standard and SSD\)](#)
- [Google Cloud documentation: Optimizing Persistent Disk and Local SSD Performance](#)

## View default system disks

In addition to the storage for user data, BlueXP also purchases cloud storage for Cloud Volumes ONTAP system data (boot data, root data, core data, and NVRAM). For planning purposes, it might help for you to review these details before you deploy Cloud Volumes ONTAP.

- [View the default disks for Cloud Volumes ONTAP system data in Google Cloud.](#)
- [Google Cloud docs: Resource quotas](#)

Google Cloud Compute Engine enforces quotas on resource usage so you should ensure that you haven't reached your limit before you deploy Cloud Volumes ONTAP.



The Connector also requires a system disk. [View details about the Connector's default configuration.](#)

### Collect networking information

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

#### Network information for a single-node system

GCP information	Your value
Region	
Zone	
VPC network	
Subnet	
Firewall policy (if using your own)	

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

GCP information	Your value
Region	
Zone for Node 1	
Zone for Node 2	
Zone for the mediator	
VPC-0 and subnet	
VPC-1 and subnet	
VPC-2 and subnet	
VPC-3 and subnet	
Firewall policy (if using your own)	

#### Network information for an HA pair in a single zone

GCP information	Your value
Region	
Zone	
VPC-0 and subnet	
VPC-1 and subnet	
VPC-2 and subnet	
VPC-3 and subnet	

GCP information	Your value
Firewall policy (if using your own)	

## Choose a write speed

BlueXP enables you to choose a write speed setting for Cloud Volumes ONTAP, except for high-availability (HA) pairs in Google Cloud. Before you choose a write speed, you should understand the differences between the normal and high settings and risks and recommendations when using high write speed. [Learn more about write speed.](#)

## Choose a volume usage profile

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

NetApp storage efficiency features provide the following benefits:

### Thin provisioning

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

### Deduplication

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

### Compression

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

## Set up Google Cloud networking for Cloud Volumes ONTAP

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

If you want to deploy an HA pair, you should [learn how HA pairs work in Google Cloud.](#)

## Requirements for Cloud Volumes ONTAP

The following requirements must be met in Google Cloud.

### Requirements specific to single node systems

If you want to deploy a single node system, ensure that your networking meets the following requirements.

#### One VPC

One Virtual Private Cloud (VPC) is required for a single node system.

## Private IP addresses

For a single node system in Google Cloud, BlueXP allocates private IP addresses to the following:

- Node
- Cluster
- Storage VM
- Data NAS LIF
- Data iSCSI LIF

You can skip creation of the storage VM (SVM) management LIF if you deploy Cloud Volumes ONTAP using the API and specify the following flag:

```
skipSvmManagementLif: true
```



A LIF is an IP address associated with a physical port. A storage VM (SVM) management LIF is required for management tools like SnapCenter.

## Requirements specific to HA pairs

If you want to deploy an HA pair, ensure that your networking meets the following requirements.

### One or multiple zones

You can ensure the high availability of your data by deploying an HA configuration across multiple or in a single zone. BlueXP will prompt you to choose multiple zones or a single zone when you create the HA pair.

- Multiple zones (recommended)

Deploying an HA configuration across three zones ensures continuous data availability if a failure occurs within a zone. Note that write performance is slightly lower compared to using a single zone, but it's minimal.

- Single zone

When deployed in a single zone, a Cloud Volumes ONTAP HA configuration uses a spread placement policy. This policy ensures that an HA configuration is protected from a single point of failure within the zone, without having to use separate zones to achieve fault isolation.

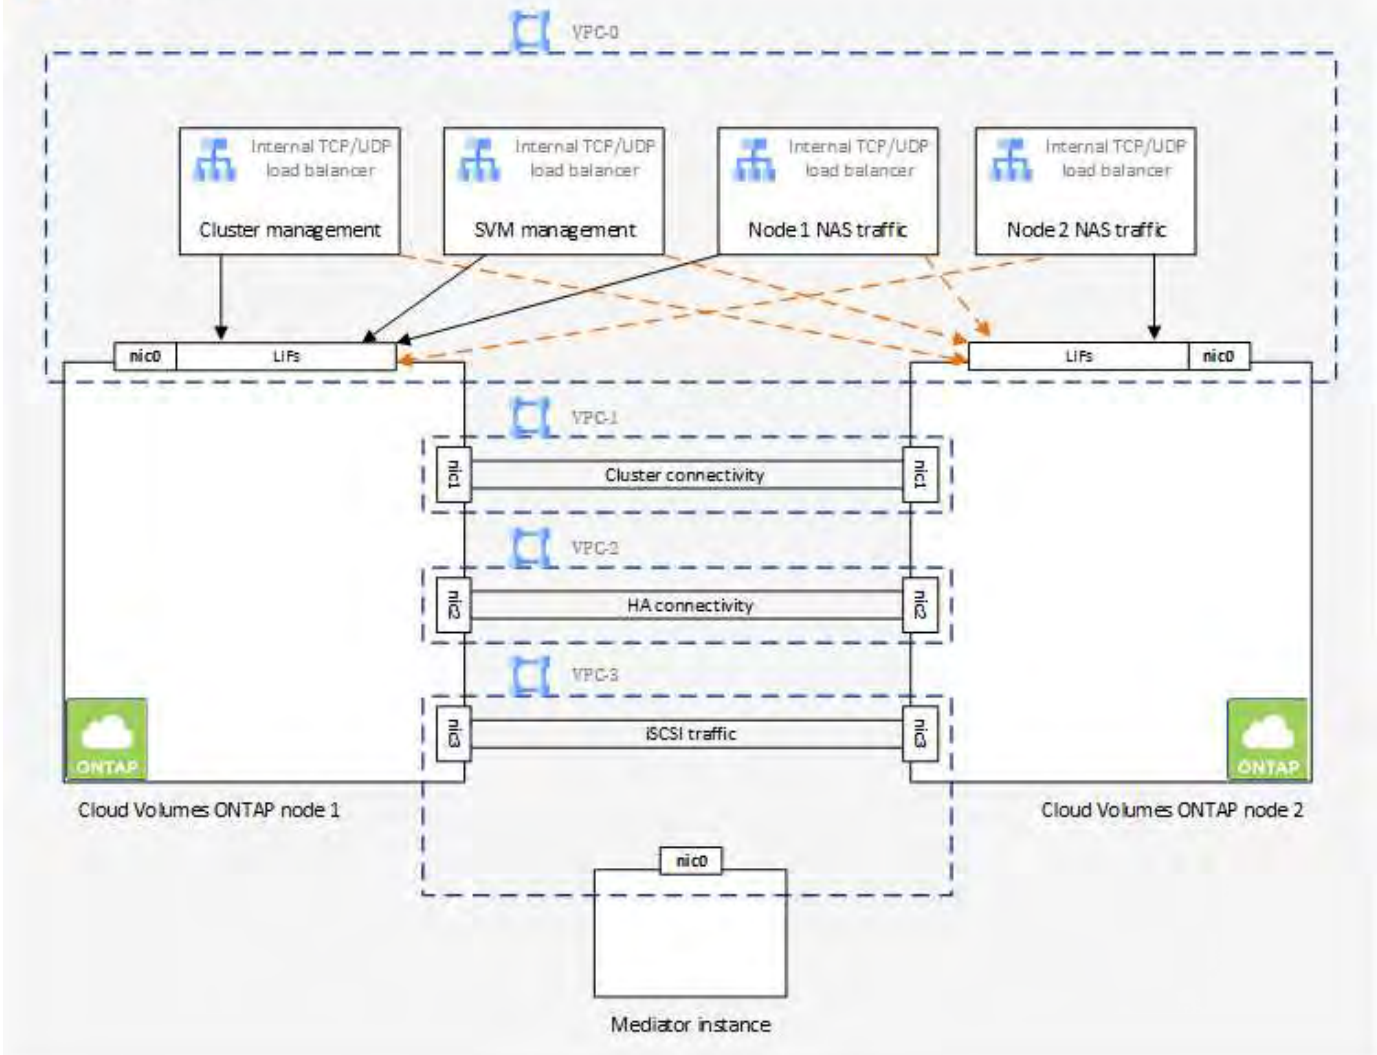
This deployment model does lower your costs because there are no data egress charges between zones.

## Four Virtual Private Clouds

Four Virtual Private Clouds (VPCs) are required for an HA configuration. Four VPCs are required because Google Cloud requires that each network interface resides in a separate VPC network.

BlueXP will prompt you to choose four VPCs when you create the HA pair:

- VPC-0 for inbound connections to the data and nodes
- VPC-1, VPC-2, and VPC-3 for internal communication between the nodes and the HA mediator



## Subnets

A private subnet is required for each VPC.

If you place the Connector in VPC-0, then you will need to enable Private Google Access on the subnet to access the APIs and to enable data tiering.

The subnets in these VPCs must have distinct CIDR ranges. They can't have overlapping CIDR ranges.

## Private IP addresses

BlueXP automatically allocates the required number of private IP addresses to Cloud Volumes ONTAP in Google Cloud. You need to make sure that your networking has enough private addresses available.

The number of LIFs that BlueXP allocates for Cloud Volumes ONTAP depends on whether you deploy a single node system or an HA pair. A LIF is an IP address associated with a physical port. An SVM management LIF is required for management tools like SnapCenter.

- **Single node**

BlueXP allocates 4 IP addresses to a single node system:

- Node management LIF
- Cluster management LIF
- iSCSI data LIF



An iSCSI LIF provides client access over the iSCSI protocol and is used by the system for other important networking workflows. These LIFs are required and should not be deleted.

- NAS LIF

You can skip creation of the storage VM (SVM) management LIF if you deploy Cloud Volumes ONTAP using the API and specify the following flag:

```
skipSvmManagementLif: true
```

#### • HA pair

BlueXP allocates 12-13 IP addresses to an HA pair:

- 2 Node management LIFs (e0a)
- 1 Cluster management LIF (e0a)
- 2 iSCSI LIFs (e0a)



An iSCSI LIF provides client access over the iSCSI protocol and is used by the system for other important networking workflows. These LIFs are required and should not be deleted.

- 1 or 2 NAS LIFs (e0a)
- 2 Cluster LIFs (e0b)
- 2 HA Interconnect IP addresses (e0c)
- 2 RSM iSCSI IP addresses (e0d)

You can skip creation of the storage VM (SVM) management LIF if you deploy Cloud Volumes ONTAP using the API and specify the following flag:

```
skipSvmManagementLif: true
```

### Internal load balancers

BlueXP automatically creates four Google Cloud internal load balancers (TCP/UDP) that manage incoming traffic to the Cloud Volumes ONTAP HA pair. No setup is required from your end. We've listed this as a requirement simply to inform you of the network traffic and to mitigate any security concerns.

One load balancer is for cluster management, one is for storage VM (SVM) management, one is for NAS traffic to node 1, and the last is for NAS traffic to node 2.

The setup for each load balancer is as follows:

- One shared private IP address
- One global health check



By default, the ports used by the health check are 63001, 63002, and 63003.

- One regional TCP backend service
- One regional UDP backend service
- One TCP forwarding rule
- One UDP forwarding rule
- Global access is disabled

Even though global access is disabled by default, enabling it post deployment is supported. We disabled it because cross region traffic will have significantly higher latencies. We wanted to ensure that you didn't have a negative experience due to accidental cross region mounts. Enabling this option is specific to your business needs.

### Shared VPCs

Cloud Volumes ONTAP and the Connector are supported in a Google Cloud shared VPC and also in standalone VPCs.

For a single node system, the VPC can be either a shared VPC or a standalone VPC.

For an HA pair, four VPCs are required. Each of those VPCs can be either shared or standalone. For example, VPC-0 could be a shared VPC, while VPC-1, VPC-2, and VPC-3 could be standalone VPCs.

A shared VPC enables you to configure and centrally manage virtual networks across multiple projects. You can set up shared VPC networks in the *host project* and deploy the Connector and Cloud Volumes ONTAP virtual machine instances in a *service project*. [Google Cloud documentation: Shared VPC overview](#).

[Review the required shared VPC permissions covered in Connector deployment](#)

### Packet mirroring in VPCs

[Packet mirroring](#) must be disabled in the Google Cloud subnet in which you deploy Cloud Volumes ONTAP.

### Outbound internet access

Cloud Volumes ONTAP systems require outbound internet access for accessing external endpoints for various functions. Cloud Volumes ONTAP can't operate properly if these endpoints are blocked in environments with strict security requirements.

The BlueXP Connector also contacts several endpoints for day-to-day operations, as well as the BlueXP web-based console. For information about the BlueXP endpoints, refer to [View endpoints contacted from the Connector](#) and [Prepare networking for using the BlueXP console](#).

### Cloud Volumes ONTAP endpoints

Cloud Volumes ONTAP uses these endpoints to communicate with various services.

Endpoints	Applicable for	Purpose	BlueXP deployment mode	Impact if endpoint is not available
https://netapp-cloud-account.auth0.com	Authentication	Used for BlueXP authentication.	Standard and restricted modes.	User authentication fails and the following services remain unavailable: <ul style="list-style-type: none"> <li>• Cloud Volumes ONTAP services</li> <li>• ONTAP services</li> <li>• Protocols and proxy services</li> </ul>
https://cloudmanager.cloud.netapp.com/tenancy	Tenancy	Used to retrieve the Cloud Volumes ONTAP resources from BlueXP tenancy to authorize resources and users.	Standard and restricted modes.	Cloud Volumes ONTAP resources and the users are not authorized.
https://support.netapp.com/aods/asupmessage https://support.netapp.com/asupprod/post/1.0/postAsup	AutoSupport	Used to send AutoSupport telemetry data to NetApp support.	Standard and restricted modes.	AutoSupport information remains undelivered.
https://www.googleapis.com/compute/v1/projects/ https://cloudresourcemanager.googleapis.com/v1/projects https://www.googleapis.com/compute/beta https://storage.googleapis.com/storage/v1 https://www.googleapis.com/storage/v1 https://iam.googleapis.com/v1 https://cloudkms.googleapis.com/v1 https://www.googleapis.com/deploymentmanager/v2/projects https://compute.googleapis.com/compute/v1	Google Cloud (Commercial use).	Communication with Google Cloud services.	Standard, restricted, and private modes.	Cloud Volumes ONTAP cannot communicate with Google Cloud service to perform specific BlueXP operations in Google Cloud.

### Connections to ONTAP systems in other networks

To replicate data between a Cloud Volumes ONTAP system in Google Cloud and ONTAP systems in other networks, you must have a VPN connection between the VPC and the other network—for example, your corporate network.

For instructions, refer to [Google Cloud documentation: Cloud VPN overview](#).

## Firewall rules

BlueXP creates Google Cloud firewall rules that include the inbound and outbound rules that Cloud Volumes ONTAP needs to operate successfully. You might want to refer to the ports for testing purposes or if you prefer to use your own firewall rules.

The firewall rules for Cloud Volumes ONTAP requires both inbound and outbound rules. If you're deploying an HA configuration, these are the firewall rules for Cloud Volumes ONTAP in VPC-0.

Note that two sets of firewall rules are required for an HA configuration:

- One set of rules for HA components in VPC-0. These rules enable data access to Cloud Volumes ONTAP.
- Another set of rules for HA components in VPC-1, VPC-2, and VPC-3. These rules are open for inbound & outbound communication between the HA components. [Learn more](#).



Looking for information about the Connector? [View firewall rules for the Connector](#)

## Inbound rules

When you create a working environment, you can choose the source filter for the predefined firewall policy during deployment:

- **Selected VPC only:** the source filter for inbound traffic is the subnet range of the VPC for the Cloud Volumes ONTAP system and the subnet range of the VPC where the Connector resides. This is the recommended option.
- **All VPCs:** the source filter for inbound traffic is the 0.0.0.0/0 IP range.

If you use your own firewall policy, ensure that you add all networks that need to communicate with Cloud Volumes ONTAP, but also ensure to add both address ranges to allow the internal Google Load Balancer to function correctly. These addresses are 130.211.0.0/22 and 35.191.0.0/16. For more information, refer to [Google Cloud documentation: Load Balancer Firewall Rules](#).

Protocol	Port	Purpose
All ICMP	All	Pinging the instance
HTTP	80	HTTP access to the ONTAP System Manager web console using the IP address of the cluster management LIF
HTTPS	443	Connectivity with the Connector and HTTPS access to the ONTAP System Manager web console using the IP address of the cluster management LIF
SSH	22	SSH access to the IP address of the cluster management LIF or a node management LIF
TCP	111	Remote procedure call for NFS
TCP	139	NetBIOS service session for CIFS
TCP	161-162	Simple network management protocol
TCP	445	Microsoft SMB/CIFS over TCP with NetBIOS framing
TCP	635	NFS mount
TCP	749	Kerberos
TCP	2049	NFS server daemon

Protocol	Port	Purpose
TCP	3260	iSCSI access through the iSCSI data LIF
TCP	4045	NFS lock daemon
TCP	4046	Network status monitor for NFS
TCP	10000	Backup using NDMP
TCP	11104	Management of intercluster communication sessions for SnapMirror
TCP	11105	SnapMirror data transfer using intercluster LIFs
TCP	63001-63050	Load balance probe ports to determine which node is healthy (required for HA pairs only)
UDP	111	Remote procedure call for NFS
UDP	161-162	Simple network management protocol
UDP	635	NFS mount
UDP	2049	NFS server daemon
UDP	4045	NFS lock daemon
UDP	4046	Network status monitor for NFS
UDP	4049	NFS rquotad protocol

### Outbound rules

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

#### Basic outbound rules

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

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

#### Advanced outbound rules

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



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

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

Service	Protocol	Port	Source	Destination	Purpose
AutoSupport	HTTPS	443	Node management LIF	support.netapp.com	AutoSupport (HTTPS is the default)
	HTTP	80	Node management LIF	support.netapp.com	AutoSupport (only if the transport protocol is changed from HTTPS to HTTP)
	TCP	3128	Node management LIF	Connector	Sending AutoSupport messages through a proxy server on the Connector, if an outbound internet connection isn't available
Configuration backups	HTTP	80	Node management LIF	http://<connector-IP-address>/occm/offbo xconfig	Send configuration backups to the Connector. <a href="#">ONTAP documentation</a>
DHCP	UDP	68	Node management LIF	DHCP	DHCP client for first-time setup
DHCPS	UDP	67	Node management LIF	DHCP	DHCP server
DNS	UDP	53	Node management LIF and data LIF (NFS, CIFS)	DNS	DNS
NDMP	TCP	1860 0–18 699	Node management LIF	Destination servers	NDMP copy
SMTP	TCP	25	Node management LIF	Mail server	SMTP alerts, can be used for AutoSupport
SNMP	TCP	161	Node management LIF	Monitor server	Monitoring by SNMP traps
	UDP	161	Node management LIF	Monitor server	Monitoring by SNMP traps
	TCP	162	Node management LIF	Monitor server	Monitoring by SNMP traps
	UDP	162	Node management LIF	Monitor server	Monitoring by SNMP traps
SnapMirror	TCP	1110 4	Intercluster LIF	ONTAP intercluster LIFs	Management of intercluster communication sessions for SnapMirror
	TCP	1110 5	Intercluster LIF	ONTAP intercluster LIFs	SnapMirror data transfer
Syslog	UDP	514	Node management LIF	Syslog server	Syslog forward messages

### Rules for VPC-1, VPC-2, and VPC-3

In Google Cloud, an HA configuration is deployed across four VPCs. The firewall rules needed for the HA

configuration in VPC-0 are [listed above for Cloud Volumes ONTAP](#).

Meanwhile, the predefined firewall rules that BlueXP creates for instances in VPC-1, VPC-2, and VPC-3 enables ingress communication over *all* protocols and ports. These rules enable communication between HA nodes.

Communication from the HA nodes to the HA mediator takes place over port 3260 (iSCSI).



To enable high write speed for new Google Cloud HA pair deployments, a maximum transmission unit (MTU) of at least 8,896 bytes is required for VPC-1, VPC-2, and VPC-3. If you choose to upgrade existing VPC-1, VPC-2, and VPC-3 to an MTU of 8,896 bytes, you must shutdown all existing HA systems using these VPCs during the configuration process.

## Requirements for the Connector

If you haven't created a Connector yet, you should review networking requirements for the Connector as well.

- [View networking requirements for the Connector](#)
- [Firewall rules in Google Cloud](#)

## Network configurations to support Connector proxy

You can use the proxy servers configured for the BlueXP Connector to enable outbound internet access from Cloud Volumes ONTAP. BlueXP supports two types of proxies:

- **Explicit proxy:** The outbound traffic from Cloud Volumes ONTAP uses the HTTP address of the proxy server specified during the Connector proxy configuration. The Connector administrator might also have configured user credentials and root CA certificates for additional authentication. If a root CA certificate is available for the explicit proxy, make sure to obtain and upload the same certificate to your Cloud Volumes ONTAP working environment using the [ONTAP CLI: security certificate install](#) command.
- **Transparent proxy:** The network is configured to automatically route outbound traffic from Cloud Volumes ONTAP through the Connector proxy. When setting up a transparent proxy, the Connector administrator needs to provide only a root CA certificate for connectivity from Cloud Volumes ONTAP, not the HTTP address of the proxy server. Make sure that you obtain and upload the same root CA certificate to your Cloud Volumes ONTAP working environment using the [ONTAP CLI: security certificate install](#) command.

For information about configuring proxy servers for the BlueXP Connector, refer to the [Configure a Connector to use a proxy server](#).

## Configure network tags for Cloud Volumes ONTAP in Google Cloud

During the Connector transparent proxy configuration, the Connector administrator adds a network tag for Google Cloud. You need to obtain and manually add the same network tag for your Cloud Volumes ONTAP configuration. This tag is necessary for the proxy server to function correctly.

1. In the Google Cloud console, locate your Cloud Volumes ONTAP working environment.
2. Go to **Details > Networking > Network tags**.
3. Add the tag used for the Connector and save the configuration.

## Related topics

- [Verify AutoSupport setup for Cloud Volumes ONTAP](#)
- [Learn about ONTAP internal ports](#).

## Set up VPC Service Controls to deploy Cloud Volumes ONTAP in Google Cloud

When choosing to lock down your Google Cloud environment with VPC Service Controls, you should understand how BlueXP and Cloud Volumes ONTAP interact with the Google Cloud APIs, as well as how to configure your service perimeter to deploy BlueXP and Cloud Volumes ONTAP.

VPC Service Controls enable you to control access to Google-managed services outside of a trusted perimeter, to block data access from untrusted locations, and to mitigate unauthorized data transfer risks. [Learn more about Google Cloud VPC Service Controls.](#)

### How NetApp services communicate with VPC Service Controls

BlueXP communicates directly with the Google Cloud APIs. This is either triggered from an external IP address outside of Google Cloud (for example, from `api.services.cloud.netapp.com`), or within Google Cloud from an internal address assigned to the BlueXP Connector.

Depending on the deployment style of the Connector, certain exceptions may have to be made for your service perimeter.

### Images

Both Cloud Volumes ONTAP and BlueXP use images from a project within GCP that is managed by NetApp. This can affect the deployment of the BlueXP Connector and Cloud Volumes ONTAP, if your organization has a policy that blocks the use of images that are not hosted within the organization.

You can deploy a Connector manually using the manual installation method, but Cloud Volumes ONTAP will also need to pull images from the NetApp project. You must provide an allowed list in order to deploy a Connector and Cloud Volumes ONTAP.

#### Deploying a Connector

The user who deploys a Connector needs to be able to reference an image hosted in the projectId `netapp-cloudmanager` and the project number `14190056516`.

#### Deploying Cloud Volumes ONTAP

- The BlueXP service account needs to reference an image hosted in the projectId `netapp-cloudmanager` and the project number `14190056516` from the service project.
- The service account for the default Google APIs Service Agent needs to reference an image hosted in the projectId `netapp-cloudmanager` and the project number `14190056516` from the service project.

Examples of the rules needed for pulling these images with VPC Service Controls are defined below.

### VPC Service Controls perimeter policies

Policies allow exceptions to the VPC Service Controls rule sets. For more information about policies, please visit the [GCP VPC Service Controls Policy Documentation](#).

To set the policies that BlueXP requires, navigate to your VPC Service Controls Perimeter within your organization and add the following policies. The fields should match the options given in the VPC Service Controls policy page. Also note that **all** rules are required and the **OR** parameters should be used in the rule set.



## Ingress rules

### Rule 1

```
From:
  Identities:
    [User Email Address]
  Source > All sources allowed
To:
  Projects =
    [Service Project]
  Services =
    Service name: iam.googleapis.com
    Service methods: All actions
    Service name: compute.googleapis.com
    Service methods:All actions
```

OR

### Rule 2

```
From:
  Identities:
    [User Email Address]
  Source > All sources allowed
To:
  Projects =
    [Host Project]
  Services =
    Service name: compute.googleapis.com
    Service methods: All actions
```

OR

### Rule 3

```
From:
  Identities:
    [Service Project Number]@cloudservices.gserviceaccount.com
  Source > All sources allowed
To:
  Projects =
    [Service Project]
    [Host Project]
  Services =
    Service name: compute.googleapis.com
    Service methods: All actions
```

### Egress rules

#### Rule 1:

```
From:
  Identities:
    [Service Project Number]@cloudservices.gserviceaccount.com
To:
  Projects =
    14190056516
  Service =
    Service name: compute.googleapis.com
    Service methods: All actions
```



The project number outlined above is the project *netapp-cloudmanager* used by NetApp to store images for the Connector and for Cloud Volumes ONTAP.

## Create a Google Cloud service account for Cloud Volumes ONTAP

Cloud Volumes ONTAP requires a Google Cloud service account for two purposes. The first is when you enable [data tiering](#) to tier cold data to low-cost object storage in Google Cloud. The second is when you enable the [BlueXP backup and recovery](#) to back up volumes to low-cost object storage.

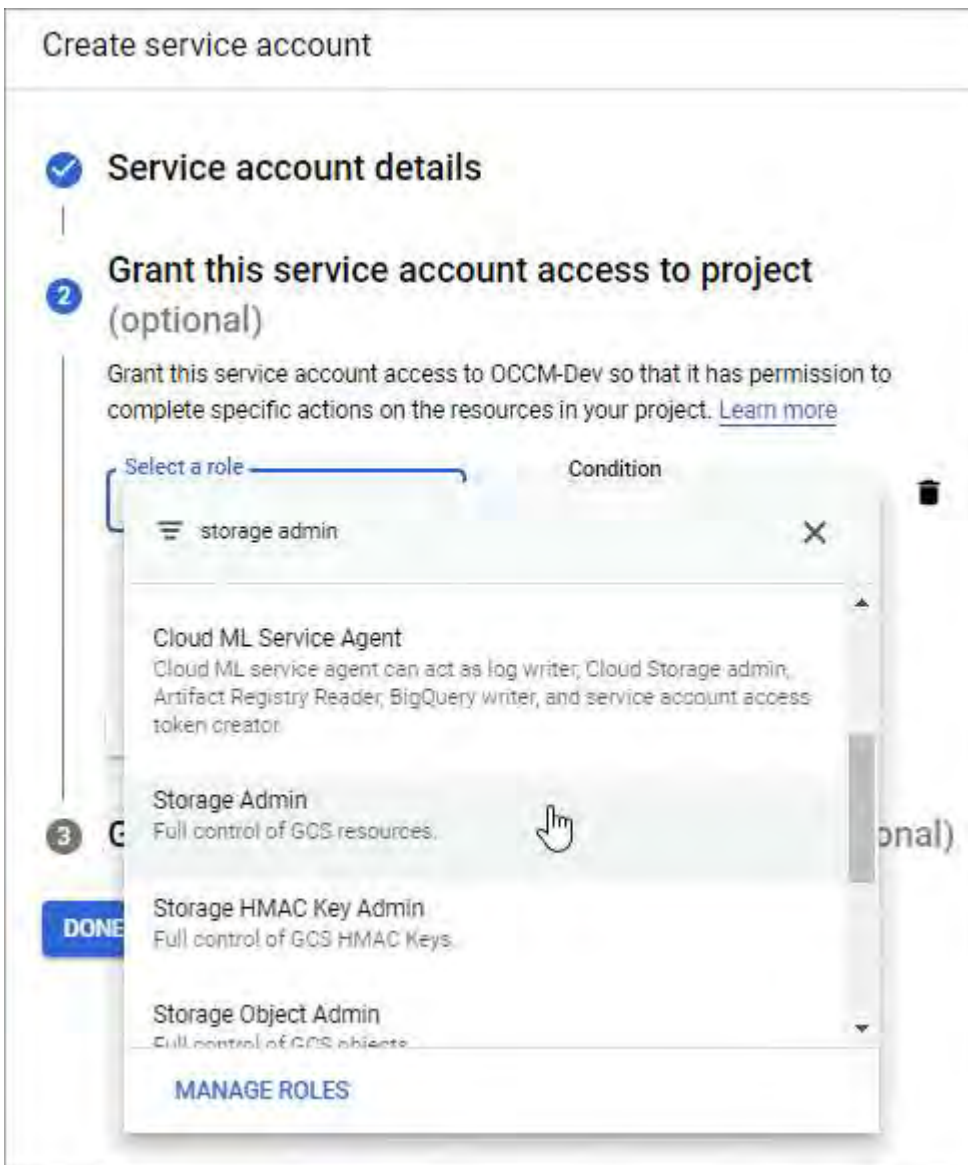
Cloud Volumes ONTAP uses the service account to access and manage one bucket for tiered data and another bucket for backups.

You can set up one service account and use it for both purposes. The service account must have the **Storage Admin** role.

### Steps

1. In the Google Cloud console, [go to the Service accounts page](#).

2. Select your project.
3. Click **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 **Storage Admin** role.



- c. **Grant users access to this service account:** Add the Connector service account as a *Service Account User* to this new service account.

This step is required for data tiering only. It's not required for BlueXP backup and recovery.

Create service account

- ✓ Service account details
- ✓ Grant this service account access to project (optional)
- 3 Grant users access to this service account (optional)  
Grant access to users or groups that need to perform actions as this service account. [Learn more](#)

Service account users role

netapp-cloud-manager@iam.gserviceaccount.com

Grant users the permissions to deploy jobs and VMs with this service account

Service account admins role

Grant users the permission to administer this service account

**DONE** CANCEL

**What's next?**

You'll need to select the service account later when you create a Cloud Volumes ONTAP working environment.

## Details and Credentials

default-project Google Cloud Project	gcp-sub2 Marketplace Subscription	<div style="border: 1px solid #0070c0; padding: 2px 5px; display: inline-block;">Edit Project</div>
---	--------------------------------------	---

### Details

Working Environment Name (Cluster Name)

Service Account

---

Service Account Name

+ Add Labels Optional Field | Up to four labels

### Credentials

User Name

Password

Confirm Password

## Using customer-managed encryption keys with Cloud Volumes ONTAP

While Google Cloud Storage always encrypts your data before it's written to disk, you can use the BlueXP API to create a Cloud Volumes ONTAP system that uses *customer-managed encryption keys*. These are keys that you generate and manage in GCP using the Cloud Key Management Service.

### Steps

1. Ensure that the BlueXP Connector service account has the correct permissions at the project level, in the project where the key is stored.

The permissions are provided in the [Connector service account permissions by default](#), but may not be applied if you use an alternate project for the Cloud Key Management Service.

The permissions are as follows:

```

- cloudkms.cryptoKeyVersions.useToEncrypt
- cloudkms.cryptoKeys.get
- cloudkms.cryptoKeys.list
- cloudkms.keyRings.list

```

2. Ensure that the service account for the [Google Compute Engine Service Agent](#) has Cloud KMS Encrypter/Decrypter permissions on the key.

The name of the service account uses the following format: "service-[service\_project\_number]@compute-system.iam.gserviceaccount.com".

[Google Cloud Documentation: Using IAM with Cloud KMS - Granting roles on a resource](#)

3. Obtain the "id" of the key by invoking the get command for the `/gcp/vsa/metadata/gcp-encryption-keys` API call or by choosing "Copy Resource Name" on the key in the GCP console.
4. If using customer-managed encryption keys and tiering data to object storage, BlueXP attempts to utilize the same keys that are used to encrypt the persistent disks. But you'll first need to enable Google Cloud Storage buckets to use the keys:
  - a. Find the Google Cloud Storage service agent by following the [Google Cloud Documentation: Getting the Cloud Storage service agent](#).
  - b. Navigate to the encryption key and assign the Google Cloud Storage service agent with Cloud KMS Encrypter/Decrypter permissions.

For more information, refer to [Google Cloud Documentation: Using customer-managed encryption keys](#)

5. Use the "GcpEncryption" parameter with your API request when creating a working environment.

### Example

```
"gcpEncryptionParameters": {  
  "key": "projects/project-1/locations/us-east4/keyRings/keyring-  
1/cryptoKeys/generatedkey1"  
}
```

Refer to the [BlueXP automation docs](#) for more details about using the "GcpEncryption" parameter.

## Set up licensing for Cloud Volumes ONTAP in Google Cloud

After you decide which licensing option you want to use with Cloud Volumes ONTAP, a few steps are required before you can choose that licensing option when creating a new working environment.

### Freemium

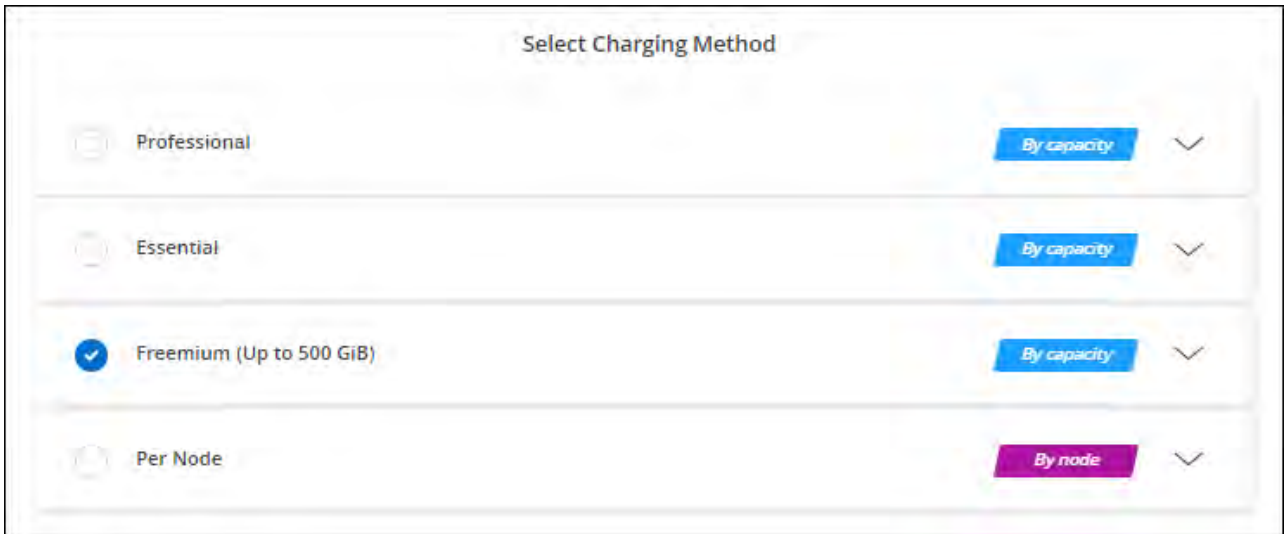
Select the Freemium offering to use Cloud Volumes ONTAP free of charge with up to 500 GiB of provisioned capacity. [Learn more about the Freemium offering](#).

### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the pay-as-you-go offering in the Google Cloud Marketplace.

You won't be charged through the marketplace subscription unless you exceed 500 GiB of provisioned capacity, at which time the system is automatically converted to the [Essentials package](#).

b. After you return to BlueXP, select **Freemium** when you reach the charging methods page.



[View step-by-step instructions to launch Cloud Volumes ONTAP in Google Cloud.](#)

## Capacity-based license

Capacity-based licensing enables you to pay for Cloud Volumes ONTAP per TiB of capacity. Capacity-based licensing is available in the form of a *package*: the Essentials or Professional package.

The Essentials and Professional packages are available with the following consumption models:

- A license (bring your own license (BYOL)) purchased from NetApp
- An hourly, pay-as-you-go (PAYGO) subscription from the Google Cloud Marketplace
- An annual contract

[Learn more about capacity-based licensing.](#)

The following sections describe how to get started with each of these consumption models.

### BYOL

Pay upfront by purchasing a license (BYOL) from NetApp to deploy Cloud Volumes ONTAP systems in any cloud provider.



NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAP](#).

### Steps

1. [Contact NetApp Sales to obtain a license](#)
2. [Add your NetApp Support Site account to BlueXP](#)

BlueXP automatically queries NetApp's licensing service to obtain details about the licenses associated with your NetApp Support Site account. If there are no errors, BlueXP automatically adds the licenses to the digital wallet.

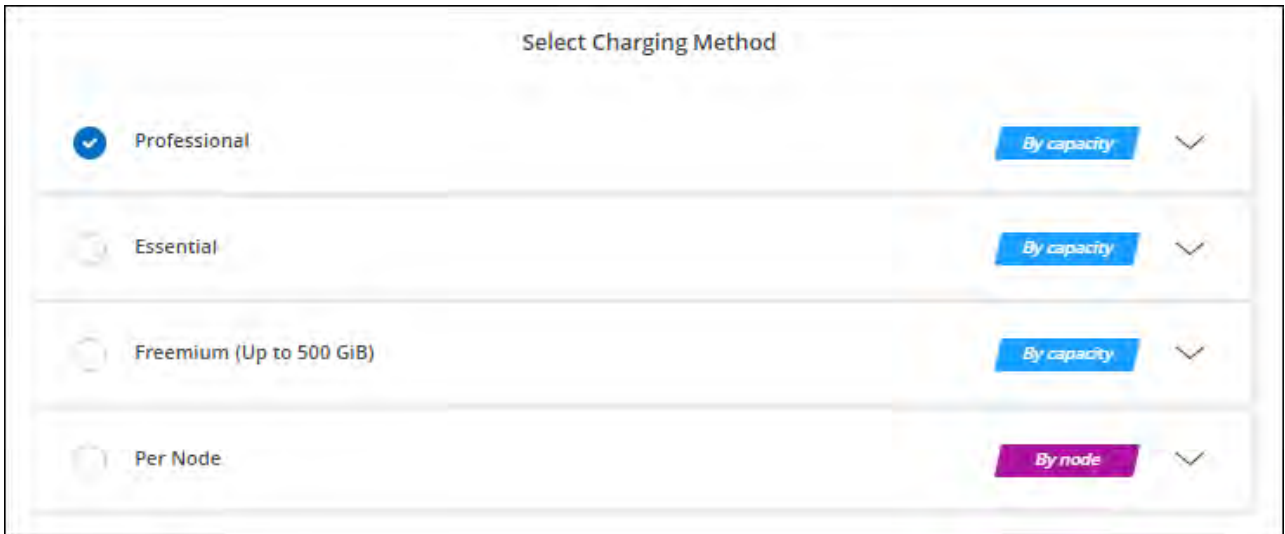
Your license must be available from the BlueXP digital wallet before you can use it with Cloud Volumes

ONTAP. If needed, you can [manually add the license to the BlueXP digital wallet](#).

3. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the pay-as-you-go offering in the Google Cloud Marketplace.

The license that you purchased from NetApp is always charged first, but you'll be charged from the hourly rate in the marketplace if you exceed your licensed capacity or if the term of your license expires.

- b. After you return to BlueXP, select a capacity-based package when you reach the charging methods page.



[View step-by-step instructions to launch Cloud Volumes ONTAP in Google Cloud.](#)

### **PAYGO subscription**


Pay hourly by subscribing to the offer from your cloud provider's marketplace.

When you create a Cloud Volumes ONTAP working environment, BlueXP prompts you to subscribe to the agreement that's available in the Google Cloud Marketplace. That subscription is then associated with the working environment for charging. You can use that same subscription for additional working environments.

### **Steps**

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the pay-as-you-go offering in the Google Cloud Marketplace.
  - b. After you return to BlueXP, select a capacity-based package when you reach the charging methods page.





[View step-by-step instructions to launch Cloud Volumes ONTAP in Google Cloud.](#)



You can manage the Google Cloud Marketplace subscriptions associated with your accounts from the Settings > Credentials page. [Learn how to manage your Google Cloud credentials and subscriptions](#)

### Annual contract

Pay for Cloud Volumes ONTAP annually by purchasing an annual contract.


### Steps

1. 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 working environment creation.

2. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. On the **Details and Credentials** page, click **Edit Credentials > Add Subscription** and then follow the prompts to subscribe to the annual plan in the Google Cloud Marketplace.
  - b. In Google Cloud, select the annual plan that was shared with your account and then click **Subscribe**.
  - c. After you return to BlueXP, select a capacity-based package when you reach the charging methods page.



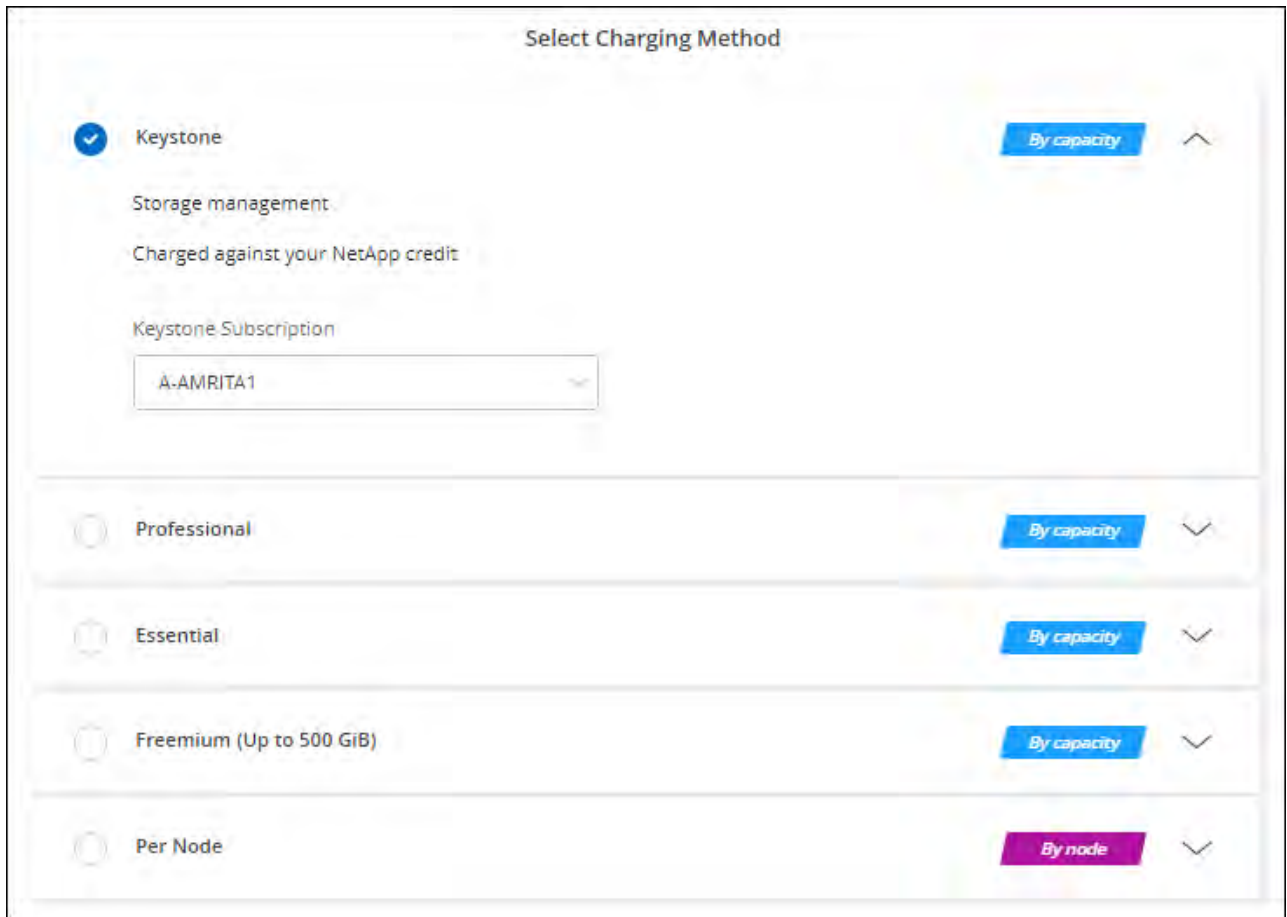
[View step-by-step instructions to launch Cloud Volumes ONTAP in Google Cloud.](#)

### Keystone Subscription

A Keystone Subscription is a pay-as-you-grow subscription-based service. [Learn more about NetApp Keystone Subscriptions.](#)

#### Steps

1. If you don't have a subscription yet, [contact NetApp](#)
2. [Contact NetApp](#) to authorize your BlueXP user account with one or more Keystone Subscriptions.
3. After NetApp authorizes your account, [link your subscriptions for use with Cloud Volumes ONTAP](#).
4. On the Canvas page, click **Add Working Environment** and follow the steps in BlueXP.
  - a. Select the Keystone Subscription charging method when prompted to choose a charging method.



[View step-by-step instructions to launch Cloud Volumes ONTAP in Google Cloud.](#)

## Launch Cloud Volumes ONTAP in Google Cloud

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

### Before you begin

You need the following to create a working environment.

- A Connector that's up and running.
  - You should have a [Connector that is associated with your project or workspace](#).
  - [You should be prepared to leave the Connector running at all times](#).
  - The service account associated with the Connector [should have the required permissions](#)
- An understanding of the configuration that you want to use.

You should have prepared by choosing a configuration and by obtaining Google Cloud networking information from your administrator. For details, refer to [Planning your Cloud Volumes ONTAP configuration](#).

- An understanding of what's required to set up licensing for Cloud Volumes ONTAP.

[Learn how to set up licensing.](#)

- Google Cloud APIs should be [enabled in your project](#):
  - Cloud Deployment Manager V2 API
  - Cloud Logging API
  - Cloud Resource Manager API
  - Compute Engine API
  - Identity and Access Management (IAM) API

## Launch a single-node system in Google Cloud


Create a working environment in BlueXP to launch Cloud Volumes ONTAP in Google Cloud.

### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the prompts.
3. **Choose a Location**: Select **Google Cloud** and **Cloud Volumes ONTAP**.
4. If you're prompted, [create a Connector](#).
5. **Details & Credentials**: Select a project, specify a cluster name, optionally select a service account, optionally add labels, and then specify credentials.

The following table describes fields for which you might need guidance:

Field	Description
Working Environment Name	BlueXP uses the working environment name to name both the Cloud Volumes ONTAP system and the Google Cloud VM instance. It also uses the name as the prefix for the predefined security group, if you select that option.
Service Account Name	If you plan to use <a href="#">data tiering</a> or <a href="#">BlueXP backup and recovery</a> with Cloud Volumes ONTAP, then you need to enable <b>Service Account</b> and select a service account that has the predefined Storage Admin role. <a href="#">Learn how to create a service account</a> .
Add Labels	Labels are metadata for your Google Cloud resources. BlueXP adds the labels to the Cloud Volumes ONTAP system and Google Cloud resources associated with the system.  You can add up to four labels from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four labels when creating a working environment.  For information about labels, refer to <a href="#">Google Cloud Documentation: Labeling Resources</a> .
User name and password	These are the credentials for the Cloud Volumes ONTAP cluster administrator account. You can use these credentials to connect to Cloud Volumes ONTAP through ONTAP System Manager or the ONTAP CLI. Keep the default <i>admin</i> user name or change it to a custom user name.

Field	Description
Edit Project	<p>Select the project where you want Cloud Volumes ONTAP to reside. The default project is the project where BlueXP resides.</p> <p>If you don't see any additional projects in the drop-down list, then you haven't yet associated the BlueXP service account with other projects. Go to the Google Cloud console, open the IAM service, and select the project. Add the service account with the BlueXP role to that project. You'll need to repeat this step for each project.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  This is the service account that you set up for BlueXP, as <a href="#">described on this page</a>. </div> <p>Click <b>Add Subscription</b> to associate the selected credentials with a subscription.</p> <p>To create a pay-as-you-go Cloud Volumes ONTAP system, you need to select a Google Cloud project that's associated with a subscription to Cloud Volumes ONTAP from the Google Cloud marketplace.</p>

The following video shows how to associate a pay-as-you-go marketplace subscription to your Google Cloud project. Alternatively, follow the steps to subscribe located in the [Associating a Marketplace subscription with Google Cloud credentials](#) section.

[Subscribe to BlueXP from the Google Cloud marketplace](#)

- Services:** Select the services that you want to use on this system. In order to select BlueXP backup and recovery, or to use BlueXP tiering, you must have specified the Service Account in step 3.



If you would like to utilize WORM and data tiering, you must disable BlueXP backup and recovery and deploy a Cloud Volumes ONTAP working environment with version 9.8 or above.

- Location & Connectivity:** Select a location, choose a firewall policy, and confirm network connectivity to Google Cloud storage for data tiering.

The following table describes fields for which you might need guidance:

Field	Description
Connectivity verification	To tier cold data to a Google Cloud Storage bucket, the subnet in which Cloud Volumes ONTAP resides must be configured for Private Google Access. For instructions, refer to <a href="#">Google Cloud Documentation: Configuring Private Google Access</a> .

Field	Description
Generated firewall policy	<p>If you let BlueXP generate the firewall policy for you, you need to choose how you'll allow traffic:</p> <ul style="list-style-type: none"> <li>• If you choose <b>Selected VPC only</b>, the source filter for inbound traffic is the subnet range of the selected VPC and the subnet range of the VPC where the Connector resides. This is the recommended option.</li> <li>• If you choose <b>All VPCs</b>, the source filter for inbound traffic is the 0.0.0.0/0 IP range.</li> </ul>
Use existing firewall policy	<p>If you use an existing firewall policy, ensure that it includes the required rules. xref:./ <a href="#">Learn about firewall rules for Cloud Volumes ONTAP</a>.</p>

- Charging Methods and NSS Account:** Specify which charging option would you like to use with this system, and then specify a NetApp Support Site account.
  - [Learn about licensing options for Cloud Volumes ONTAP](#).
  - [Learn how to set up licensing](#).
- Preconfigured Packages:** Select one of the packages to quickly deploy a Cloud Volumes ONTAP system, or click **Create my own configuration**.

If you choose one of the packages, you only need to specify a volume and then review and approve the configuration.

- Licensing:** Change the Cloud Volumes ONTAP version as needed and select a machine type.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then BlueXP updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.13.1 and 9.13.1 P4 is available. The update does not occur from one release to another—for example, from 9.13 to 9.14.

- Underlying Storage Resources:** Choose settings for the initial aggregate: a disk type and the size for each disk.

The disk type is for the initial volume. You can choose a different disk type for subsequent volumes.

The disk size is for all disks in the initial aggregate and for any additional aggregates that BlueXP creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.

For help choosing a disk type and size, refer to [Size your system in Google Cloud](#).

- Flash Cache, Write Speed & WORM:**

- Enable **Flash Cache**, if desired.



Starting with Cloud Volumes ONTAP 9.13.1, *Flash Cache* is supported on the n2-standard-16, n2-standard-32, n2-standard-48, and n2-standard-64 instance types. You cannot disable Flash Cache after deployment.

- Choose **Normal** or **High** write speed, if desired.

[Learn more about write speed.](#)



High write speed and a higher maximum transmission unit (MTU) of 8,896 bytes are available through the **High** write speed option. In addition, the higher MTU of 8,896 requires the selection of VPC-1, VPC-2 and VPC-3 for the deployment. For more information on VPC-1, VPC-2, and VPC-3, refer to [Rules for VPC-1, VPC-2, and VPC-3](#).

c. Activate write once, read many (WORM) storage, if desired.

WORM can't be enabled if data tiering was enabled for Cloud Volumes ONTAP versions 9.7 and below. Reverting or downgrading to Cloud Volumes ONTAP 9.8 is blocked after enabling WORM and tiering.

[Learn more about WORM storage.](#)

d. If you activate WORM storage, select the retention period.

13. **Data Tiering in Google Cloud Platform:** Choose whether to enable data tiering on the initial aggregate, choose a storage class for the tiered data, and then either select a service account that has the predefined Storage Admin role (required for Cloud Volumes ONTAP 9.7 or later), or select a Google Cloud account (required for Cloud Volumes ONTAP 9.6).

Note the following:

- BlueXP sets the service account on the Cloud Volumes ONTAP instance. This service account provides permissions for data tiering to a Google Cloud Storage bucket. Be sure to add the Connector service account as a user of the tiering service account, otherwise, you can't select it from BlueXP
- For help with adding a Google Cloud account, refer to [Setting up and adding Google Cloud accounts for data tiering with 9.6](#).
- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates, but you'll need to turn off the system and add a service account from the Google Cloud console.

[Learn more about data tiering.](#)

14. **Create Volume:** Enter details for the new volume or click **Skip**.

[Learn about supported client protocols and versions.](#)

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

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, BlueXP enters a value that provides access to all instances in the subnet.

Field	Description
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, BlueXP automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

### Volume Details, Protection & Protocol

#### Details & Protection

Volume Name:  Size (GB):

Snapshot Policy:

Default Policy

#### Protocol

NFS   
 CIFS   
 iSCSI

---

Share name:  Permissions:

Users / Groups:

Valid users and groups separated by a semicolon

15. **CIFS Setup:** If you chose the CIFS protocol, set up a CIFS server.



Field	Description
DNS Primary and Secondary IP Address	<p>The IP addresses of the DNS servers that provide name resolution for the CIFS server.</p> <p>The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.</p> <p>If you're configuring Google Managed Active Directory, AD can be accessed by default with the 169.254.169.254 IP address.</p>
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	<p>The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.</p> <p>To configure Google Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter <b>OU=Computers,OU=Cloud</b> in this field.</p> <p><a href="#">Google Cloud Documentation: Organizational Units in Google Managed Microsoft AD</a></p>
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	<p>Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. For information, refer to the <a href="#">BlueXP automation docs</a> for details.</p> <p>Note that you can configure an NTP server only when creating a CIFS server. It's not configurable after you create the CIFS server.</p>

16. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and change the volume tiering policy, if needed.

For more information, refer to [Choose a volume usage profile](#), [Data tiering overview](#), and [KB: What Inline Storage Efficiency features are supported with CVO?](#)

17. **Review & Approve:** Review and confirm your selections.
- Review details about the configuration.
  - Click **More information** to review details about support and the Google Cloud resources that BlueXP will purchase.
  - Select the **I understand...** check boxes.
  - Click **Go**.

## Result

BlueXP deploys the Cloud Volumes ONTAP system. You can track the progress in the timeline.

If you experience any issues deploying the Cloud Volumes ONTAP system, review the failure message. You can also select the working environment and click **Re-create environment**.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

### After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.
- If you want to apply quotas to volumes, use ONTAP System Manager or the ONTAP CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

## Launch an HA pair in Google Cloud


Create a working environment in BlueXP to launch Cloud Volumes ONTAP in Google Cloud.

### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** and follow the prompts.
3. **Choose a Location:** Select **Google Cloud** and **Cloud Volumes ONTAP HA**.
4. **Details & Credentials:** Select a project, specify a cluster name, optionally select a Service Account, optionally add labels, and then specify credentials.

The following table describes fields for which you might need guidance:

Field	Description
Working Environment Name	BlueXP uses the working environment name to name both the Cloud Volumes ONTAP system and the Google Cloud VM instance. It also uses the name as the prefix for the predefined security group, if you select that option.
Service Account Name	If you plan to use the <a href="#">BlueXP tiering</a> or <a href="#">BlueXP backup and recovery</a> services, you need to enable the <b>Service Account</b> switch and then select the Service Account that has the predefined Storage Admin role.
Add Labels	Labels are metadata for your Google Cloud resources. BlueXP adds the labels to the Cloud Volumes ONTAP system and Google Cloud resources associated with the system.  You can add up to four labels from the user interface when creating a working environment, and then you can add more after its created. Note that the API does not limit you to four labels when creating a working environment.  For information about labels, refer to <a href="#">Google Cloud Documentation: Labeling Resources</a> .
User name and password	These are the credentials for the Cloud Volumes ONTAP cluster administrator account. You can use these credentials to connect to Cloud Volumes ONTAP through ONTAP System Manager or the ONTAP CLI. Keep the default <i>admin</i> user name or change it to a custom user name.

Field	Description
Edit Project	<p>Select the project where you want Cloud Volumes ONTAP to reside. The default project is the project where BlueXP resides.</p> <p>If you don't see any additional projects in the drop-down list, then you haven't yet associated the BlueXP service account with other projects. Go to the Google Cloud console, open the IAM service, and select the project. Add the service account with the BlueXP role to that project. You'll need to repeat this step for each project.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin: 10px 0;">  This is the service account that you set up for BlueXP, <a href="#">as described on this page</a>. </div> <p>Click <b>Add Subscription</b> to associate the selected credentials with a subscription.</p> <p>To create a pay-as-you-go Cloud Volumes ONTAP system, you need to select a Google Cloud project that's associated with a subscription to Cloud Volumes ONTAP from the Google Cloud Marketplace.</p>

The following video shows how to associate a pay-as-you-go marketplace subscription to your Google Cloud project. Alternatively, follow the steps to subscribe located in the [Associating a marketplace subscription with Google Cloud credentials](#) section.

[Subscribe to BlueXP from the Google Cloud marketplace](#)

- Services:** Select the services that you want to use on this system. In order to select BlueXP backup and recovery, or to use BlueXP Tiering, you must have specified the Service Account in step 3.



If you would like to utilize WORM and data tiering, you must disable BlueXP backup and recovery and deploy a Cloud Volumes ONTAP working environment with version 9.8 or above.

- HA Deployment Models:** Choose multiple zones (recommended) or a single zone for the HA configuration. Then select a region and zones.

[Learn more about HA deployment models.](#)

- Connectivity:** Select four different VPCs for the HA configuration, a subnet in each VPC, and then choose a firewall policy.

[Learn more about networking requirements.](#)

The following table describes fields for which you might need guidance:

Field	Description
Generated policy	<p>If you let BlueXP generate the firewall policy for you, you need to choose how you'll allow traffic:</p> <ul style="list-style-type: none"> <li>• If you choose <b>Selected VPC only</b>, the source filter for inbound traffic is the subnet range of the selected VPC and the subnet range of the VPC where the Connector resides. This is the recommended option.</li> <li>• If you choose <b>All VPCs</b>, the source filter for inbound traffic is the 0.0.0.0/0 IP range.</li> </ul>
Use existing	<p>If you use an existing firewall policy, ensure that it includes the required rules. <a href="#">Learn about firewall rules for Cloud Volumes ONTAP.</a></p>

- Charging Methods and NSS Account:** Specify which charging option would you like to use with this system, and then specify a NetApp Support Site account.
  - [Learn about licensing options for Cloud Volumes ONTAP.](#)
  - [Learn how to set up licensing.](#)
- Preconfigured Packages:** Select one of the packages to quickly deploy a Cloud Volumes ONTAP system, or click **Create my own configuration**.

If you choose one of the packages, you only need to specify a volume and then review and approve the configuration.

- Licensing:** Change the Cloud Volumes ONTAP version as needed and select a machine type.



If a newer Release Candidate, General Availability, or patch release is available for the selected version, then BlueXP updates the system to that version when creating the working environment. For example, the update occurs if you select Cloud Volumes ONTAP 9.13.1 and 9.13.1 P4 is available. The update does not occur from one release to another—for example, from 9.13 to 9.14.

- Underlying Storage Resources:** Choose settings for the initial aggregate: a disk type and the size for each disk.

The disk type is for the initial volume. You can choose a different disk type for subsequent volumes.

The disk size is for all disks in the initial aggregate and for any additional aggregates that BlueXP creates when you use the simple provisioning option. You can create aggregates that use a different disk size by using the advanced allocation option.

For help choosing a disk type and size, refer to [Size your system in Google Cloud.](#)

- Flash Cache, Write Speed & WORM:**

- Enable **Flash Cache**, if desired.



Starting with Cloud Volumes ONTAP 9.13.1, *Flash Cache* is supported on the n2-standard-16, n2-standard-32, n2-standard-48, and n2-standard-64 instance types. You cannot disable Flash Cache after deployment.

- Choose **Normal** or **High** write speed, if desired.

[Learn more about write speed.](#)



High write speed and a higher maximum transmission unit (MTU) of 8,896 bytes are available through the **High** write speed option with the n2-standard-16, n2-standard-32, n2-standard-48, and n2-standard-64 instance types. In addition, the higher MTU of 8,896 requires the selection of VPC-1, VPC-2 and VPC-3 for the deployment. High write speed and an MTU of 8,896 are feature-dependent and cannot be disabled individually within a configured instance. For more information on VPC-1, VPC-2, and VPC-3, refer to [Rules for VPC-1, VPC-2, and VPC-3](#).

c. Activate write once, read many (WORM) storage, if desired.

WORM can't be enabled if data tiering was enabled for Cloud Volumes ONTAP versions 9.7 and below. Reverting or downgrading to Cloud Volumes ONTAP 9.8 is blocked after enabling WORM and tiering.

[Learn more about WORM storage.](#)

d. If you activate WORM storage, select the retention period.

13. **Data Tiering in Google Cloud:** Choose whether to enable data tiering on the initial aggregate, choose a storage class for the tiered data, and then select a service account that has the predefined Storage Admin role.

Note the following:

- BlueXP sets the service account on the Cloud Volumes ONTAP instance. This service account provides permissions for data tiering to a Google Cloud Storage bucket. Be sure to add the Connector service account as a user of the tiering service account, otherwise, you can't select it from BlueXP.
- You can choose a specific volume tiering policy when you create or edit a volume.
- If you disable data tiering, you can enable it on subsequent aggregates, but you'll need to turn off the system and add a service account from the Google Cloud console.

[Learn more about data tiering.](#)

14. **Create Volume:** Enter details for the new volume or click **Skip**.

[Learn about supported client protocols and versions.](#)

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

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Access control (for NFS only)	An export policy defines the clients in the subnet that can access the volume. By default, BlueXP enters a value that provides access to all instances in the subnet.

Field	Description
Permissions and Users / Groups (for CIFS only)	These fields enable you to control the level of access to a share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.
Advanced options (for NFS only)	Select an NFS version for the volume: either NFSv3 or NFSv4.
Initiator group and IQN (for iSCSI only)	<p>iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices.</p> <p>Initiator groups are tables of iSCSI host node names and control which initiators have access to which LUNs.</p> <p>iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).</p> <p>When you create an iSCSI volume, BlueXP automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, <a href="#">use the IQN to connect to the LUN from your hosts</a>.</p>

The following image shows the Volume page filled out for the CIFS protocol:

### Volume Details, Protection & Protocol

#### Details & Protection

Volume Name:  Size (GB):

Snapshot Policy:

Default Policy

#### Protocol

NFS   
 CIFS   
 iSCSI

Share name:  Permissions:

Users / Groups:

Valid users and groups separated by a semicolon

15. **CIFS Setup:** If you chose the CIFS protocol, set up a CIFS server.

Field	Description
DNS Primary and Secondary IP Address	<p>The IP addresses of the DNS servers that provide name resolution for the CIFS server.</p> <p>The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.</p> <p>If you're configuring Google Managed Active Directory, AD can be accessed by default with the 169.254.169.254 IP address.</p>
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	<p>The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.</p> <p>To configure Google Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter <b>OU=Computers,OU=Cloud</b> in this field.</p> <p><a href="#">Google Cloud Documentation: Organizational Units in Google Managed Microsoft AD</a></p>
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
NTP Server	<p>Select <b>Use Active Directory Domain</b> to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. Refer to the <a href="#">BlueXP automation docs</a> for details.</p> <p>Note that you can configure an NTP server only when creating a CIFS server. It's not configurable after you create the CIFS server.</p>

16. **Usage Profile, Disk Type, and Tiering Policy:** Choose whether you want to enable storage efficiency features and change the volume tiering policy, if needed.

For more information, refer to [Choose a volume usage profile](#), [Data tiering overview](#), and [KB: What Inline Storage Efficiency features are supported with CVO?](#)

17. **Review & Approve:** Review and confirm your selections.
- Review details about the configuration.
  - Click **More information** to review details about support and the Google Cloud resources that BlueXP will purchase.
  - Select the **I understand...** check boxes.
  - Click **Go**.

## Result

BlueXP deploys the Cloud Volumes ONTAP system. You can track the progress in the timeline.

If you experience any issues deploying the Cloud Volumes ONTAP system, review the failure message. You can also select the working environment and click **Re-create environment**.

For additional help, go to [NetApp Cloud Volumes ONTAP Support](#).

#### After you finish

- If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.
- If you want to apply quotas to volumes, use ONTAP System Manager or the ONTAP CLI.

Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

## Google Cloud Platform Image Verification

### Learn how Google Cloud image is verified in Cloud Volumes ONTAP

Google Cloud image verification complies with enhanced NetApp security requirements. Changes have been made to the script generating the images to sign the image along the way using private keys specifically generated for this task. You can verify the integrity of the Google Cloud image by using the signed digest and public certificate for Google Cloud which can be downloaded via [NSS](#) for a specific release.



Google Cloud image verification is supported on Cloud Volumes ONTAP software version 9.13.0 or greater.

### Convert Google Cloud image to raw format for Cloud Volumes ONTAP

The image being used to deploy new instances, upgrades, or being used in existing images will be shared with the clients through [the NetApp Support Site \(NSS\)](#). The signed digest, and the certificates will be available to download through the NSS portal. Make sure you are downloading the digest and certificates for the right release corresponding to the image shared by NetApp Support. For instance, 9.13.0 images will have a 9.13.0 signed digest and certificates available on NSS.

#### Why is this step needed?

The images from Google Cloud cannot be downloaded directly. In order to verify the image against the signed digest and the certificates, you need to have a mechanism to compare the two files and download the image. To do so, you must export/convert the image into a disk.raw format and save the results in a storage bucket in Google Cloud. The disk.raw file is tarred and gzipped in the process.

The user/service-account will need privileges to perform the following:

- Access to Google storage bucket
- Write to Google Storage bucket
- Create cloud build jobs (used during export process)
- Access to the desired image
- Create export image tasks



To verify the image, it must be converted to a disk.raw format and then downloaded.

#### **Use Google Cloud command line to export Google Cloud image**

The preferred way to export an image to Cloud Storage is to use the [gcloud compute images export command](#). This command takes the provided image and converts it to a disk.raw file which gets tarred and gzipped. The generated file is saved at the destination URL and can then be downloaded for verification.

The user/account must have privileges to access and write to the desired bucket, export the image, and cloud builds (used by Google to export the image) to execute this operation.

#### **Export Google Cloud image using gcloud**

## Click to display

```
$ gcloud compute images export \  
  --destination-uri DESTINATION_URI \  
  --image IMAGE_NAME  
  
# For our example:  
$ gcloud compute images export \  
  --destination-uri gs://vsa-dev-bucket1/example-user-exportimage-  
gcp-demo \  
  --image example-user-20230120115139  
  
## DEMO ##  
# Step 1 - Optional: Checking access and listing objects in the  
destination bucket  
$ gsutil ls gs://example-user-export-image-bucket/  
  
# Step 2 - Exporting the desired image to the bucket  
$ gcloud compute images export --image example-user-export-image-demo  
--destination-uri gs://example-user-export-image-bucket/export-  
demo.tar.gz  
Created [https://cloudbuild.googleapis.com/v1/projects/example-demo-  
project/locations/us-central1/builds/xxxxxxxxxxxxx].  
Logs are available at [https://console.cloud.google.com/cloud-  
build/builds;region=us-central1/xxxxxxxxxxxxx?project=xxxxxxxxxxxxx].  
[image-export]: 2023-01-25T18:13:48Z Fetching image "example-user-  
export-image-demo" from project "example-demo-project".  
[image-export]: 2023-01-25T18:13:49Z Validating workflow  
[image-export]: 2023-01-25T18:13:49Z Validating step "setup-disks"  
[image-export]: 2023-01-25T18:13:49Z Validating step "image-export-  
export-disk"  
[image-export.image-export-export-disk]: 2023-01-25T18:13:49Z  
Validating step "setup-disks"  
[image-export.image-export-export-disk]: 2023-01-25T18:13:49Z  
Validating step "run-image-export-export-disk"  
[image-export.image-export-export-disk]: 2023-01-25T18:13:50Z  
Validating step "wait-for-inst-image-export-export-disk"  
[image-export.image-export-export-disk]: 2023-01-25T18:13:50Z  
Validating step "copy-image-object"  
[image-export.image-export-export-disk]: 2023-01-25T18:13:50Z  
Validating step "delete-inst"  
[image-export]: 2023-01-25T18:13:51Z Validation Complete  
[image-export]: 2023-01-25T18:13:51Z Workflow Project: example-demo-  
project  
[image-export]: 2023-01-25T18:13:51Z Workflow Zone: us-central1-c
```

```
[image-export]: 2023-01-25T18:13:51Z Workflow GCSPath: gs://example-
demo-project-example-bkt-us/
[image-export]: 2023-01-25T18:13:51Z Example scratch path:
https://console.cloud.google.com/storage/browser/example-demo-project-
example-bkt-us/example-image-export-20230125-18:13:49-r88px
[image-export]: 2023-01-25T18:13:51Z Uploading sources
[image-export]: 2023-01-25T18:13:51Z Running workflow
[image-export]: 2023-01-25T18:13:51Z Running step "setup-disks"
(CreateDisks)
[image-export.setup-disks]: 2023-01-25T18:13:51Z CreateDisks: Creating
disk "disk-image-export-image-export-r88px".
[image-export]: 2023-01-25T18:14:02Z Step "setup-disks" (CreateDisks)
successfully finished.
[image-export]: 2023-01-25T18:14:02Z Running step "image-export-export-
disk" (IncludeWorkflow)
[image-export.image-export-export-disk]: 2023-01-25T18:14:02Z Running
step "setup-disks" (CreateDisks)
[image-export.image-export-export-disk.setup-disks]: 2023-01-
25T18:14:02Z CreateDisks: Creating disk "disk-image-export-export-disk-
image-export-image-export--r88px".
[image-export.image-export-export-disk]: 2023-01-25T18:14:02Z Step
"setup-disks" (CreateDisks) successfully finished.
[image-export.image-export-export-disk]: 2023-01-25T18:14:02Z Running
step "run-image-export-export-disk" (CreateInstances)
[image-export.image-export-export-disk.run-image-export-export-disk]:
2023-01-25T18:14:02Z CreateInstances: Creating instance "inst-image-
export-export-disk-image-export-image-export--r88px".
[image-export.image-export-export-disk]: 2023-01-25T18:14:08Z Step
"run-image-export-export-disk" (CreateInstances) successfully finished.
[image-export.image-export-export-disk.run-image-export-export-disk]:
2023-01-25T18:14:08Z CreateInstances: Streaming instance "inst-image-
export-export-disk-image-export-image-export--r88px" serial port 1
output to https://storage.cloud.google.com/example-demo-project-
example-bkt-us/example-image-export-20230125-18:13:49-r88px/logs/inst-
image-export-export-disk-image-export-image-export--r88px-serial-
port1.log
[image-export.image-export-export-disk]: 2023-01-25T18:14:08Z Running
step "wait-for-inst-image-export-export-disk" (WaitForInstancesSignal)
[image-export.image-export-export-disk.wait-for-inst-image-export-
export-disk]: 2023-01-25T18:14:08Z WaitForInstancesSignal: Instance
"inst-image-export-export-disk-image-export-image-export--r88px":
watching serial port 1, SuccessMatch: "ExportSuccess", FailureMatch:
["ExportFailed:"] (this is not an error), StatusMatch: "GCEExport:".
[image-export.image-export-export-disk.wait-for-inst-image-export-
export-disk]: 2023-01-25T18:14:29Z WaitForInstancesSignal: Instance
"inst-image-export-export-disk-image-export-image-export--r88px":
```

```
StatusMatch found: "GCEExport: <serial-output key:'source-size-gb'  
value:'10'>"  
[image-export.image-export-export-disk.wait-for-inst-image-export-  
export-disk]: 2023-01-25T18:14:29Z WaitForInstancesSignal: Instance  
"inst-image-export-export-disk-image-export-image-export--r88px":  
StatusMatch found: "GCEExport: Running export tool."  
[image-export.image-export-export-disk.wait-for-inst-image-export-  
export-disk]: 2023-01-25T18:14:29Z WaitForInstancesSignal: Instance  
"inst-image-export-export-disk-image-export-image-export--r88px":  
StatusMatch found: "GCEExport: Disk /dev/sdb is 10 GiB, compressed size  
will most likely be much smaller."  
[image-export.image-export-export-disk.wait-for-inst-image-export-  
export-disk]: 2023-01-25T18:14:29Z WaitForInstancesSignal: Instance  
"inst-image-export-export-disk-image-export-image-export--r88px":  
StatusMatch found: "GCEExport: Beginning export process..."  
[image-export.image-export-export-disk.wait-for-inst-image-export-  
export-disk]: 2023-01-25T18:14:29Z WaitForInstancesSignal: Instance  
"inst-image-export-export-disk-image-export-image-export--r88px":  
StatusMatch found: "GCEExport: Copying \"/dev/sdb\" to gs://example-  
demo-project-example-bkt-us/example-image-export-20230125-18:13:49-  
r88px/outs/image-export-export-disk.tar.gz."  
[image-export.image-export-export-disk.wait-for-inst-image-export-  
export-disk]: 2023-01-25T18:14:29Z WaitForInstancesSignal: Instance  
"inst-image-export-export-disk-image-export-image-export--r88px":  
StatusMatch found: "GCEExport: Using \"/root/upload\" as the buffer  
prefix, 1.0 GiB as the buffer size, and 4 as the number of workers."  
[image-export.image-export-export-disk.wait-for-inst-image-export-  
export-disk]: 2023-01-25T18:14:29Z WaitForInstancesSignal: Instance  
"inst-image-export-export-disk-image-export-image-export--r88px":  
StatusMatch found: "GCEExport: Creating gzipped image of \"/dev/sdb\"." "  
[image-export.image-export-export-disk.wait-for-inst-image-export-  
export-disk]: 2023-01-25T18:14:29Z WaitForInstancesSignal: Instance  
"inst-image-export-export-disk-image-export-image-export--r88px":  
StatusMatch found: "GCEExport: Read 1.0 GiB of 10 GiB (212 MiB/sec),  
total written size: 992 MiB (198 MiB/sec)"  
[image-export.image-export-export-disk.wait-for-inst-image-export-  
export-disk]: 2023-01-25T18:14:59Z WaitForInstancesSignal: Instance  
"inst-image-export-export-disk-image-export-image-export--r88px":  
StatusMatch found: "GCEExport: Read 8.0 GiB of 10 GiB (237 MiB/sec),  
total written size: 1.5 GiB (17 MiB/sec)"  
[image-export.image-export-export-disk.wait-for-inst-image-export-  
export-disk]: 2023-01-25T18:15:19Z WaitForInstancesSignal: Instance  
"inst-image-export-export-disk-image-export-image-export--r88px":  
StatusMatch found: "GCEExport: Finished creating gzipped image of  
\"/dev/sdb\" in 48.956433327s [213 MiB/s] with a compression ratio of  
6."
```

```

[image-export.image-export-export-disk.wait-for-inst-image-export-export-disk]: 2023-01-25T18:15:19Z WaitForInstancesSignal: Instance "inst-image-export-export-disk-image-export-image-export--r88px": StatusMatch found: "GCEExport: Finished export in 48.957347731s"
[image-export.image-export-export-disk.wait-for-inst-image-export-export-disk]: 2023-01-25T18:15:19Z WaitForInstancesSignal: Instance "inst-image-export-export-disk-image-export-image-export--r88px": StatusMatch found: "GCEExport: <serial-output key:'target-size-gb' value:'2'>"
[image-export.image-export-export-disk.wait-for-inst-image-export-export-disk]: 2023-01-25T18:15:19Z WaitForInstancesSignal: Instance "inst-image-export-export-disk-image-export-image-export--r88px": SuccessMatch found "ExportSuccess"
[image-export.image-export-export-disk]: 2023-01-25T18:15:19Z Step "wait-for-inst-image-export-export-disk" (WaitForInstancesSignal) successfully finished.
[image-export.image-export-export-disk]: 2023-01-25T18:15:19Z Running step "copy-image-object" (CopyGCSObjects)
[image-export.image-export-export-disk]: 2023-01-25T18:15:19Z Running step "delete-inst" (DeleteResources)
[image-export.image-export-export-disk.delete-inst]: 2023-01-25T18:15:19Z DeleteResources: Deleting instance "inst-image-export-export-disk".
[image-export.image-export-export-disk]: 2023-01-25T18:15:19Z Step "copy-image-object" (CopyGCSObjects) successfully finished.
[image-export.image-export-export-disk]: 2023-01-25T18:15:34Z Step "delete-inst" (DeleteResources) successfully finished.
[image-export]: 2023-01-25T18:15:34Z Step "image-export-export-disk" (IncludeWorkflow) successfully finished.
[image-export]: 2023-01-25T18:15:34Z Serial-output value -> source-size-gb:10
[image-export]: 2023-01-25T18:15:34Z Serial-output value -> target-size-gb:2
[image-export]: 2023-01-25T18:15:34Z Workflow "image-export" cleaning up (this may take up to 2 minutes).
[image-export]: 2023-01-25T18:15:35Z Workflow "image-export" finished cleanup.

# Step 3 - Validating the image was successfully exported
$ gsutil ls gs://example-user-export-image-bucket/
gs://example-user-export-image-bucket/export-demo.tar.gz

# Step 4 - Download the exported image
$ gcloud storage cp gs://BUCKET_NAME/OBJECT_NAME SAVE_TO_LOCATION

```

```
$ gcloud storage cp gs://example-user-export-image-bucket/export-  
demo.tar.gz CVO_GCP_Signed_Digest.tar.gz  
Copying gs://example-user-export-image-bucket/export-demo.tar.gz to  
file://CVO_GCP_Signed_Digest.tar.gz  
Completed files 1/1 | 1.5GiB/1.5GiB | 185.0MiB/s
```

```
Average throughput: 213.3MiB/s
```

```
$ ls -l  
total 1565036  
-rw-r--r-- 1 example-user example-user 1602589949 Jan 25 18:44  
CVO_GCP_Signed_Digest.tar.gz
```

## Extract zipped files

```
# Extracting files from the digest  
$ tar -xf CVO_GCP_Signed_Digest.tar.gz
```



For more information on how to export an image through Google Cloud, refer to [Google Cloud doc on Exporting an image](#).

## Image signature verification

### Google Cloud image signature verification for Cloud Volumes ONTAP

To verify the exported Google Cloud signed image, you must download the image digest file from the NSS to validate the disk.raw file and digest file contents.

### Signed image verification workflow summary

The following is an overview of the Google Cloud signed image verification workflow process.

- From the [NSS](#), download the Google Cloud archive containing the following files:
  - Signed digest (.sig)
  - Certificate containing the public key (.pem)
  - Certificate chain (.pem)

# Cloud Volumes ONTAP 9.15.0P1

Date Posted : 17-May-2024

Cloud Volumes ONTAP

### Non-Restricted Countries

If you are upgrading to ONTAP 9.15.0P1, and you are in "Non-restricted Countries", please download the image with NetApp Volume Encryption.

[DOWNLOAD 9150P1\\_V\\_IMAGE.TGZ \[2.58 GB\]](#)  
View and download checksums

[DOWNLOAD 9150P1\\_V\\_IMAGE.TGZ.PEM \[451 B\]](#)  
View and download checksums

[DOWNLOAD 9150P1\\_V\\_IMAGE.TGZ.SIG \[256 B\]](#)  
View and download checksums

Cloud Volumes ONTAP

### Restricted Countries

If you are unsure whether your company complied with all applicable legal requirements on encryption technology, download the image without NetApp Volume Encryption.

[DOWNLOAD 9150P1\\_V\\_NODAR\\_IMAGE.TGZ \[2.58 GB\]](#)  
View and download checksums

[DOWNLOAD 9150P1\\_V\\_NODAR\\_IMAGE.TGZ.PEM \[451 B\]](#)  
View and download checksums

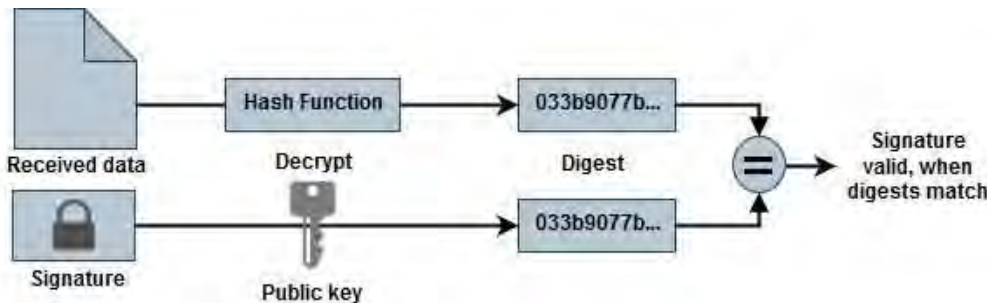
[DOWNLOAD 9150P1\\_V\\_NODAR\\_IMAGE.TGZ.SIG \[256 B\]](#)  
View and download checksums

Cloud Volumes ONTAP

[DOWNLOAD GCP-9-15-0P1\\_PKG.TAR.GZ \[7.49 KB\]](#)  
View and download checksums

[DOWNLOAD AZURE-9-15-0P1\\_PKG.TAR.GZ \[7.64 KB\]](#)  
View and download checksums

- Download the converted disk.raw file
- Validate the certificate using the certificate chain
- Validate the signed digest using the certificate contain the public key
  - Decrypt the signed digest using the public key to extract the digest of the image file
  - Create a digest of the downloaded disk.raw file
  - Compare the two digest file for validation



Verify the Google Cloud image disk.raw file for Cloud Volumes ONTAP using OpenSSL

You can verify the Google Cloud downloaded disk.raw file against the digest file contents available through the [NSS](#) using OpenSSL.



The OpenSSL commands to validate the image are compatible with Linux, macOS, and Windows machines.

## Steps

1. Verify the certificate using OpenSSL.

## Click to display

```
# Step 1 - Optional, but recommended: Verify the certificate using
OpenSSL

# Step 1.1 - Copy the Certificate and certificate chain to a
directory
$ openssl version
LibreSSL 3.3.6
$ ls -l
total 48
-rw-r--r--@ 1 example-user  engr  8537 Jan 19 15:42 Certificate-
Chain-GCP-CVO-20230119-0XXXXX.pem
-rw-r--r--@ 1 example-user  engr  2365 Jan 19 15:42 Certificate-GCP-
CVO-20230119-0XXXXX.pem

# Step 1.2 - Get the OSCP URL
$ oscp_url=$(openssl x509 -noout -ocsp_uri -in <Certificate-
Chain.pem>)
$ oscp_url=$(openssl x509 -noout -ocsp_uri -in Certificate-Chain-
GCP-CVO-20230119-0XXXXX.pem)
$ echo $oscp_url
http://ocsp.entrust.net

# Step 1.3 - Generate an OSCP request for the certificate
$ openssl ocsf -issuer <Certificate-Chain.pem> -CAfile <Certificate-
Chain.pem> -cert <Certificate.pem> -reqout <request.der>
$ openssl ocsf -issuer Certificate-Chain-GCP-CVO-20230119-0XXXXX.pem
-CAfile Certificate-Chain-GCP-CVO-20230119-0XXXXX.pem -cert
Certificate-GCP-CVO-20230119-0XXXXX.pem -reqout req.der

# Step 1.4 - Optional: Check the new file "req.der" has been
generated
$ ls -l
total 56
-rw-r--r--@ 1 example-user  engr  8537 Jan 19 15:42 Certificate-
Chain-GCP-CVO-20230119-0XXXXX.pem
-rw-r--r--@ 1 example-user  engr  2365 Jan 19 15:42 Certificate-GCP-
CVO-20230119-0XXXXX.pem
-rw-r--r--  1 example-user  engr   120 Jan 19 16:50 req.der

# Step 1.5 - Connect to the OSCP Manager using openssl to send the
OCSP request
$ openssl ocsf -issuer <Certificate-Chain.pem> -CAfile <Certificate-
Chain.pem> -cert <Certificate.pem> -url ${ocsp_url} -resp_text
-respout <response.der>
```



```
$ openssl ocspl -issuer Certificate-Chain-GCP-CVO-20230119-0XXXXX.pem
-CAfile Certificate-Chain-GCP-CVO-20230119-0XXXXX.pem -cert
Certificate-GCP-CVO-20230119-0XXXXX.pem -url ${ocsp_url} -resp_text
-respout resp.der
```

OCSP Response Data:

OCSP Response Status: successful (0x0)

Response Type: Basic OCSP Response

Version: 1 (0x0)

Responder Id: C = US, O = "Entrust, Inc.", CN = Entrust Extended  
Validation Code Signing CA - EVCS2

Produced At: Jan 19 15:14:00 2023 GMT

Responses:

Certificate ID:

Hash Algorithm: sha1

Issuer Name Hash: 69FA640329AB84E27220FE0927647B8194B91F2A

Issuer Key Hash: CE894F8251AA15A28462CA312361D261F8FE78

Serial Number: 5994B3D01D26D594BD1D0FA7098C6FF5

Cert Status: good

This Update: Jan 19 15:00:00 2023 GMT

Next Update: Jan 26 14:59:59 2023 GMT

Signature Algorithm: sha512WithRSAEncryption

0b:b6:61:e4:03:5f:98:6f:10:1c:9a:f7:5f:6f:c7:e3:f4:72:  
f2:30:f4:86:88:9a:b9:ba:1e:d6:f6:47:af:dc:ea:e4:cd:31:  
af:e3:7a:20:35:9e:60:db:28:9c:7f:2e:17:7b:a5:11:40:4f:  
1e:72:f7:f8:ef:e3:23:43:1b:bb:28:1a:6f:c6:9c:c5:0c:14:  
d3:5d:bd:9b:6b:28:fb:94:5e:8a:ef:40:20:72:a4:41:df:55:  
cf:f3:db:1b:39:e0:30:63:c9:c7:1f:38:7e:7f:ec:f4:25:7b:  
1e:95:4c:70:6c:83:17:c3:db:b2:47:e1:38:53:ee:0a:55:c0:  
15:6a:82:20:b2:ea:59:eb:9c:ea:7e:97:aa:50:d7:bc:28:60:  
8c:d4:21:92:1c:13:19:b4:e0:66:cb:59:ed:2e:f8:dc:7b:49:  
e3:40:f2:b6:dc:d7:2d:2e:dd:21:82:07:bb:3a:55:99:f7:59:  
5d:4a:4d:ca:e7:8f:1c:d3:9a:3f:17:7b:7a:c4:57:b2:57:a8:  
b4:c0:a5:02:bd:59:9c:50:32:ff:16:b1:65:3a:9c:8c:70:3b:  
9e:be:bc:4f:f9:86:97:b1:62:3c:b2:a9:46:08:be:6b:1b:3c:  
24:14:59:28:c6:ae:e8:d5:64:b2:f8:cc:28:24:5c:b2:c8:d8:  
5a:af:9d:55:48:96:f6:3e:c6:bf:a6:0c:a4:c0:ab:d6:57:03:  
2b:72:43:b0:6a:9f:52:ef:43:bb:14:6a:ce:66:cc:6c:4e:66:  
17:20:a3:64:e0:c6:d1:82:0a:d7:41:8a:cc:17:fd:21:b5:c6:  
d2:3a:af:55:2e:2a:b8:c7:21:41:69:e1:44:ab:a1:dd:df:6d:  
15:99:90:cc:a0:74:1e:e5:2e:07:3f:50:e6:72:a6:b9:ae:fc:  
44:15:eb:81:3d:1a:f8:17:b6:0b:ff:05:76:9d:30:06:40:72:  
cf:d5:c4:6f:8b:c9:14:76:09:6b:3d:6a:70:2c:5a:c4:51:92:  
e5:cd:84:b6:f9:d9:d5:bc:8d:72:b7:7c:13:9c:41:89:a8:97:  
6f:4a:11:5f:8f:b6:c9:b5:df:00:7e:97:20:e7:29:2e:2b:12:  
77:dc:e2:63:48:87:42:49:1d:fc:d0:94:a8:8d:18:f9:07:85:

```

e4:d0:3e:9a:4a:d7:d5:d0:02:51:c3:51:1c:73:12:96:2d:75:
22:83:a6:70:5a:4a:2b:f2:98:d9:ae:1b:57:53:3d:3b:58:82:
38:fc:fa:cb:57:43:3f:3e:7e:e0:6d:5b:d6:fc:67:7e:07:7e:
fb:a3:76:43:26:8f:d1:42:d6:a6:33:4e:9e:e0:a0:51:b4:c4:
bc:e3:10:0d:bf:23:6c:4b
WARNING: no nonce in response
Response Verify OK
Certificate-GCP-CVO-20230119-0XXXXX.pem: good
  This Update: Jan 19 15:00:00 2023 GMT
  Next Update: Jan 26 14:59:59 2023 GMT

# Step 1.5 - Optional: Check the response file "response.der" has
been generated. Verify its contents.
$ ls -l
total 64
-rw-r--r--@ 1 example-user  engr  8537 Jan 19 15:42 Certificate-
Chain-GCP-CVO-20230119-0XXXXX.pem
-rw-r--r--@ 1 example-user  engr  2365 Jan 19 15:42 Certificate-GCP-
CVO-20230119-0XXXXX.pem
-rw-r--r--  1 example-user  engr   120 Jan 19 16:50 req.der
-rw-r--r--  1 example-user  engr   806 Jan 19 16:51 resp.der

# Step 1.6 - Verify the chain of trust and expiration dates against
the local host
$ openssl version -d
OPENSSLDIR: "/private/etc/ssl"
$ OPENSSLDIR=$(openssl version -d | cut -d '"' -f2)
$ echo $OPENSSLDIR
/private/etc/ssl

$ openssl verify -untrusted <Certificate-Chain.pem> -CApath <OpenSSL
dir> <Certificate.pem>
$ openssl verify -untrusted Certificate-Chain-GCP-CVO-20230119-
0XXXXX.pem -CApath ${OPENSSLDIR} Certificate-GCP-CVO-20230119-
0XXXXX.pem
Certificate-GCP-CVO-20230119-0XXXXX.pem: OK

```

2. Place the downloaded disk.raw file, the signature, and certificates in a directory.
3. Extract the public key from the certificate using OpenSSL.
4. Decrypt the signature using the extracted public key and verify the contents of the downloaded disk.raw file.

## Click to display

```
# Step 1 - Place the downloaded disk.raw, the signature and the
certificates in a directory
$ ls -l
-rw-r--r--@ 1 example-user  staff  Jan 19 15:42 Certificate-Chain-
GCP-CVO-20230119-0XXXXX.pem
-rw-r--r--@ 1 example-user  staff  Jan 19 15:42 Certificate-GCP-CVO-
20230119-0XXXXX.pem
-rw-r--r--@ 1 example-user  staff  Jan 19 15:42 GCP_CVO_20230119-
XXXXXX_digest.sig
-rw-r--r--@ 1 example-user  staff  Jan 19 16:39 disk.raw

# Step 2 - Extract the public key from the certificate
$ openssl x509 -pubkey -noout -in (certificate.pem) >
(public_key.pem)
$ openssl x509 -pubkey -noout -in Certificate-GCP-CVO-20230119-
0XXXXX.pem > CVO-GCP-pubkey.pem

$ ls -l
-rw-r--r--@ 1 example-user  staff  Jan 19 15:42 Certificate-Chain-
GCP-CVO-20230119-0XXXXX.pem
-rw-r--r--@ 1 example-user  staff  Jan 19 15:42 Certificate-GCP-CVO-
20230119-0XXXXX.pem
-rw-r--r--@ 1 example-user  staff  Jan 19 17:02 CVO-GCP-pubkey.pem
-rw-r--r--@ 1 example-user  staff  Jan 19 15:42 GCP_CVO_20230119-
XXXXXX_digest.sig
-rw-r--r--@ 1 example-user  staff  Jan 19 16:39 disk.raw

# Step 3 - Decrypt the signature using the extracted public key and
verify the contents of the downloaded disk.raw
$ openssl dgst -verify (public_key) -keyform PEM -sha256 -signature
(signed digest) -binary (downloaded or obtained disk.raw)
$ openssl dgst -verify CVO-GCP-pubkey.pem -keyform PEM -sha256
-signature GCP_CVO_20230119-XXXXXX_digest.sig -binary disk.raw
Verified OK

# A failed response would look like this
$ openssl dgst -verify CVO-GCP-pubkey.pem -keyform PEM -sha256
-signature GCP_CVO_20230119-XXXXXX_digest.sig -binary
../sample_file.txt
Verification Failure
```

# Use Cloud Volumes ONTAP

## License management

### Manage capacity-based licensing for Cloud Volumes ONTAP

Manage your capacity-based licenses from the BlueXP digital wallet to ensure that your NetApp account has enough capacity for your Cloud Volumes ONTAP systems.

*Capacity-based licenses* enable you to pay for Cloud Volumes ONTAP per TiB of capacity.

The *BlueXP digital wallet* enables you to manage licenses for Cloud Volumes ONTAP from a single location. You can add new licenses and update existing licenses.



While the actual usage and metering for the products and services managed in BlueXP are always calculated in GiB and TiB, the terms GB/GiB and TB/TiB are used interchangeably. This is reflected in the Cloud Marketplace listings, price quotes, listing descriptions, and in other supporting documentation

[Learn more about Cloud Volumes ONTAP licenses.](#)

### How licenses are added to the BlueXP digital wallet

After you purchase a license from your NetApp sales representative, NetApp will send you an email with the serial number and additional licensing details.

In the meantime, BlueXP automatically queries NetApp's licensing service to obtain details about the licenses associated with your NetApp Support Site account. If there are no errors, BlueXP automatically adds the licenses to the digital wallet.

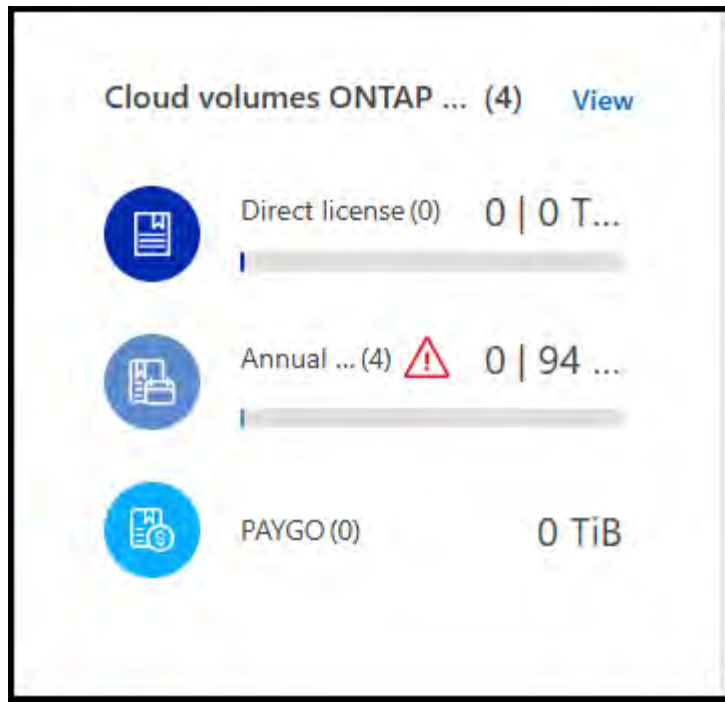
If BlueXP can't add the license, you'll need to manually add them to the digital wallet yourself. For example, if the Connector is installed in a location that doesn't have internet access, you'll need to add the licenses yourself. [Learn how to add purchased licenses to your account.](#)

### View the consumed capacity in your account

The BlueXP digital wallet shows you the total consumed capacity in your account and the consumed capacity by licensing package. This can help you understand how you're being charged and whether you need to purchase additional capacity.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. On the **Overview** tab, the Cloud Volumes ONTAP tile displays the current capacity provisioned for your account.



- *Direct license* is the total provisioned capacity of all Cloud Volumes ONTAP systems in your NetApp account. The charging is based on each volume's provisioned size, regardless of local, used, stored, or effective space within the volume.
- *Annual contract* is the total licensed capacity (bring your own license (BYOL) or Marketplace Contract) that you purchased from NetApp.
- *PAYGO* is the total provisioned capacity using cloud marketplace subscriptions. Charging via PAYGO is used only if the consumed capacity is higher than the licensed capacity or if there is no BYOL license available in the BlueXP digital wallet.

3. Select **View** to see the consumed capacity for each of your licensing packages.
4. Select the **Licenses** tab to see details for each package license that you have purchased.


To better understand the capacities that display for the Essentials package, you should be familiar with how charging works. [Learn about charging for the Essentials package.](#)

5. Select the **Subscriptions** tab to see the consumed capacity by license consumption model. This tab includes both PAYGO and annual contract licenses.

You'll only see the subscriptions that are associated with the organization that you are that you're currently viewing.

6. As you view the information about your subscriptions, you can interact with the details in the table as follows:
  - Expand a row to view more details.

Provider	Name	Type	Service	Start Date	Status	
aws	AWS subscription	PAYGO	NetApp BlueXP	Apr 04, 2024	Subscribed	⋮
NetApp BlueXP			N/A	N/A		
Product Title			Term	Auto Renew		

- Select  to choose which columns appear in the table. Note that the Term and Auto Renew columns don't appear by default. The Auto Renew column displays renewal information for Azure contracts only.

### Viewing package details

You can view details about the capacity used per package by switching to legacy mode on the Cloud Volumes ONTAP page.

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. On the **Overview** tab, the Cloud Volumes ONTAP tile displays the current capacity provisioned for your account.
3. Select **View** to see the provisioned capacity for each of your licensing packages.
4. Select **Switch to advanced view**.

5. View the details of the package you want to see.

## Change charging methods

Capacity-based licensing is available in the form of a *package*. When you create a Cloud Volumes ONTAP working environment, you can choose from several licensing packages based on your business needs. If your needs change after you create the working environment, you can change the package at any time. For example, you might change from the Essentials package to the Professional package.

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

### About this task

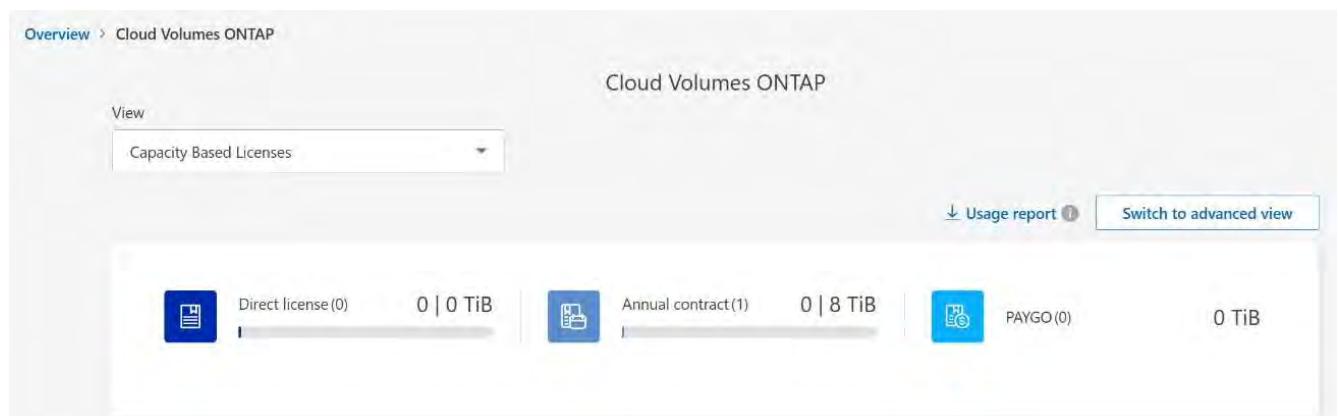
- Changing the charging method doesn't affect whether you're charged through a license purchased from NetApp (BYOL) or from your cloud provider's marketplace pay-as-you-go (PAYGO) subscription.

BlueXP always attempts to charge against a license first. If a license isn't available, it charges against a marketplace subscription. No "conversion" is required for BYOL to marketplace subscription or vice versa.

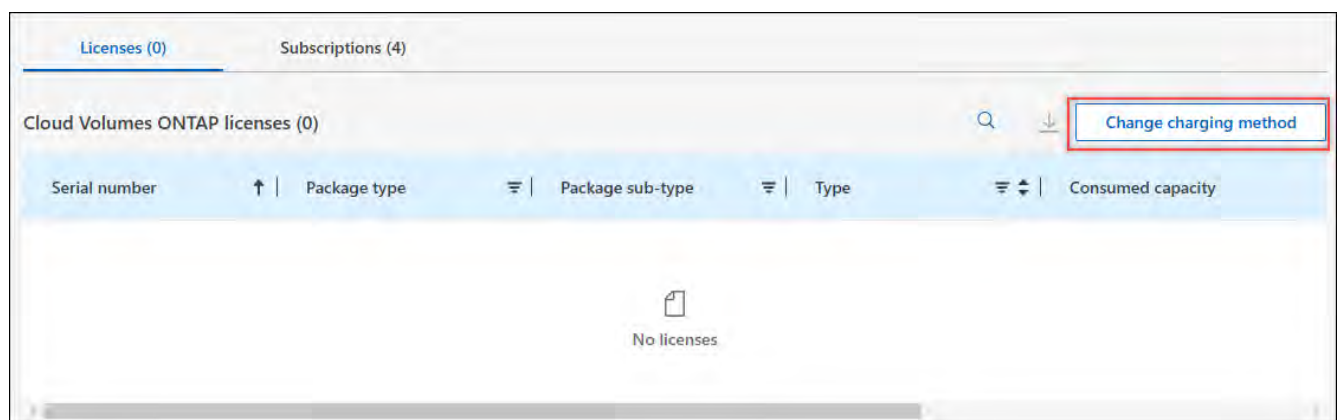
- If you have a private offer or contract from your cloud provider's marketplace, changing to a charging method that's not included in your contract will result in charging against BYOL (if you purchased a license from NetApp) or PAYGO.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Switch to advanced view**.



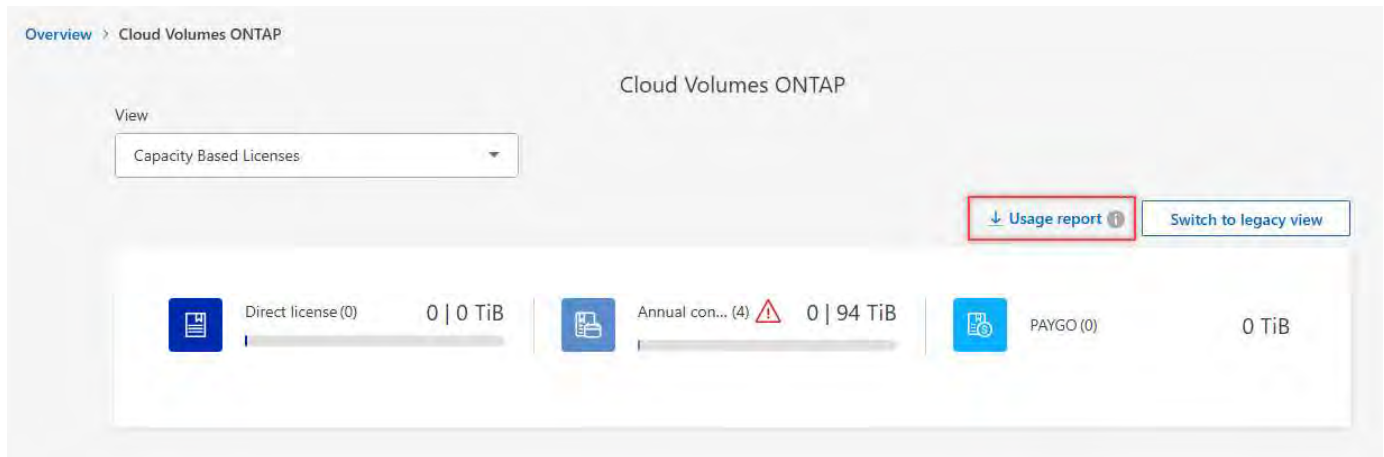
5. Scroll down to the **Capacity-based license** table and select **Change charging method**.



6. On the **Change charging method** pop-up, select a working environment, choose the new charging method, and then confirm your understanding that changing the package type will affect service charges.
7. Select **Change charging method**.

## Download usage reports

You can download four usage reports from the BlueXP digital wallet. These usage reports provide capacity details of your subscriptions and tell you how you're being charged for the resources in your Cloud Volumes ONTAP subscriptions. The downloadable reports capture data at a point in time and can be easily shared with others.



The following reports are available for download. Capacity values shown are in TiB.

- **High-level usage:** This report includes the following information:
  - Total consumed capacity
  - Total precommitted capacity
  - Total BYOL capacity
  - Total Marketplace contracts capacity
  - Total PAYGO capacity
- **Cloud Volumes ONTAP package usage:** This report includes the following information for each package:
  - Total consumed capacity
  - Total precommitted capacity
  - Total BYOL capacity
  - Total Marketplace contracts capacity
  - Total PAYGO capacity
- **Storage VMs usage:** This report shows how charged capacity is broken down across Cloud Volumes ONTAP systems and storage virtual machines (SVMs). This information is only available in the report. It contains the following information:
  - Working environment ID and name (appears as the UUID)
  - Cloud
  - NetApp account ID
  - Working environment configuration



- SVM name
- Provisioned capacity
- Charged capacity roundup
- Marketplace billing term
- Cloud Volumes ONTAP package or feature
- Charging SaaS Marketplace subscription name
- Charging SaaS Marketplace subscription ID
- Workload type
- **Volumes usage:** This report shows how charged capacity is broken down by volumes in a working environment. This information is not available on any screen in the digital wallet. It includes the following information:
  - Working environment ID and name (appears as the UUID)
  - SVN name
  - Volume ID
  - Volume type
  - Volume provisioned capacity



FlexClone volumes aren't included in this report because these types of volumes don't incur charges.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. On the **Overview** tab, select **View** from the Cloud Volumes ONTAP tile.
3. Select **Usage report**.

The usage report downloads.

4. Open the downloaded file to access the reports.

## Manage Keystone subscriptions through BlueXP

Manage your Keystone subscriptions from the BlueXP digital wallet by enabling subscriptions for use with Cloud Volumes ONTAP and by requesting changes to the committed capacity for your subscription's service levels. Requesting additional capacity for a service level provides more storage for on-premises ONTAP clusters or for Cloud Volumes ONTAP systems.

NetApp Keystone is a flexible pay-as-you-grow subscription-based service that delivers a hybrid cloud experience for customers who prefer OpEx to CapEx or leasing.

[Learn more about Keystone](#)

### Authorize your account

Before you can use and manage Keystone subscriptions in BlueXP, you need to contact NetApp to authorize

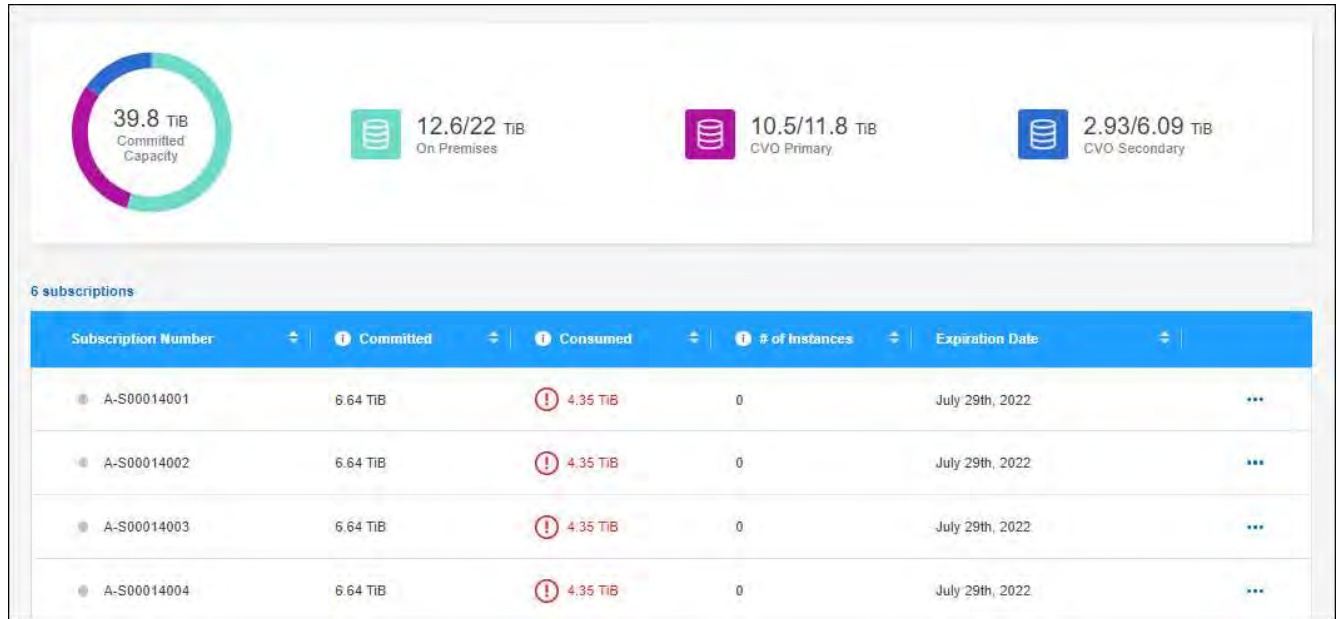
your BlueXP user account with your Keystone subscriptions.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Keystone Subscriptions**.
3. If you see the **Welcome to NetApp Keystone** page, send an email to the address listed on the page.

A NetApp representative will process your request by authorizing your user account to access the subscriptions.

4. Come back to the **Keystone Subscriptions** tab to view your subscriptions.




### Link a subscription

After NetApp authorizes your account, you can link Keystone subscriptions for use with Cloud Volumes ONTAP. This action enables users to select the subscription as the charging method for new Cloud Volumes ONTAP systems.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Keystone Subscriptions**.
3. For the subscription that you want to link, click **...** and select **Link**.



## Result

The subscription is now linked to your BlueXP organization or account and available to select when creating a Cloud Volumes ONTAP working environment.

## Request more or less committed capacity

If you want to change the committed capacity for your subscription's service levels, you can send a request to NetApp directly from BlueXP. Requesting additional capacity for a service level provides more storage for on-premises clusters or for Cloud Volumes ONTAP systems.

## Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Keystone Subscriptions**.
3. For the subscription that you want adjust the capacity, click **...** and select **View detail and edit**.
4. Enter the requested committed capacity for one or more subscriptions.

### Subscription Modification for A-S00014001

Service Level	Current Committed Capacity	Current Consumed Capacity	Requested Committed Capacity
Extreme	0.977 TiB	0.293 TiB	<input type="text" value="Enter amount"/> TiB
Premium	0.977 TiB	0.488 TiB	<input type="text" value="Enter amount"/> TiB
Performance	0 TiB	0 TiB	<input type="text" value="Enter amount"/> TiB
Standard	0.732 TiB	0.439 TiB	<input type="text" value="Enter amount"/> TiB
Value	0.977 TiB	<span>!</span> 0.879 TiB	<input type="text" value="Enter amount"/> TiB
Data Tiering	0 TiB	0 TiB	<input type="text" value="Enter amount"/> TiB
CVO Primary	1.96 TiB	<span>!</span> 1.76 TiB	<input type="text" value="3"/> TiB
CVO Secondary	1.02 TiB	0.488 TiB	<input type="text" value="Enter amount"/> TiB

#### Additional Information

Is there anything else we should know about your request?  
Please be as descriptive as possible.

5. Scroll down, enter any additional details for the request, and then click **Submit**.

## Result

Your request creates a ticket in NetApp's system for processing.

## Monitor usage

The BlueXP digital advisor dashboard enables you to monitor Keystone subscription usage and to generate reports.

[Learn more about monitoring subscription usage](#)

## Unlink a subscription

If you no longer want to use a Keystone subscription with BlueXP, you can unlink the subscription. Note that you can only unlink a subscription that isn't attached to an existing Cloud Volumes ONTAP subscription.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Keystone**.
3. For the subscription that you want to unlink, click **...** and select **Unlink**.

### Result

The subscription is unlinked from your BlueXP organization or account and no longer available to select when creating a Cloud Volumes ONTAP working environment.

## Manage node-based licensing for Cloud Volumes ONTAP

Manage node-based licenses in the BlueXP digital wallet to ensure that each Cloud Volumes ONTAP system has a valid license with the required capacity.

*Node-based licenses* are the previous generation licensing model (and not available for new customers):

- Bring your own license (BYOL) licenses purchased from NetApp
- Hourly pay-as-you-go (PAYGO) subscriptions from your cloud provider's marketplace

The *BlueXP digital wallet* enables you to manage licenses for Cloud Volumes ONTAP from a single location. You can add new licenses and update existing licenses.

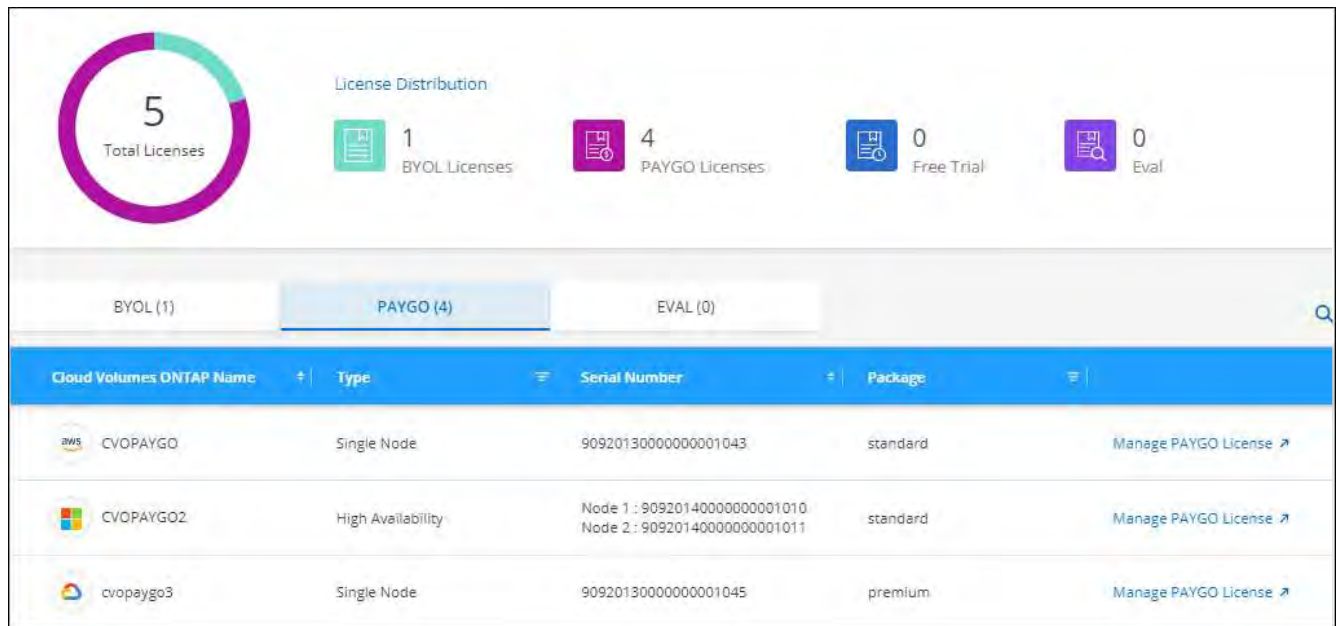
[Learn more about Cloud Volumes ONTAP licenses.](#)

## Manage PAYGO licenses

The BlueXP digital wallet page enables you to view details about each of your PAYGO Cloud Volumes ONTAP systems, including the serial number and PAYGO license type.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. Click **PAYGO**.
6. View details in the table about each of your PAYGO licenses.



7. If needed, click **Manage PAYGO License** to change the PAYGO license or to change the instance type.

## Manage BYOL licenses

Manage licenses that you purchased directly from NetApp by adding and removing system licenses and extra capacity licenses.



NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAP](#).

## Add unassigned licenses

Add a node-based license to the BlueXP digital wallet so that you can select the license when you create a new Cloud Volumes ONTAP system. The digital wallet identifies these licenses as *unassigned*.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. Click **Unassigned**.
6. Click **Add Unassigned Licenses**.
7. Enter the serial number of the license or upload the license file.

If you don't have the license file yet, refer to the section below.

8. Click **Add License**.

### Result

BlueXP adds the license to the digital wallet. The license will be identified as unassigned until you associate it with a new Cloud Volumes ONTAP system. After that happens, the license moves to the **BYOL** tab in the digital wallet.

## Exchange unassigned node-based licenses

If you have an unassigned node-based license for Cloud Volumes ONTAP that you haven't used, you can exchange the license by converting it to a BlueXP backup and recovery license, a BlueXP classification license, or a BlueXP tiering license.

Exchanging the license revokes the Cloud Volumes ONTAP license and creates a dollar-equivalent license for the service:

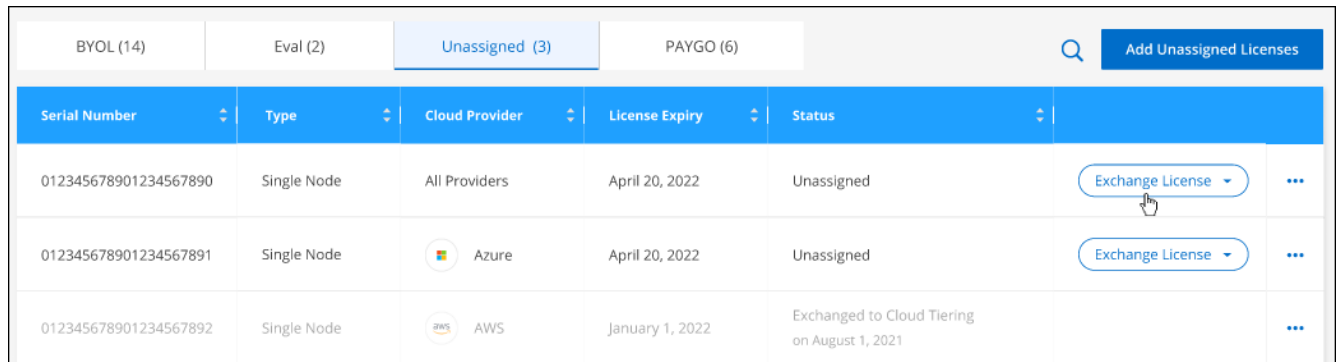
- Licensing for a Cloud Volumes ONTAP HA pair is converted to a 51 TiB direct license
- Licensing for a Cloud Volumes ONTAP single node is converted to a 32 TiB direct license

The converted license has the same expiration date as the Cloud Volumes ONTAP license.

[View walkthrough of how to exchange node-based licenses.](#)

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. Click **Unassigned**.
6. Click **Exchange License**.



Serial Number	Type	Cloud Provider	License Expiry	Status	
012345678901234567890	Single Node	All Providers	April 20, 2022	Unassigned	Exchange License ▾ ...
012345678901234567891	Single Node	Azure	April 20, 2022	Unassigned	Exchange License ▾ ...
012345678901234567892	Single Node	AWS	January 1, 2022	Exchanged to Cloud Tiering on August 1, 2021	...

7. Select the service that you'd like to exchange the license with.
8. If you're prompted, select an additional license for the HA pair.
9. Read the legal consent and click **Agree**.

### Result

BlueXP converts the unassigned license to the service that you selected. You can view the new license in the **Data Services Licenses** tab.

### Obtain a system license file

In most cases, BlueXP can automatically obtain your license file using your NetApp Support Site account. But if it can't, then you'll need to manually upload the license file. If you don't have the license file, you can obtain it from netapp.com.

### Steps

1. Go to the [NetApp License File Generator](#) and log in using your NetApp Support Site credentials.
2. Enter your password, choose your product, enter the serial number, confirm that you have read and accepted the privacy policy, and then click **Submit**.

### Example

3. Choose whether you want to receive the serialnumber.NLF JSON file through email or direct download.

### Update a system license

When you renew a BYOL subscription by contacting a NetApp representative, BlueXP automatically obtains the new license from NetApp and installs it on the Cloud Volumes ONTAP system. If BlueXP can't access the license file over the secure internet connection, you can obtain the file yourself and then manually upload the file to BlueXP.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. In the **BYOL** tab, expand the details for a Cloud Volumes ONTAP system.
6. Click the action menu next to the system license and select **Update License**.
7. Upload the license file (or files if you have an HA pair).
8. Click **Update License**.



## Result

BlueXP updates the license on the Cloud Volumes ONTAP system.

## Manage extra capacity licenses

You can purchase extra capacity licenses for a Cloud Volumes ONTAP BYOL system to allocate more than the 368 TiB of capacity that's provided with a BYOL system license. For example, you might purchase one extra license capacity to allocate up to 736 TiB of capacity to Cloud Volumes ONTAP. Or you could purchase three extra capacity licenses to get up to 1.4 PiB.

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

## Add capacity licenses

Purchase an extra capacity license by contacting us through the chat icon in the lower-right of BlueXP. After you purchase the license, you can apply it to a Cloud Volumes ONTAP system.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. In the **BYOL** tab, expand the details for a Cloud Volumes ONTAP system.
6. Click **Add Capacity License**.
7. Enter the serial number or upload the license file (or files if you have an HA pair).
8. Click **Add Capacity License**.

## Update capacity licenses

If you extended the term of an extra capacity license, you'll need to update the license in BlueXP.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. In the **BYOL** tab, expand the details for a Cloud Volumes ONTAP system.
6. Click the action menu next to the capacity license and select **Update License**.
7. Upload the license file (or files if you have an HA pair).
8. Click **Update License**.

## Remove capacity licenses

If an extra capacity license expired and is no longer in use, then you can remove it at any time.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.



2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. In the **BYOL** tab, expand the details for a Cloud Volumes ONTAP system.
6. Click the action menu next to the capacity license and select **Remove License**.
7. Click **Remove**.

## Change between PAYGO and BYOL

Converting a system from PAYGO by-node licensing to BYOL by-node licensing (and vice versa) isn't supported. If you want to switch between a pay-as-you-go subscription and a BYOL subscription, then you need to deploy a new system and replicate data from the existing system to the new system.

### Steps

1. Create a new Cloud Volumes ONTAP working environment.
2. Set up a one-time data replication between the systems for each volume that you need to replicate.

[Learn how to replicate data between systems](#)

3. Terminate the Cloud Volumes ONTAP system that you no longer need by deleting the original working environment.

[Learn how to delete a Cloud Volumes ONTAP working environment.](#)

### Related links

<https://docs.netapp.com/us-en/bluexp-cloud-volumes-ontap/concept-licensing.html#end-of-availability-of-node-based-licenses>>End of availability of node-based licenses</a><br> [xref:./task-convert-node-capacity.html](#)>Convert node-based licenses to capacity based</a>

## Volume and LUN administration

### Create a FlexVol volume on a Cloud Volumes ONTAP system

If you need more storage after you launch your initial Cloud Volumes ONTAP system, you can create new FlexVol volumes for NFS, CIFS, or iSCSI from BlueXP.

BlueXP provides several ways to create a new volume:

- Specify details for a new volume and let BlueXP handle the underlying data aggregates for you. [Learn more](#)
- Create a volume on a data aggregate of your choice. [Learn more](#)
- Create a volume on the second node in an HA configuration. [Learn more](#)

### Before you begin

A few notes about volume provisioning:

- When you create an iSCSI volume, BlueXP automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, [use](#)

[the IQN to connect to the LUN from your hosts.](#)

- You can create additional LUNs from ONTAP System Manager or the ONTAP CLI.
- If you want to use CIFS in AWS, you must have set up DNS and Active Directory. For details, refer to [Networking requirements for Cloud Volumes ONTAP for AWS](#).
- If your Cloud Volumes ONTAP configuration supports the Amazon EBS Elastic Volumes feature, you might want to [learn more about what happens when you create a volume](#).

## Create a volume

The most common way to create a volume is to specify the type of volume that you need and then BlueXP handles the disk allocation for you. But you also have the option to choose the specific aggregate on which you want to create the volume.

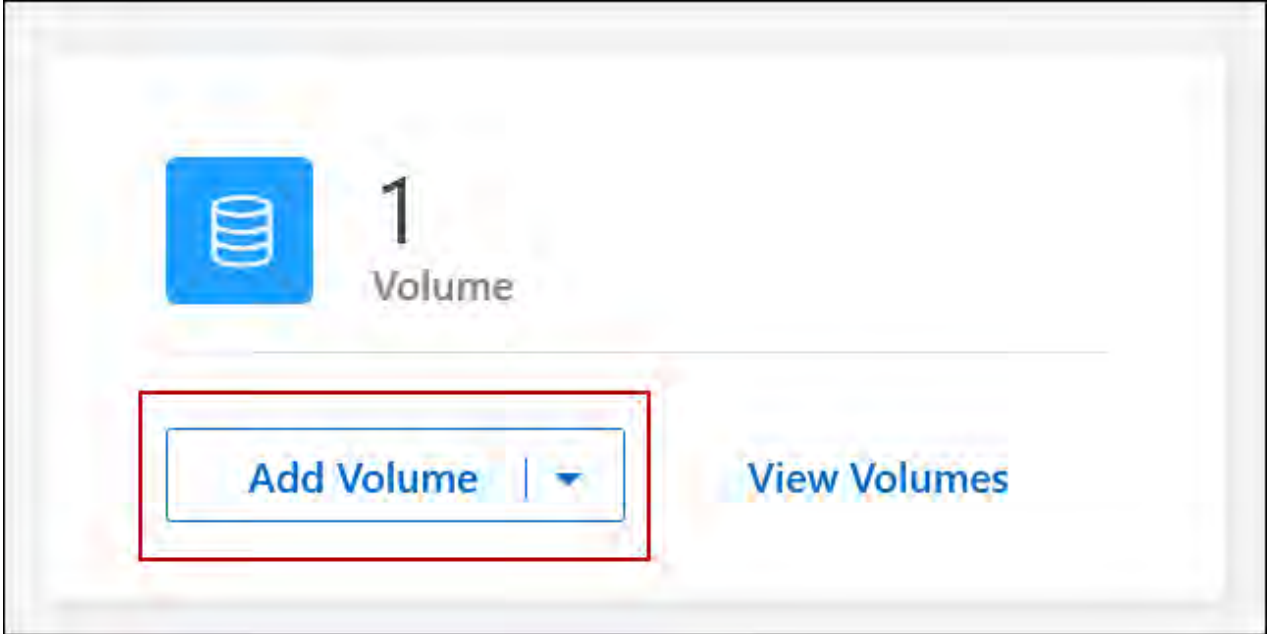
### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, double-click the name of the Cloud Volumes ONTAP system on which you want to provision a FlexVol volume.
3. Create a new volume by letting BlueXP handle the disk allocation for you, or choose a specific aggregate for the volume.

Choosing a specific aggregate is recommended only if you have a good understanding of the data aggregates on your Cloud Volumes ONTAP system.

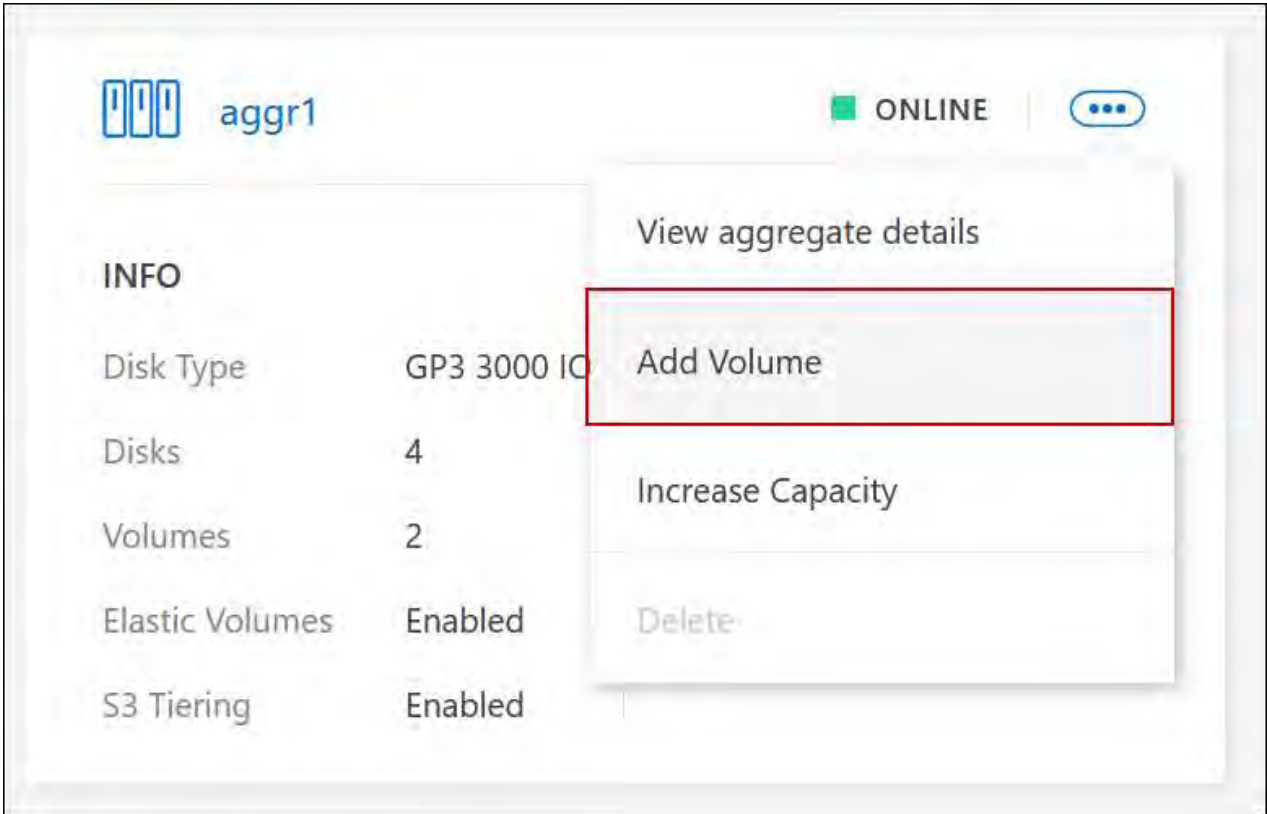
### Any aggregate

On the Overview tab, navigate to the Volumes tile, and click **Add Volume**.



### Specific aggregate

On the Aggregates tab, navigate to the desired aggregate tile. Click the menu icon, and then click **Add Volume**.



4. Follow the steps in the wizard to create the volume.

a. **Details, Protection, and Tags:** Enter basic details about the volume and select a Snapshot policy.

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

Field	Description
Volume Name	The identifiable name you can enter for the new volume.
Volume Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Storage VM (SVM)	A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an SVM or a vserver. Cloud Volumes ONTAP is configured with one storage VM by default, but some configurations support additional storage VMs. You can specify the Storage VM for the new volume.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.

- b. **Protocol:** Choose a protocol for the volume (NFS, CIFS, or iSCSI) and then provide the required information.

If you select CIFS and a server isn't set up, BlueXP prompts you to set up CIFS connectivity after you click **Next**.

[Learn about supported client protocols and versions.](#)

The following sections describe fields for which you might need guidance. The descriptions are organized by protocol.

## NFS

### Access control

Choose a custom export policy to make the volume available to clients.

### Export policy

Defines the clients in the subnet that can access the volume. By default, BlueXP enters a value that provides access to all instances in the subnet.

## CIFS

### Permissions and users/groups

Enables you to control the level of access to an SMB share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.

### DNS Primary and Secondary IP Address

The IP addresses of the DNS servers that provide name resolution for the CIFS server. The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.

If you're configuring Google Managed Active Directory, AD can be accessed by default with the 169.254.169.254 IP address.

### Active Directory Domain to join

The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.

### Credentials authorized to join the domain

The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.

### CIFS server NetBIOS name

A CIFS server name that is unique in the AD domain.

### Organizational Unit

The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.

- To configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter **OU=Computers,OU=corp** in this field.
- To configure Azure AD Domain Services as the AD server for Cloud Volumes ONTAP, enter **OU=AADDC Computers** or **OU=AADDC Users** in this field.  
[Azure Documentation: Create an Organizational Unit \(OU\) in an Azure AD Domain Services managed domain](#)
- To configure Google Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter **OU=Computers,OU=Cloud** in this field.  
[Google Cloud Documentation: Organizational Units in Google Managed Microsoft AD](#)

### DNS Domain

The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.

### **NTP Server**

Select **Use Active Directory Domain** to configure an NTP server using the Active Directory DNS. If you need to configure an NTP server using a different address, then you should use the API. For information, refer to the [BlueXP automation docs](#).

Note that you can configure an NTP server only when creating a CIFS server. It's not configurable after you create the CIFS server.

### **iSCSI**

#### **LUN**

iSCSI storage targets are called LUNs (logical units) and are presented to hosts as standard block devices. When you create an iSCSI volume, BlueXP automatically creates a LUN for you. We've made it simple by creating just one LUN per volume, so there's no management involved. After you create the volume, [use the IQN to connect to the LUN from your hosts](#).

#### **Initiator group**

Initiator groups (igroups) specify which hosts can access specified LUNs on the storage system

#### **Host initiator (IQN)**

iSCSI targets connect to the network through standard Ethernet network adapters (NICs), TCP offload engine (TOE) cards with software initiators, converged network adapters (CNAs) or dedicated host bus adapters (HBAs) and are identified by iSCSI qualified names (IQNs).

- c. **Disk Type:** Choose an underlying disk type for the volume based on your performance needs and cost requirements.
  - [Sizing your system in AWS](#)
  - [Sizing your system in Azure](#)
  - [Sizing your system in Google Cloud](#)
- d. **Usage Profile & Tiering Policy:** Choose whether to enable or disable storage efficiency features on the volume and then select a [volume tiering policy](#).

ONTAP includes several storage efficiency features that can reduce the total amount of storage that you need. NetApp storage efficiency features provide the following benefits:

#### **Thin provisioning**

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

#### **Deduplication**

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

#### **Compression**

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

- e. **Review:** Review details about the volume and then click **Add**.

## Result

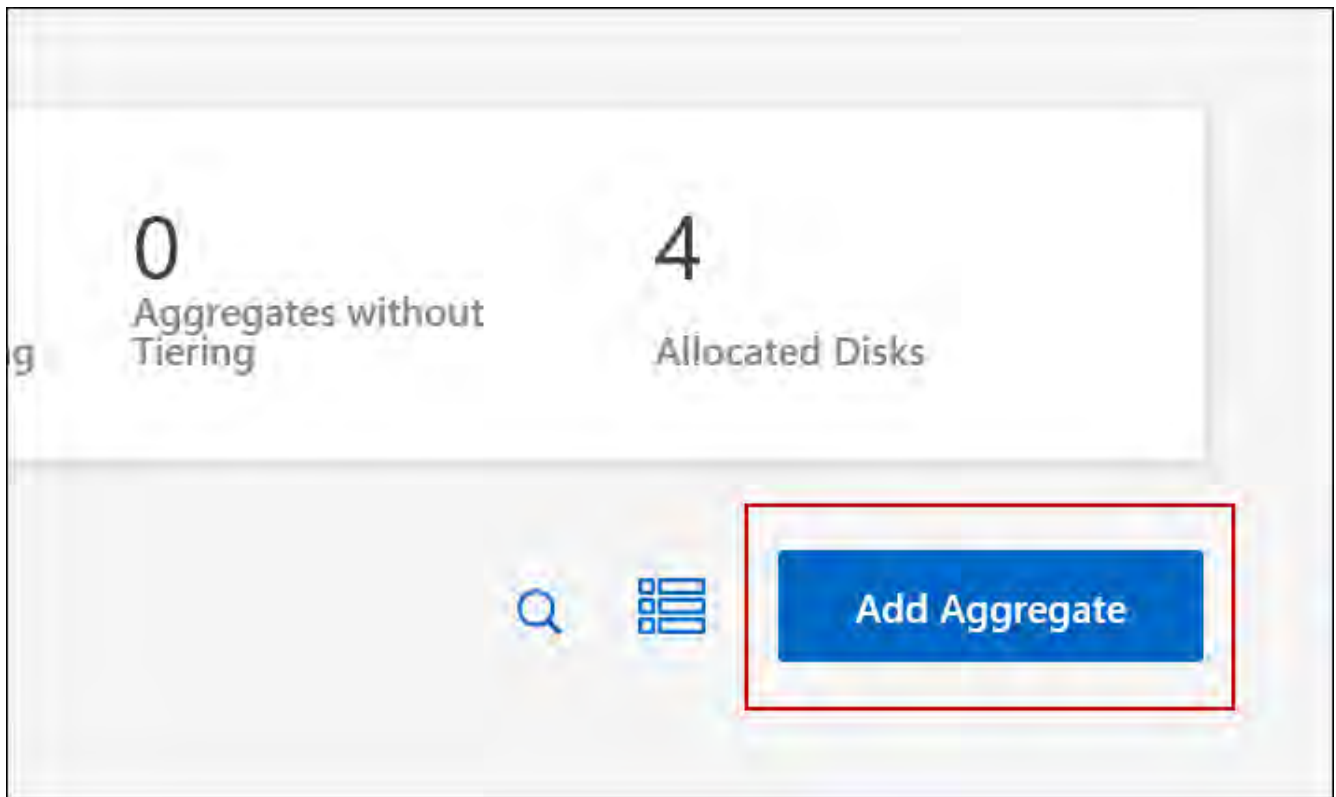
BlueXP creates the volume on the Cloud Volumes ONTAP system.

## Create a volume on the second node in an HA configuration

By default, BlueXP creates volumes on the first node in an HA configuration. If you need an active-active configuration, in which both nodes serve data to clients, you must create aggregates and volumes on the second node.

## Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, double-click the name of the Cloud Volumes ONTAP working environment on which you want to manage aggregates.
3. On the Aggregates tab, click **Add Aggregate**.
4. From the *Add Aggregate* screen, create the aggregate.



5. For Home Node, choose the second node in the HA pair.
6. After BlueXP creates the aggregate, select it and then click **Create volume**.
7. Enter details for the new volume, and then click **Create**.

## Result

BlueXP creates the volume on the second node in the HA pair.



For HA pairs deployed in multiple AWS Availability Zones, you must mount the volume to clients by using the floating IP address of the node on which the volume resides.

## After you create a volume

If you provisioned a CIFS share, give users or groups permissions to the files and folders and verify that those users can access the share and create a file.

If you want to apply quotas to volumes, you must use ONTAP System Manager or the ONTAP CLI. Quotas enable you to restrict or track the disk space and number of files used by a user, group, or qtree.

## Manage volumes on Cloud Volumes ONTAP systems

BlueXP enables you to manage volumes and CIFS servers. It also prompts you to move volumes to avoid capacity issues.

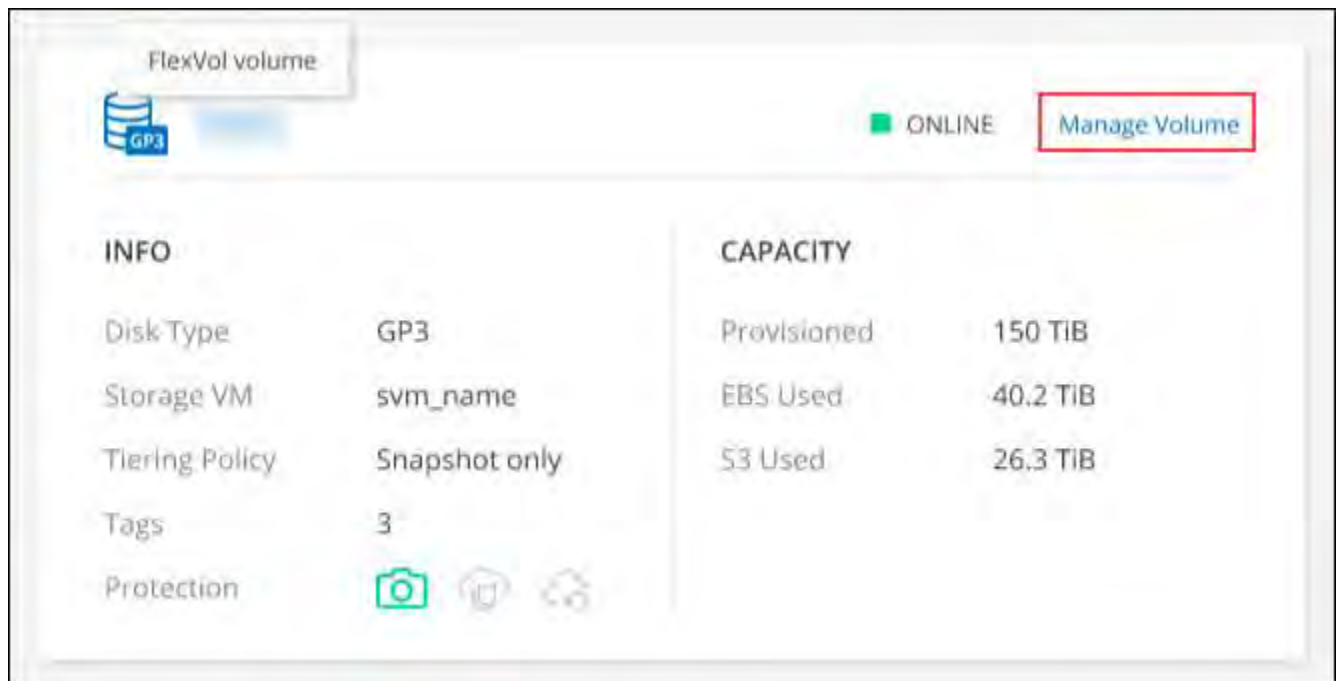
You can manage volumes in BlueXP Standard View or through ONTAP System Manager that is included within BlueXP for advanced volume management. The Standard View provides a limited set of options to modify your volumes. System Manager provides advanced level of management, such as cloning, resizing, changing settings for anti-ransomware, analytics, protection, and activity tracking, and moving volumes across tiers. For information, refer to [Administer Cloud Volumes ONTAP using System Manager](#).

### Manage volumes

By using the Standard View of BlueXP, you can manage volumes according to your storage needs. You can view, edit, clone, restore, and delete volumes.


#### Steps



1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, double-click the Cloud Volumes ONTAP working environment on which you want to manage volumes.
3. In the working environment, click the **Volumes** tab.



4. On the Volumes tab, navigate to the desired volume title and then click **Manage volume** to access the Manage Volumes right-side panel.



Task	Action
View information about a volume	Under Volume Actions in the Manage volumes panel, click <b>View volume details</b> .
Get the NFS mount command	<ol style="list-style-type: none"> <li>Under Volume Actions in the Manage volumes panel, click <b>Mount Command</b>.</li> <li>Click <b>Copy</b>.</li> </ol>
Clone a volume	<ol style="list-style-type: none"> <li>Under Volume Actions in the Manage volumes panel, click <b>Clone the volume</b>.</li> <li>Modify the clone name as needed, and then click <b>Clone</b>.</li> </ol> <p>This process creates a FlexClone volume. A FlexClone volume is a writable, point-in-time copy that is space-efficient because it uses a small amount of space for metadata, and then only consumes additional space as data is changed or added.</p> <p>To learn more about FlexClone volumes, refer to the <a href="#">ONTAP 9 Logical Storage Management Guide</a>.</p>
Edit a volume (read-write volumes only)	<ol style="list-style-type: none"> <li>Under Volume Actions in the Manage volumes panel, click <b>Edit volume settings</b></li> <li>Modify the volume's Snapshot policy, NFS protocol version, NFS access control list (export policy), or share permissions, and then click <b>Apply</b>.</li> </ol> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>If you need custom Snapshot policies, you can create them by using ONTAP System Manager.</p> </div>
Delete a volume	<ol style="list-style-type: none"> <li>Under Volume Actions in the Manage volumes panel, click <b>Delete the volume</b>.</li> <li>Under the Delete Volume window, enter the name of the volume you want to delete.</li> <li>Click <b>Delete</b> again to confirm.</li> </ol>
Create a Snapshot copy on demand	<ol style="list-style-type: none"> <li>Under Protection Actions in the Manage Volumes panel, click <b>Create a Snapshot copy</b>.</li> <li>Change the name, if needed, and then click <b>Create</b>.</li> </ol>
Restore data from a Snapshot copy to a new volume	<ol style="list-style-type: none"> <li>Under Protection Actions in the Manage Volumes panel, click <b>Restore from Snapshot copy</b>.</li> <li>Select a Snapshot copy, enter a name for the new volume, and then click <b>Restore</b>.</li> </ol>


Task	Action
Change the underlying disk type	<ol style="list-style-type: none"> <li>Under Advanced Actions in the Manage Volumes panel, click <b>Change Disk Type</b>.</li> <li>Select the disk type, and then click <b>Change</b>.</li> </ol> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  BlueXP moves the volume to an existing aggregate that uses the selected disk type or it creates a new aggregate for the volume. </div>
Change the tiering policy	<ol style="list-style-type: none"> <li>Under Advanced Actions in the Manage Volumes panel, click <b>Change Tiering Policy</b>.</li> <li>Select a different policy and click <b>Change</b>.</li> </ol> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  BlueXP moves the volume to an existing aggregate that uses the selected disk type with tiering, or it creates a new aggregate for the volume. </div>
Delete a volume	<ol style="list-style-type: none"> <li>Select a volume, and then click <b>Delete</b>.</li> <li>Type the name of the volume in the dialog.</li> <li>Click <b>Delete</b> again to confirm.</li> </ol>


## Resize a volume

By default, a volume automatically grows to a maximum size when it's out of space. The default value is 1,000, which means the volume can grow to 11 times its size. This value is configurable in the Connector's settings.

If you need to resize your volume, you can do it from ONTAP System Manager in BlueXP.

### Steps

- Click the System Manager view to resize a volume through ONTAP System Manager. Refer to [How to get started](#).
- From the left navigation menu, select **Storage > Volumes**.
- From the list of volumes, identify the one that you should resize.
- Click the options icon .
- Select **Resize**.
- On the **Resize Volume** screen, edit the capacity and Snapshot reserve percentage as required. You can compare the existing, available space with the modified capacity.
- Click **Save**.



Be sure to take your system's capacity limits into consideration as you resize volumes. Go to the [Cloud Volumes ONTAP Release Notes](#) for more information.

### Modify the CIFS server

If you change your DNS servers or Active Directory domain, you need to modify the CIFS server in Cloud Volumes ONTAP so that it can continue to serve storage to clients.

#### Steps

1. From the Overview tab of the working environment, click the Feature tab under the right-side panel.
2. Under the CIFS Setup field, click the **pencil icon** to display the CIFS Setup window.
3. Specify settings for the CIFS server:

Task	Action
Select Storage VM (SVM)	Selecting the Cloud Volume ONTAP storage virtual machine (SVM) displays its configured CIFS information.
Active Directory Domain to join	The FQDN of the Active Directory (AD) domain that you want the CIFS server to join.
Credentials authorized to join the domain	The name and password of a Windows account with sufficient privileges to add computers to the specified Organizational Unit (OU) within the AD domain.

Task	Action
DNS Primary and Secondary IP Address	<p>The IP addresses of the DNS servers that provide name resolution for the CIFS server.</p> <p>The listed DNS servers must contain the service location records (SRV) needed to locate the Active Directory LDAP servers and domain controllers for the domain that the CIFS server will join.</p> <p>If you're configuring Google Managed Active Directory, AD can be accessed by default with the 169.254.169.254 IP address.</p>
DNS Domain	The DNS domain for the Cloud Volumes ONTAP storage virtual machine (SVM). In most cases, the domain is the same as the AD domain.
CIFS server NetBIOS name	A CIFS server name that is unique in the AD domain.
Organizational Unit	<p>The organizational unit within the AD domain to associate with the CIFS server. The default is CN=Computers.</p> <ul style="list-style-type: none"> <li>• To configure AWS Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter <b>OU=Computers,OU=corp</b> in this field.</li> <li>• To configure Azure AD Domain Services as the AD server for Cloud Volumes ONTAP, enter <b>OU=AADDC Computers</b> or <b>OU=AADDC Users</b> in this field. <a href="#">Azure Documentation: Create an Organizational Unit (OU) in an Azure AD Domain Services managed domain</a></li> <li>• To configure Google Managed Microsoft AD as the AD server for Cloud Volumes ONTAP, enter <b>OU=Computers,OU=Cloud</b> in this field. <a href="#">Google Cloud Documentation: Organizational Units in Google Managed Microsoft AD</a></li> </ul>

4. Click **Set**.

### Result

Cloud Volumes ONTAP updates the CIFS server with the changes.

### Move a volume

Move volumes for capacity utilization, improved performance, and to satisfy service-level agreements.

You can move a volume in ONTAP System Manager by selecting a volume and the destination aggregate, starting the volume move operation, and optionally monitoring the volume move job. When using System Manager, a volume move operation finishes automatically.

### Steps

1. Use ONTAP System Manager or the ONTAP CLI to move the volumes to the aggregate.

In most situations, you can use System Manager to move volumes.

For instructions, refer to the [ONTAP 9 Volume Move Express Guide](#).

## Move a volume when BlueXP displays an Action Required message

BlueXP might display an Action Required message that says moving a volume is necessary to avoid capacity issues, but that you need to correct the issue yourself. If this happens, you need to identify how to correct the issue and then move one or more volumes.



BlueXP displays these Action Required messages when an aggregate has reached 90% used capacity. If data tiering is enabled, the messages display when an aggregate has reached 80% used capacity. By default, 10% free space is reserved for data tiering. [Learn more about the free space ratio for data tiering.](#)

### Steps

1. [Identify how to correct capacity issues.](#)
2. Based on your analysis, move volumes to avoid capacity issues:
  - [Move volumes to another system to avoid capacity issues.](#)
  - [Move volumes to another aggregate to avoid capacity issues.](#)

### Identify how to correct capacity issues

If BlueXP can't provide recommendations for moving a volume to avoid capacity issues, you must identify the volumes that you need to move and whether you should move them to another aggregate on the same system or to another system.

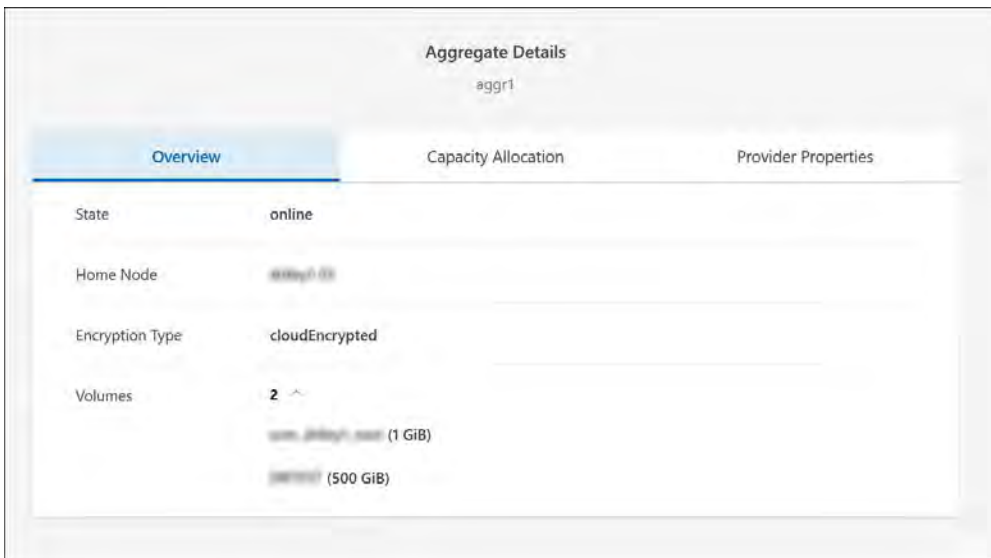
### Steps

1. View the advanced information in the Action Required message to identify the aggregate that has reached its capacity limit.

For example, the advanced information should say something similar to the following: Aggregate aggr1 has reached its capacity limit.

2. Identify one or more volumes to move out of the aggregate:
  - a. In the working environment, click the **Aggregates tab**.
  - b. Navigate to the desired aggregate tile, and then click the ... (**ellipses icon**) > **View aggregate details**.
  - c. Under the Overview tab of the Aggregate Details screen, review the size of each volume and choose one or more volumes to move out of the aggregate.

You should choose volumes that are large enough to free space in the aggregate so that you avoid additional capacity issues in the future.



3. If the system has not reached the disk limit, you should move the volumes to an existing aggregate or a new aggregate on the same system.

For information, refer to [Move volumes to another aggregate to avoid capacity issues](#).

4. If the system has reached the disk limit, do any of the following:
  - a. Delete any unused volumes.
  - b. Rearrange volumes to free space on an aggregate.

For information, refer to [Move volumes to another aggregate to avoid capacity issues](#).

- c. Move two or more volumes to another system that has space.

For information, refer to [Move volumes to another aggregate to avoid capacity issues](#).

### Move volumes to another system to avoid capacity issues

You can move one or more volumes to another Cloud Volumes ONTAP system to avoid capacity issues. You might need to do this if the system reached its disk limit.

#### About this task

You can follow the steps in this task to correct the following Action Required message:

Moving a volume is necessary to avoid capacity issues; however, BlueXP cannot perform this action for you because the system has reached the disk limit.

#### Steps

1. Identify a Cloud Volumes ONTAP system that has available capacity, or deploy a new system.
2. Drag and drop the source working environment on the target working environment to perform a one-time data replication of the volume.

For information, refer to [Replicating data between systems](#).

3. Go to the Replication Status page, and then break the SnapMirror relationship to convert the replicated

volume from a data protection volume to a read/write volume.

For information, refer to [Managing data replication schedules and relationships](#).

4. Configure the volume for data access.

For information about configuring a destination volume for data access, refer to the [ONTAP 9 Volume Disaster Recovery Express Guide](#).

5. Delete the original volume.

For information, refer to [Manage volumes](#).

### Move volumes to another aggregate to avoid capacity issues

You can move one or more volumes to another aggregate to avoid capacity issues.

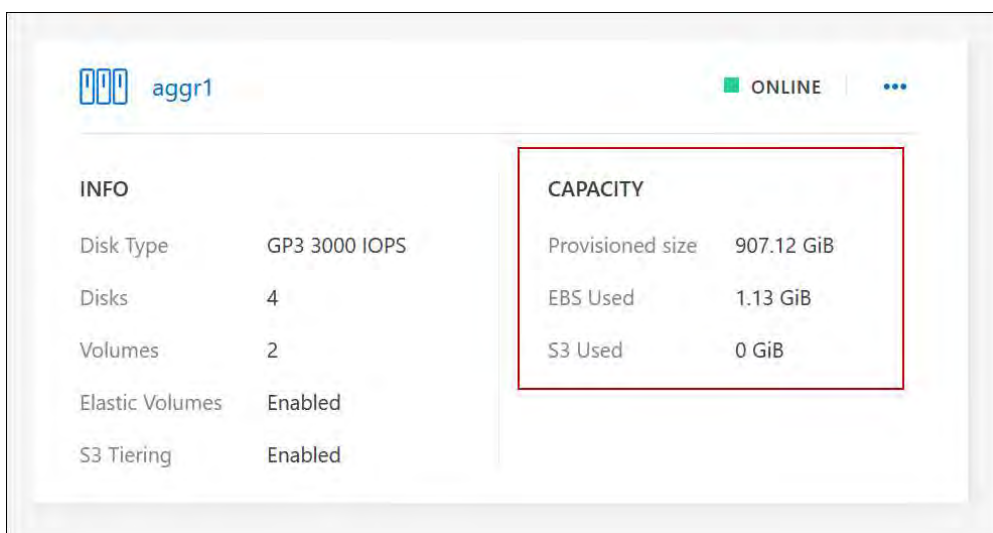
#### About this task

You can follow the steps in this task to correct the following Action Required message:

Moving two or more volumes is necessary to avoid capacity issues; however, BlueXP cannot perform this action for you.

#### Steps

1. Verify whether an existing aggregate has available capacity for the volumes that you need to move:
  - a. In the working environment, click the **Aggregates tab**.
  - b. Navigate to the desired aggregate tile, and then click the ... (**ellipses icon**) > **View aggregate details**.
  - c. Under the aggregate tile, view the available capacity (provisioned size minus used aggregate capacity).



2. If needed, add disks to an existing aggregate:
  - a. Select the aggregate, then click the ... (**ellipses icon**) > **Add Disks**.
  - b. Select the number of disks to add, and then click **Add**.
3. If no aggregates have available capacity, create a new aggregate.

For information, refer to [Creating aggregates](#).

4. Use ONTAP System Manager or the ONTAP CLI to move the volumes to the aggregate.
5. In most situations, you can use System Manager to move volumes.

For instructions, refer to the [ONTAP 9 Volume Move Express Guide](#).

### Reasons why a volume move might perform slowly

Moving a volume might take longer than you expect if any of the following conditions are true for Cloud Volumes ONTAP:

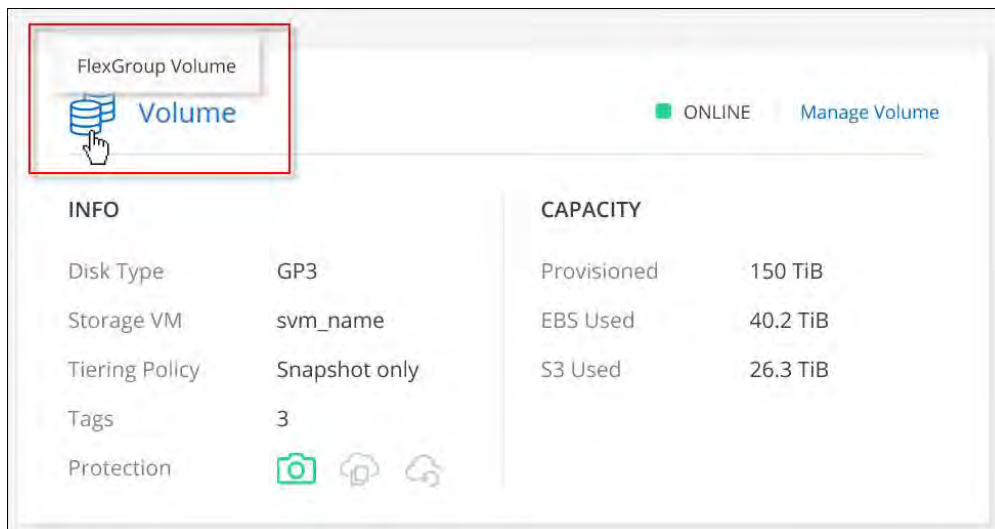
- The volume is a clone.
- The volume is a parent of a clone.
- The source or destination aggregate has a single Throughput Optimized HDD (st1) disk.
- One of the aggregates uses an older naming scheme for objects. Both aggregates have to use the same name format.

An older naming scheme is used if data tiering was enabled on an aggregate in the 9.4 release or earlier.

- The encryption settings don't match on the source and destination aggregates, or a rekey is in progress.
- The *-tiering-policy* option was specified on the volume move to change the tiering policy.
- The *-generate-destination-key* option was specified on the volume move.

### View FlexGroup Volumes

You can view FlexGroup volumes created through ONTAP System Manager or the ONTAP CLI directly through the Volumes tab within BlueXP. Identical to the information provided for FlexVol volumes, BlueXP provides detailed information for created FlexGroup volumes through a dedicated Volumes tile. Under the Volumes tile, you can identify each FlexGroup volume group through the icon's hover text. Additionally, you can identify and sort FlexGroup volumes under the volumes list view through the Volume Style column.



Currently, you can only view existing FlexGroup volumes under BlueXP. The ability to create FlexGroup volumes in BlueXP is not available but planned for a future release.



## Tier inactive Cloud Volumes ONTAP data to a low-cost object storage

You can reduce storage costs for Cloud Volumes ONTAP by combining an SSD or HDD performance tier for hot data with an object storage capacity tier for inactive data. Data tiering is powered by FabricPool technology. For a high-level overview, refer to [Data tiering overview](#).

To set up data tiering, you need to do the following:

1

### Choose a supported configuration

Most configurations are supported. If you have a Cloud Volumes ONTAP system running the most recent version, then you are good to go. [Learn more](#).

2

### Ensure connectivity between Cloud Volumes ONTAP and object storage

- For AWS, you'll need a VPC Endpoint to S3. [Learn more](#).
- For Azure, you won't need to do anything as long as BlueXP has the required permissions. [Learn more](#).
- For Google Cloud, you need to configure the subnet for Private Google Access and set up a service account. [Learn more](#).

3

### Ensure that you have an aggregate with tiering enabled

Data tiering should be enabled on an aggregate to enable it on a volume. You should be aware of the requirements for new volumes and for existing volumes. [Learn more](#).

4

### Choose a tiering policy when creating, modifying, or replicating a volume

BlueXP prompts you to choose a tiering policy when you create, modify, or replicate a volume.

- [Tier data from read-write volumes](#)
- [Tier data from data protection volumes](#)

#### What's not required for data tiering?

- You don't need to install a feature license to enable data tiering.
- You don't need to create an object store for the capacity tier. BlueXP does that for you.
- You don't need to enable data tiering at the system level.



BlueXP creates an object store for cold data when it creates the system, [as long as there are no connectivity or permissions issues](#). After that, you just need to enable data tiering on volumes (and in some cases, [on aggregates](#)).

## Configurations that support data tiering

You can enable data tiering when using specific configurations and features.

## Support in AWS

- Data tiering is supported in AWS beginning with Cloud Volumes ONTAP 9.2.
- The performance tier can be General Purpose SSDs (gp3 or gp2) or Provisioned IOPS SSDs (io1).



We do not recommend tiering data to object storage when using Throughput Optimized HDDs (st1).

- The inactive data is tiered to Amazon S3 buckets. Tiering to other providers is not supported.

## Support in Azure

- Data tiering is supported in Azure as follows:
  - Version 9.4 in with single node systems
  - Version 9.6 in with HA pairs
- The performance tier can be Premium SSD managed disks, Standard SSD managed disks, or Standard HDD managed disks.
- The inactive data is tiered to Microsoft Azure Blob. Tiering to other providers is not supported.

## Support in Google Cloud

- Data tiering is supported in Google Cloud beginning with Cloud Volumes ONTAP 9.6.
- The performance tier can be either SSD persistent disks, balanced persistent disks, or standard persistent disks.
- The inactive data is tiered to Google Cloud Storage. Tiering to other providers is not supported.

## Feature interoperability

- Data tiering is supported with encryption technologies.
- Thin provisioning must be enabled on volumes.

## Requirements

Depending on your cloud provider, certain connections and permissions must be set up so that Cloud Volumes ONTAP can tier cold data to object storage.

### Requirements to tier cold data to AWS S3

Ensure that Cloud Volumes ONTAP has a connection to S3. The best way to provide that connection is by creating a VPC Endpoint to the S3 service. For instructions, refer to the [AWS Documentation: Creating a Gateway Endpoint](#).

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

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

## Requirements to tier cold data to Azure Blob storage

You don't need to set up a connection between the performance tier and the capacity tier as long as BlueXP has the required permissions. BlueXP enables a VNet service endpoint for you if the custom role for the Connector has these permissions:

```
"Microsoft.Network/virtualNetworks/subnets/write",  
"Microsoft.Network/routeTables/join/action",
```

The custom role includes the permissions by default. [View Azure permission for the Connector](#)

## Requirements to tier cold data to a Google Cloud Storage bucket

- The subnet in which Cloud Volumes ONTAP resides must be configured for Private Google Access. For instructions, refer to [Google Cloud Documentation: Configuring Private Google Access](#).
- A service account must be attached to Cloud Volumes ONTAP.

[Learn how to set up this service account.](#)

You're prompted to select this service account when you create a Cloud Volumes ONTAP working environment.

If you don't select a service account during deployment, you'll need to shut down Cloud Volumes ONTAP, go to the Google Cloud console, and then attach the service account to the Cloud Volumes ONTAP instances. You can then enable data tiering as described in the next section.

- To encrypt the bucket with customer-managed encryption keys, enable the Google Cloud storage bucket to use the key.

[Learn how to use customer-managed encryption keys with Cloud Volumes ONTAP.](#)

## Enable data tiering after implementing the requirements

BlueXP creates an object store for cold data when the system is created, as long as there are no connectivity or permissions issues. If you didn't implement the requirements listed above until after you created the system, then you'll need to manually enable tiering through the API or ONTAP System Manager, which creates the object store.



The ability to enable tiering through the BlueXP user interface will be available in a future Cloud Volumes ONTAP release.

## Ensure that tiering is enabled on aggregates

Data tiering must be enabled on an aggregate in order to enable data tiering on a volume. You should be aware of the requirements for new volumes and for existing volumes.

### • New volumes

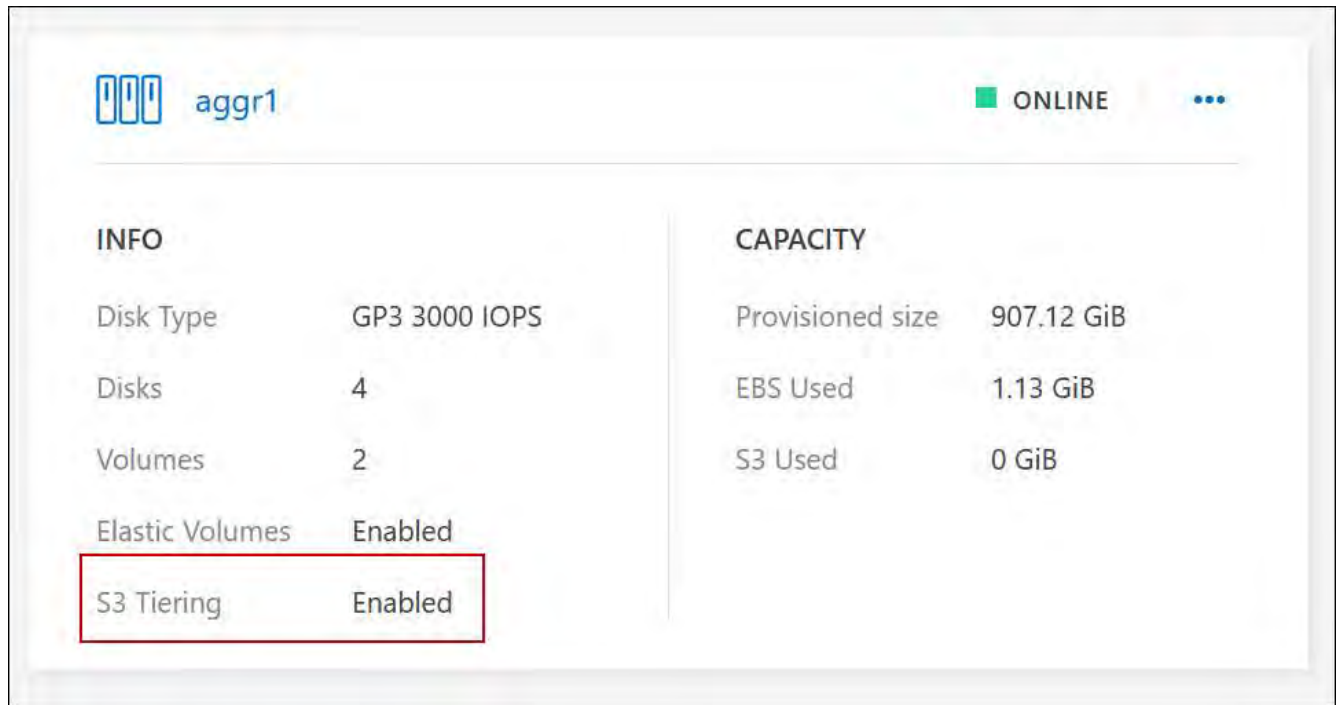
If you're enabling data tiering on a new volume, then you don't need to worry about enabling data tiering on an aggregate. BlueXP creates the volume on an existing aggregate that has tiering enabled, or it creates a new aggregate for the volume if a data tiering-enabled aggregate doesn't already exist.

- **Existing volumes**

To enable data tiering on an existing volume, ensure it is enabled on the underlying aggregate. If data tiering isn't enabled on the existing aggregate, then you'll need to use ONTAP System Manager to attach an existing aggregate to the object store.

**Steps to confirm whether tiering is enabled on an aggregate**

1. Open the working environment in BlueXP.
2. Click the Aggregates tab.
3. Navigate to the desired tile and verify whether tiering is enabled or disabled on the aggregate.



**Steps to enable tiering on an aggregate**

1. In ONTAP System Manager, click **Storage > Tiers**.
2. Click the action menu for the aggregate and select **Attach Cloud Tiers**.
3. Select the cloud tier to attach and click **Save**.

**What's next?**

You can now enable data tiering on new and existing volumes, as explained in the next section.

**Tier data from read-write volumes**

Cloud Volumes ONTAP can tier inactive data on read-write volumes to cost-effective object storage, freeing up the performance tier for hot data.

**Steps**

1. In Volumes tab under the working environment, create a new volume or change the tier of an existing volume:

Task	Action
Create a new volume	Click <b>Add New Volume</b> .
Modify an existing volume	Select the desired volume tile, click <b>Manage volume</b> to access the Manage Volumes right-side panel, and then click <b>Advanced actions</b> and <b>Change tiering policy</b> under the right panel.

2. Select a tiering policy.

For a description of these policies, refer to [Data tiering overview](#).

### Example

#### Change Tiering Policy

Volume\_1

**Tiering Policy**

**Auto** - Tiers cold Snapshot copies and cold user data from the active file system to object storage.  
Minimum cooling days: 31 (2-183)

**All** - Immediately tiers all data (not including metadata) to object storage.

**Snapshot Only** - Tiers cold Snapshot copies to object storage.

**None** - Data tiering is disabled.

**S3 Storage classes** Standard-Infrequent Access

**S3 Storage Encryption Key** aws/s3

**!** This action is non-disruptive and changing the tier impacts cost, performance, and maximum capacity. Refer to [BlueXP documentation](#) for more details.

BlueXP creates a new aggregate for the volume if a data tiering-enabled aggregate does not already exist.

### Tier data from data protection volumes


Cloud Volumes ONTAP can tier data from a data protection volume to a capacity tier. If you activate the destination volume, the data gradually moves to the performance tier as it is read.

### Steps

1. From the left navigation menu, select **Storage > Canvas**.

2. On the Canvas page, select the working environment that contains the source volume, and then drag it to the working environment to which you want to replicate the volume.
3. Follow the prompts until you reach the tiering page and enable data tiering to object storage.

### Example



For help with replicating data, refer to [Replicating data to and from the cloud](#).

### Change the storage class for tiered data

After you deploy Cloud Volumes ONTAP, you can reduce your storage costs by changing the storage class for inactive data that hasn't been accessed for 30 days. The access costs are higher if you do access the data, so you must take that into consideration before you change the storage class.

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

For information about supported storage classes, refer to [Data tiering overview](#).

### Steps

1. From the working environment, click the menu icon and then click **Storage Classes** or **Blob Storage Tiering**.
2. Choose a storage class and then click **Save**.

### Change the free space ratio for data tiering

The free space ratio for data tiering defines how much free space is required on Cloud Volumes ONTAP SSDs/HDDs when tiering data to object storage. The default setting is 10% free space, but you can tweak the setting based on your requirements.

For example, you might choose less than 10% free space to ensure that you are utilizing the purchased capacity. BlueXP can then purchase additional disks for you when additional capacity is required (up until you reach the disk limit for the aggregate).



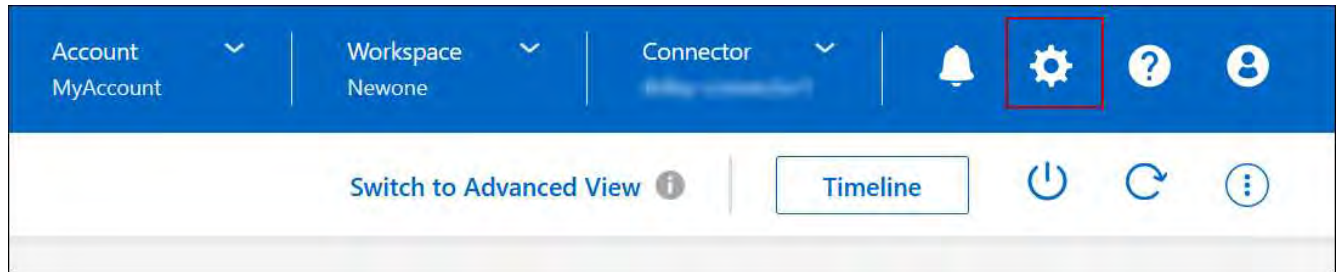
If there isn't sufficient space, then Cloud Volumes ONTAP can't move the data and you might experience performance degradation. Any change should be done with caution. If you're unsure, reach out to NetApp Support for guidance.

The ratio is important for disaster recovery scenarios because as data is read from the object store, Cloud Volumes ONTAP moves the data to SSDs/HDDs to provide better performance. If there isn't sufficient space, then Cloud Volumes ONTAP can't move the data. Take this into consideration when changing the ratio so that you can meet your business requirements.

### Steps

1. In the upper right of the BlueXP console, click the **Settings** icon, and select **Cloud Volumes ONTAP**

## Settings.



2. Under **Capacity**, click **Aggregate Capacity Thresholds - Free Space Ratio for Data Tiering**.
3. Change the free space ratio based on your requirements and click **Save**.

### Change the cooling period for the auto tiering policy

If you enabled data tiering on a Cloud Volumes ONTAP volume using the *auto* tiering policy, you can adjust the default cooling period based on your business needs. This action is supported using ONTAP CLI and API only.

The cooling period is the number of days that user data in a volume must remain inactive before it is considered "cold" and moved to object storage.

The default cooling period for the auto tiering policy is 31 days. You can change the cooling period as follows:

- 9.8 or later: 2 days to 183 days
- 9.7 or earlier: 2 days to 63 days

### Step

1. Use the *minimumCoolingDays* parameter with your API request when creating a volume or modifying an existing volume.

### Remove an S3 bucket on decommissioning a working environment

You can delete an S3 bucket with the data tiered from a Cloud Volumes ONTAP working environment when you decommission the environment.

You can delete the S3 bucket only if:

- The Cloud Volume ONTAP working environment is deleted from BlueXP.
- All objects are deleted from the bucket and the S3 bucket is empty.

When you decommission a Cloud Volumes ONTAP working environment, the S3 bucket that was created for the environment is not deleted automatically. Instead, it remains in an orphaned state to prevent any accidental data loss. You can delete the objects in the bucket, then remove the S3 bucket itself, or keep it for later use. Refer to [ONTAP CLI: vservers object-store-server bucket delete](#).

### Connect to a LUN on Cloud Volumes ONTAP from your host system

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

Note the following:

- BlueXP's automatic capacity management doesn't apply to LUNs. When BlueXP creates a LUN, it disables the autogrow feature.
- You can create additional LUNs from ONTAP System Manager or the ONTAP CLI.

### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, double-click the Cloud Volumes ONTAP working environment on which you want to manage volumes.
3. In the working environment, click the **Volumes** tab.
4. On the Volumes tab, navigate to the desired volume title and then click **Manage volume** to access the Manage Volumes right-side panel.
5. Click **Target iQN**.
6. Click **Copy** to copy the iQN name.
7. Set up an iSCSI connection from the host to the LUN.
  - [ONTAP 9 iSCSI express configuration for Red Hat Enterprise Linux: Starting the iSCSI sessions with the target](#)
  - [ONTAP 9 iSCSI express configuration for Windows: Starting iSCSI sessions with the target](#)
  - [ONTAP SAN host configuration](#)

## Accelerate data access with FlexCache volumes on a Cloud Volumes ONTAP system

A FlexCache volume is a storage volume that caches SMB and NFS read data from an origin (or source) volume. Subsequent reads to the cached data result in faster access to that data.

You can use FlexCache volumes to speed up access to data or to offload traffic from heavily accessed volumes. FlexCache volumes help improve performance, especially when clients need to access the same data repeatedly, because the data can be served directly without having to access the origin volume. FlexCache volumes work well for system workloads that are read-intensive.

BlueXP provides management of FlexCache volumes with the [BlueXP volume caching](#) service.

You can also use the ONTAP CLI or ONTAP System Manager to create and manage FlexCache volumes:

- [FlexCache Volumes for Faster Data Access Power Guide](#)
- [Creating FlexCache volumes in System Manager](#)





## Work with FlexCache when the origin is encrypted

When configuring FlexCache on a Cloud Volumes ONTAP system where the origin volume is encrypted, additional steps are required, to ensure that the FlexCache volume can properly access and cache the encrypted data.

### Before you begin

1. **Encryption setup:** Ensure that the source volume is fully encrypted and operational. For Cloud Volumes ONTAP systems, this involves integrating with cloud-specific key management services. For AWS, this typically means using AWS Key Management Service (KMS). For information, refer to [Manage keys with AWS Key Management Service](#). For Azure, you need to set up Azure Key Vault for NetApp Volume Encryption (NVE). For information, refer to [Manage keys with Azure Key Vault](#). For Google Cloud, it is Google Cloud Key Management Service. For information, refer to [Manage keys with Google's Cloud Key Management Service](#).
2. **Key management services:** Before creating a FlexCache volume, verify that the key management services are configured correctly on the Cloud Volumes ONTAP system. This configuration is essential for the FlexCache volume to decrypt the data from the origin volume.
3. **Licensing:** Confirm that a valid FlexCache license is available and activated on the Cloud Volumes ONTAP system.
4. **ONTAP version:** Ensure that the ONTAP version of your Cloud Volumes ONTAP system supports FlexCache with encrypted volumes. Refer to the latest [ONTAP release notes](#) or compatibility matrix for more information.
5. **Network Configuration:** Ensure that the network configuration allows for seamless communication between the origin volume and the FlexCache volume. This includes proper routing and DNS resolution in a cloud environment.

### Steps

Create a FlexCache volume on your Cloud Volumes ONTAP system with an encrypted source volume. For

detailed steps and additional considerations, refer to the following sections:

- [FlexCache Volumes for Faster Data Access Power Guide](#)
- [Creating FlexCache volumes in System Manager](#)

## Aggregate administration

### Create an aggregate for Cloud Volumes ONTAP systems

You can create aggregates yourself or let BlueXP do it for you when it creates volumes. The benefit of creating aggregates yourself is that you can choose the underlying disk size, which enables you to size your aggregate for the capacity or the performance that you need.



All disks and aggregates must be created and deleted directly from BlueXP. You should not perform these actions from another management tool. Doing so can impact system stability, hamper the ability to add disks in the future, and potentially generate redundant cloud provider fees.

#### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, double-click the name of the Cloud Volumes ONTAP instance on which you want to manage aggregates.
3. On the Aggregates tab, click **Add Aggregate** and then specify details for the aggregate.

## AWS


- If you're prompted to choose a disk type and disk size, refer to [Plan your Cloud Volumes ONTAP configuration in AWS](#).
- If you're prompted to enter the aggregate's capacity size, then you're creating an aggregate on a configuration that supports the Amazon EBS Elastic Volumes feature. The following screenshot shows an example of a new aggregate comprised of gp3 disks.

1 Disk Type   2 Aggregate details   3 Tiering Data   4 Review



### Select Disk Type



Disk Type

GP3 - General Purpose SSD Dynamic Performance

 **General Purpose SSD (gp3) Disk Properties**

**Description:** General purpose SSD volume that balances price and performance (performance level is independent of storage capacity)

IOPS Value    Throughput MB/s 

12000    250 

[Learn more about support for Elastic Volumes.](#)

## Azure

For help with disk type and disk size, refer to [Plan your Cloud Volumes ONTAP configuration in Azure](#).

## Google Cloud

For help with disk type and disk size, refer to [Plan your Cloud Volumes ONTAP configuration in Google Cloud](#).

4. Click **Go**, and then click **Approve and Purchase**.

## Manage aggregates for Cloud Volumes ONTAP clusters

Manage aggregates yourself by adding disks, viewing information about the aggregates, and by deleting them.



All disks and aggregates must be created and deleted directly from BlueXP. You should not perform these actions from another management tool. Doing so can impact system stability, hamper the ability to add disks in the future, and potentially generate redundant cloud provider fees.

### Before you begin

If you want to delete an aggregate, you must have first deleted the volumes in the aggregate.

### About this task

If an aggregate is running out of space, you can move volumes to another aggregate by using ONTAP System Manager.


### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, double-click the Cloud Volumes ONTAP working environment on which you want to manage aggregates.
3. In the working environment, click the **Aggregates** tab.
4. On the Aggregates tab, navigate to the desired title and then click the ... (ellipses icon).

INFO		CAPACITY	
Disk Type	GP3 3000 IOPS	Provisioned size	907.12 GiB
Disks	4	EBS Used	1.13 GiB
Volumes	2	S3 Used	0 GiB
Elastic Volumes	Enabled		
S3 Tiering	Enabled		

5. Manage your aggregates:

Task	Action
View information about an aggregate	Under the ... (ellipses icon) menu, click <b>View aggregate details</b> .

Task	Action
Create a volume on a specific aggregate	Under the ... (ellipses icon) menu, click <b>Add volume</b> .
Add disks to an aggregate	<p>a. Under the ... (ellipses icon) menu, click <b>Add disks</b>.</p> <p>b. Select the number of disks that you want to add and click <b>Add</b>.</p> <p> All disks in an aggregate must be the same size.</p>
Increase the capacity of an aggregate that supports Amazon EBS Elastic Volumes	<p>a. Under the ... (ellipses icon) menu, click <b>Increase capacity</b>.</p> <p>b. Enter the additional capacity that you'd like to add and then click <b>Increase</b>.</p> <p>Note that you must increase the capacity of the aggregate by a minimum of 256 GiB or 10% of the aggregate's size.</p> <p>For example, if you have a 1.77 TiB aggregate, 10% is 181 GiB. That's lower than 256 GiB, so the size of the aggregate must be increased by the 256 GiB minimum.</p>
Delete an aggregate	<p>a. Select an aggregate tile that does not contain any volumes click the ... (ellipses icon) &gt; <b>Delete</b>.</p> <p>b. Click <b>Delete</b> again to confirm.</p>

## Manage the Cloud Volumes ONTAP aggregate capacity on a Connector

Each Connector has settings that determines how it manages aggregate capacity for Cloud Volumes ONTAP.

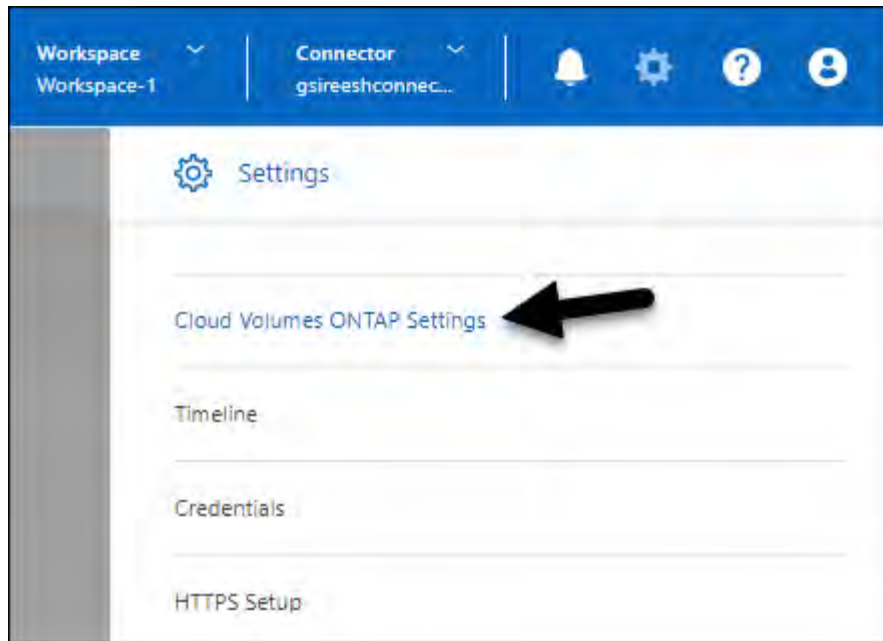
These settings affect all Cloud Volumes ONTAP systems managed by a Connector. If you have another Connector, it can be configured differently.

### Required permissions

BlueXP Organization or Account admin privileges are required to modify Cloud Volumes ONTAP Settings.

### Steps

1. In the upper right of the BlueXP console, click the Settings icon, and select **Cloud Volumes ONTAP Settings**.



2. Under **Capacity**, modify any of the following settings:

### **Capacity Management Mode**

Choose whether BlueXP notifies you of storage capacity decisions or whether BlueXP automatically manages capacity requirements for you.

[Learn how Capacity Management Mode works.](#)

### **Aggregate Capacity Threshold - Free Space Ratio**

This ratio is a key parameter in capacity management decisions, and understanding its impact is essential regardless of whether you are in an automatic or manual mode of capacity management. It is recommended to set this threshold with consideration of your specific storage needs and anticipated growth to maintain a balance between resource utilization and cost.

In the manual mode, if the free space ratio on an aggregate drops below the specified threshold, it triggers a notification, alerting you that you should take actions to address the low free space ratio. It is important to monitor these notifications and manually manage the aggregate capacity to avoid service disruption and ensure optimal performance.

The free space ratio is calculated as follows:

$(\text{aggregate capacity} - \text{total used capacity on the aggregate}) / \text{aggregate capacity}$

Refer to [Automatic capacity management](#) to learn how capacity is automatically managed in Cloud Volumes ONTAP.

### **Aggregate Capacity Thresholds - Free Space Ratio for Data Tiering**

Defines how much free space is required on the performance tier (disks) when tiering data to a capacity tier (object storage).

The ratio is important for disaster recovery scenarios. As data is read from the capacity tier, Cloud Volumes ONTAP moves data to the performance tier to provide better performance. If there isn't sufficient space, then Cloud Volumes ONTAP can't move the data.

3. Click **Save**.

# Storage VM administration

## Manage storage VMs for Cloud Volumes ONTAP

A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*. Cloud Volumes ONTAP is configured with one storage VM by default, but some configurations support additional storage VMs.

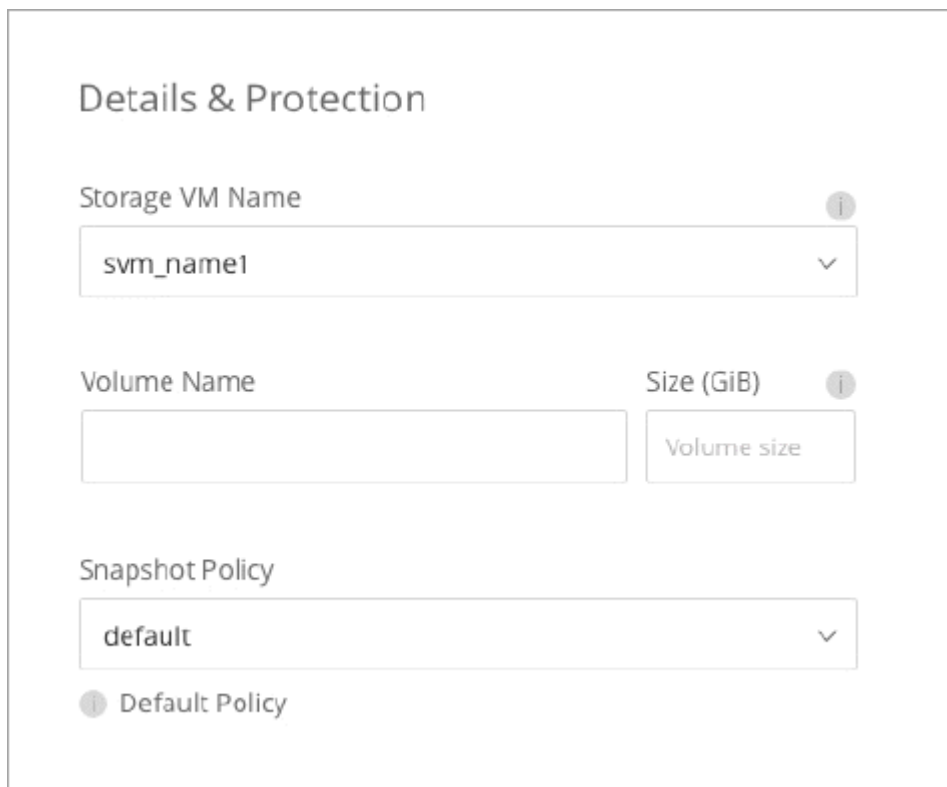
### Supported number of storage VMs

Multiple storage VMs are supported with certain configurations. Go to the [Cloud Volumes ONTAP Release Notes](#) to verify the supported number of storage VMs for your version of Cloud Volumes ONTAP.

### Work with multiple storage VMs

BlueXP supports any additional storage VMs that you create from ONTAP System Manager or the ONTAP CLI.

For example, the following image shows how you can choose a storage VM when you create a volume.



**Details & Protection**

Storage VM Name ⓘ  
svm\_name1 ▼

Volume Name      Size (GiB) ⓘ  
     

Snapshot Policy  
default ▼

ⓘ Default Policy

And the following image shows how you can choose a storage VM when replicating a volume to another system.

Destination Volume Name  
volume\_copy

Destination Storage VM Name  
svm\_name1

Destination Aggregate  
Automatically select the best aggregate

### Modify the name of the default storage VM

BlueXP automatically names the single storage VM that it creates for Cloud Volumes ONTAP. From ONTAP System Manager, the ONTAP CLI, or API, you can modify the name of the storage VM if you have strict naming standards. For example, you might want the name to match how you name the storage VMs for your ONTAP clusters.

### Manage data-serving storage VMs for Cloud Volumes ONTAP in AWS

A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*. Cloud Volumes ONTAP is configured with one storage VM by default, but some configurations support additional storage VMs.

To create additional data-serving storage VMs, you need to allocate IP addresses in AWS and then run ONTAP commands based on your Cloud Volumes ONTAP configuration.

### Supported number of storage VMs

Multiple storage VMs are supported with specific Cloud Volumes ONTAP configurations starting with the 9.7 release. Go to the [Cloud Volumes ONTAP Release Notes](#) to verify the supported number of storage VMs for your version of Cloud Volumes ONTAP.

All other Cloud Volumes ONTAP configurations support one data-serving storage VM and one destination storage VM used for disaster recovery. You can activate the destination storage VM for data access if there's an outage on the source storage VM.

### Verify limits for your configuration

Each EC2 instance supports a maximum number of private IPv4 addresses per network interface. You need to verify the limit before you allocate IP addresses in AWS for the new storage VM.

### Steps

1. Go to the [Storage limits section in the Cloud Volumes ONTAP Release Notes](#).



2. Identify the maximum number of IP addresses per interface for your instance type.
3. Make note of this number because you'll need it in the next section when you allocate IP addresses in AWS.

## Allocate IP addresses in AWS

Private IPv4 addresses must be assigned to port e0a in AWS before you create LIFs for the new storage VM.

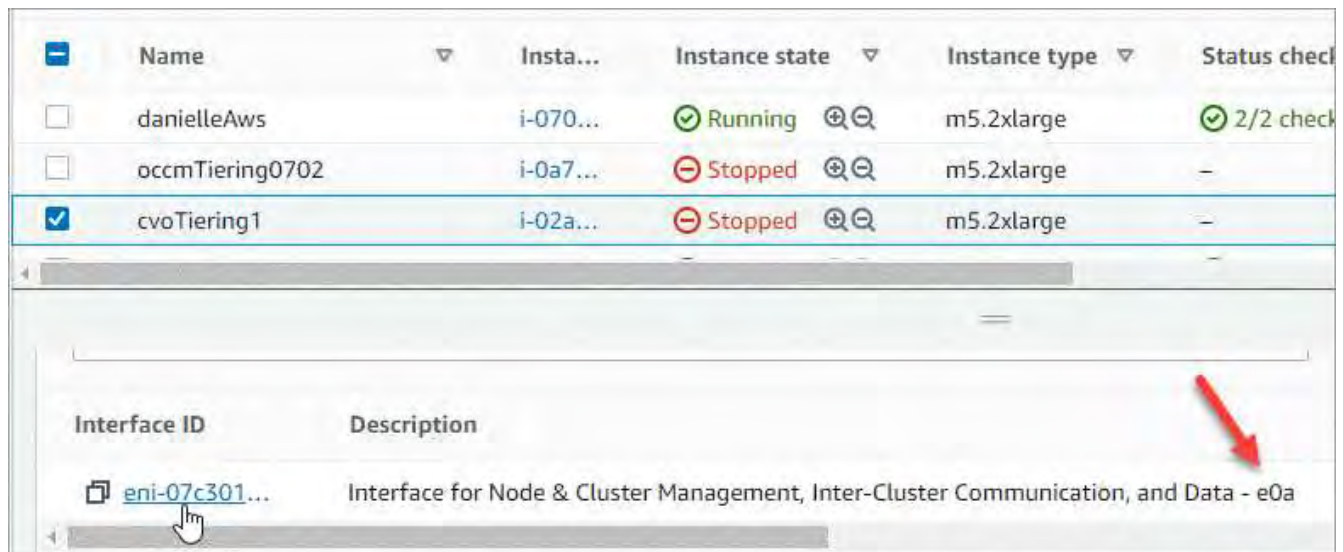
Note that an optional management LIF for a storage VM requires a private IP address on a single node system and on an HA pair in a single AZ. This management LIF provides a connection to management tools like SnapCenter.

### Steps

1. Log in to AWS and open the EC2 service.
2. Select the Cloud Volumes ONTAP instance and click **Networking**.

If you're creating a storage VM on an HA pair, select node 1.

3. Scroll down to **Network interfaces** and click the **Interface ID** for port e0a.



4. Select the network interface and click **Actions > Manage IP addresses**.
5. Expand the list of IP addresses for e0a.
6. Verify the IP addresses:
  - a. Count the number of allocated IP addresses to confirm that the port has room for additional IPs.

You should have identified the maximum number of supported IP addresses per interface in the previous section of this page.
  - b. Optional: Go to the ONTAP CLI for Cloud Volumes ONTAP and run **network interface show** to confirm that each of these IP addresses are in use.

If an IP address isn't in use, then you can use it with the new storage VM.

7. Back in the AWS Console, click **Assign new IP address** to assign additional IP addresses based on the amount that you need for the new storage VM.

- Single node system: One unused secondary private IP is required.

An optional secondary private IP is required if you want to create a management LIF on the storage VM.

- HA pair in a single AZ: One unused secondary private IP is required on node 1.

An optional secondary private IP is required if you want to create a management LIF on the storage VM.

- HA pair in multiple AZs: One unused secondary private IP is required on each node.

8. If you're allocating the IP address on an HA pair in a single AZ, enable **Allow secondary private IPv4 addresses to be reassigned**.

9. Click **Save**.

10. If you have an HA pair in multiple AZs, then you'll need to repeat these steps for node 2.

### Create a storage VM on a single node system

These steps create a new storage VM on a single node system. One private IP address is required to create a NAS LIF and another optional private IP address is needed if you want to create a management LIF.

#### Steps

1. Create the storage VM and a route to the storage VM.

```
vserver create -rootvolume-security-style unix -rootvolume root_svm_2  
-snapshot-policy default -vserver svm_2 -aggregate aggr1
```

```
network route create -destination 0.0.0.0/0 -vserver svm_2 -gateway  
subnet_gateway
```

2. Create a NAS LIF.

```
network interface create -auto-revert true -vserver svm_2 -service  
-policy default-data-files -home-port e0a -address private_ip_x -netmask  
node1Mask -lif ip_nas_2 -home-node cvo-node
```

Where *private\_ip\_x* is an unused secondary private IP on e0a.

3. Optional: Create a storage VM management LIF.

```
network interface create -auto-revert true -vserver svm_2 -service  
-policy default-management -home-port e0a -address private_ip_y -netmask  
node1Mask -lif ip_svm_mgmt_2 -home-node cvo-node
```

Where *private\_ip\_y* is another unused secondary private IP on e0a.

#### 4. Assign one or more aggregates to the storage VM.

```
vserver add-aggregates -vserver svm_2 -aggregates aggr1,aggr2
```

This step is required because the new storage VM needs access to at least one aggregate before you can create volumes on the storage VM.

### Create a storage VM on an HA pair in a single AZ

These steps create a new storage VM on an HA pair in a single AZ. One private IP address is required to create a NAS LIF and another optional private IP address is needed if you want to create a management LIF.

Both of these LIFs get allocated on node 1. The private IP addresses can move between nodes if failures occur.

#### Steps

##### 1. Create the storage VM and a route to the storage VM.

```
vserver create -rootvolume-security-style unix -rootvolume root_svm_2  
-snapshot-policy default -vserver svm_2 -aggregate aggr1
```

```
network route create -destination 0.0.0.0/0 -vserver svm_2 -gateway  
subnet_gateway
```

##### 2. Create a NAS LIF on node 1.

```
network interface create -auto-revert true -vserver svm_2 -service  
-policy default-data-files -home-port e0a -address private_ip_x -netmask  
node1Mask -lif ip_nas_2 -home-node cvo-node1
```

Where *private\_ip\_x* is an unused secondary private IP on e0a of cvo-node1. This IP address can be relocated to the e0a of cvo-node2 in case of takeover because the service policy default-data-files indicates that IPs can migrate to the partner node.

##### 3. Optional: Create a storage VM management LIF on node 1.

```
network interface create -auto-revert true -vserver svm_2 -service  
-policy default-management -home-port e0a -address private_ip_y -netmask  
node1Mask -lif ip_svm_mgmt_2 -home-node cvo-node1
```

Where *private\_ip\_y* is another unused secondary private IP on e0a.

##### 4. Assign one or more aggregates to the storage VM.

```
vserver add-aggregates -vserver svm_2 -aggregates aggr1,aggr2
```

This step is required because the new storage VM needs access to at least one aggregate before you can create volumes on the storage VM.

5. If you're running Cloud Volumes ONTAP 9.11.1 or later, modify the network service policies for the storage VM.

Modifying the services is required because it ensures that Cloud Volumes ONTAP can use the iSCSI LIF for outbound management connections.

```
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service data-fpolicy-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-ad-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-dns-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-ldap-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-nis-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service data-fpolicy-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-ad-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-dns-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-ldap-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-nis-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service data-fpolicy-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-ad-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-dns-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-ldap-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-nis-client
```

## Create a storage VM on an HA pair in multiple AZs

These steps create a new storage VM on an HA pair in multiple AZs.

A *floating* IP address is required for a NAS LIF and is optional for a management LIF. These floating IP addresses don't require you to allocate private IPs in AWS. Instead, the floating IPs are automatically configured in the AWS route table to point to a specific node's ENI in the same VPC.

In order for floating IPs to work with ONTAP, a private IP address must be configured on every storage VM on each node. This is reflected in the steps below where an iSCSI LIF is created on node 1 and on node 2.

## Steps

1. Create the storage VM and a route to the storage VM.

```
vserver create -rootvolume-security-style unix -rootvolume root_svm_2  
-snapshot-policy default -vserver svm_2 -aggregate aggr1
```

```
network route create -destination 0.0.0.0/0 -vserver svm_2 -gateway  
subnet_gateway
```

2. Create a NAS LIF on node 1.

```
network interface create -auto-revert true -vserver svm_2 -service  
-policy default-data-files -home-port e0a -address floating_ip -netmask  
node1Mask -lif ip_nas_floating_2 -home-node cvo-node1
```

- The floating IP address must be outside of the CIDR blocks for all VPCs in the AWS region in which you deploy the HA configuration. 192.168.209.27 is an example floating IP address. [Learn more about choosing a floating IP address.](#)
- `-service-policy default-data-files` indicates that IPs can migrate to the partner node.

3. Optional: Create a storage VM management LIF on node 1.

```
network interface create -auto-revert true -vserver svm_2 -service  
-policy default-management -home-port e0a -address floating_ip -netmask  
node1Mask -lif ip_svm_mgmt_2 -home-node cvo-node1
```

4. Create an iSCSI LIF on node 1.

```
network interface create -vserver svm_2 -service-policy default-data-  
blocks -home-port e0a -address private_ip -netmask node1Mask -lif  
ip_node1_iscsi_2 -home-node cvo-node1
```

- This iSCSI LIF is required to support LIF migration of the floating IPs in the storage VM. It doesn't have to be an iSCSI LIF, but it can't be configured to migrate between nodes.
- `-service-policy default-data-block` indicates that an IP address does not migrate between nodes.

- *private\_ip* is an unused secondary private IP address on eth0 (e0a) of cvo\_node1.

5. Create an iSCSI LIF on node 2.

```
network interface create -vserver svm_2 -service-policy default-data-blocks -home-port e0a -address private_ip -netmaskNode2Mask -lif ip_node2_iscsi_2 -home-node cvo-node2
```

- This iSCSI LIF is required to support LIF migration of the floating IPs in the storage VM. It doesn't have to be an iSCSI LIF, but it can't be configured to migrate between nodes.
- `-service-policy default-data-block` indicates that an IP address does not migrate between nodes.
- *private\_ip* is an unused secondary private IP address on eth0 (e0a) of cvo\_node2.

6. Assign one or more aggregates to the storage VM.

```
vserver add-aggregates -vserver svm_2 -aggregates aggr1,aggr2
```

This step is required because the new storage VM needs access to at least one aggregate before you can create volumes on the storage VM.

7. If you're running Cloud Volumes ONTAP 9.11.1 or later, modify the network service policies for the storage VM.

Modifying the services is required because it ensures that Cloud Volumes ONTAP can use the iSCSI LIF for outbound management connections.

```

network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service data-fpolicy-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-ad-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-dns-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-ldap-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-nis-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service data-fpolicy-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-ad-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-dns-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-ldap-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-nis-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service data-fpolicy-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-ad-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-dns-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-ldap-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-nis-client

```

## Manage data-serving storage VMs for Cloud Volumes ONTAP in Azure

A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*. Cloud Volumes ONTAP is configured with one storage VM by default, but you can create additional storage VMs when running Cloud Volumes ONTAP in Azure.

To create and manage additional data-serving storage VMs in Azure, you should use the BlueXP APIs. This is because the APIs automate the process of creating the storage VMs and configuring the required network interfaces. When creating the storage VMs, BlueXP configures the required LIF services, as well as an iSCSI LIF that's required for outbound SMB/CIFS communications from the storage VM.

For information about running Cloud Volumes ONTAP API calls, refer to [Your first API call](#).

## Supported number of storage VMs

Beginning with Cloud Volumes ONTAP 9.9.0, based on your license, multiple storage VMs are supported with specific configurations. Refer to the [Cloud Volumes ONTAP Release Notes](#) to verify the supported number of storage VMs for your version of Cloud Volumes ONTAP.

All versions of Cloud Volumes ONTAP prior to 9.9.0 support one data-serving storage VM and one destination storage VM used for disaster recovery. You can activate the destination storage VM for data access if there's an outage on the source storage VM.

## Create a storage VM

Based on your configuration and license type, you can create multiple storage VMs on a single node system or in a high-availability (HA) configuration by using the BlueXP APIs.

### About this task

When you create storage VMs using the APIs, along with configuring the required network interfaces, BlueXP also modifies the `default-data-files` policies on the data storage VMs by removing the following services from the NAS data LIF and adding them to the iSCSI data LIF that's used for outbound management connections:

- `data-fpolicy-client`
- `management-ad-client`
- `management-dns-client`
- `management-ldap-client`
- `management-nis-client`

### Before you begin

The Connector requires specific permissions to create storage VMs for Cloud Volumes ONTAP. The required permissions are included in [the policies provided by NetApp](#).

### Single node system

Use the following API call to create a storage VM on a single node system.

```
POST /azure/vsa/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{ "svmName": "myNewSvm1"  
  "svmPassword": "optional, the API takes the cluster password if not  
provided"  
  "mgmtLif": "optional, to create an additional management LIF, if you  
want to use the storage VM for management purposes"}
```

### HA pair

Use the following API call to create a storage VM on an HA pair:



POST /azure/ha/working-environments/{workingEnvironmentId}/svm

Include the following parameters in the request body:

```
{ "svmName": "NewSvmName"
  "svmPassword": "optional value, the API takes the cluster password if
not provided"
  "mgmtLif": "optional value, to create an additional management LIF, if
you want to use the storage VM for management purposes"}
```

## Manage storage VMs on single node systems and HA pairs

Using the BlueXP APIs, you can rename and delete storage VMs in both single node and HA configurations.

### Before you begin

The Connector requires specific permissions to manage storage VMs for Cloud Volumes ONTAP. The required permissions are included in [the policies provided by NetApp](#).

### Rename a storage VM

To rename a storage VM, you should provide the names of the existing storage VM and new storage VM as parameters.

### Steps

- Use the following API call to rename a storage VM on a single node system:

```
PUT /azure/vsa/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{
  "svmNewName": "NewSvmName",
  "svmName": "OldSvmName"
}
```

- Use the following API call to rename a storage VM on an HA pair:

```
PUT /azure/ha/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{
  "svmNewName": "NewSvmName",
  "svmName": "OldSvmName"
}
```

## Delete a storage VM

In a single node or HA configuration, you can remove a storage VM if it doesn't have any active volumes.

### Steps

- Use the following API call to delete a storage VM on a single node system:

```
DELETE /azure/vsa/working-environments/{workingEnvironmentId}/svm/{svmName}
```

- Use the following API call to delete a storage VM on an HA pair:

```
DELETE /azure/ha/working-environments/{workingEnvironmentId}/svm/{svmName}
```

### Related information

- [Prepare to use the API](#)
- [Cloud Volumes ONTAP workflows](#)
- [Get required identifiers](#)
- [Use the BlueXP REST APIs](#)

## Manage data-serving storage VMs for Cloud Volumes ONTAP in Google Cloud

A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*. Cloud Volumes ONTAP is configured with one storage VM by default, but some configurations support additional storage VMs.

To create and manage additional data-serving storage VMs in Google Cloud, you should use the BlueXP APIs. This is because the APIs automate the process of creating the storage VMs and configuring the required network interfaces. When creating the storage VMs, BlueXP configures the required LIF services, as well as an iSCSI LIF that's required for outbound SMB/CIFS communications from the storage VM.

For information about running Cloud Volumes ONTAP API calls, refer to [Your first API call](#).

### Supported number of storage VMs

Beginning with Cloud Volumes ONTAP 9.11.1, based on your license, multiple storage VMs are supported with specific configurations. Refer to the [Cloud Volumes ONTAP Release Notes](#) to verify the supported number of storage VMs for your version of Cloud Volumes ONTAP.

All versions of Cloud Volumes ONTAP prior to 9.11.1 support one data-serving storage VM and one destination storage VM used for disaster recovery. You can activate the destination storage VM for data access if there's an outage on the source storage VM.

### Create a storage VM

Based on your configuration and license type, you can create multiple storage VMs on a single node system or in a high-availability (HA) configuration by using the BlueXP APIs.

### About this task

When you create storage VMs using the APIs, along with configuring the required network interfaces, BlueXP also modifies the `default-data-files` policies on the data storage VMs by removing the following services

from the NAS data LIF and adding them to the iSCSI data LIF that's used for outbound management connections:

- data-fpolicy-client
- management-ad-client
- management-dns-client
- management-ldap-client
- management-nis-client

### Before you begin

The Connector requires specific permissions to create storage VMs for Cloud Volumes ONTAP HA pairs. The required permissions are included in [the policies provided by NetApp](#).

### Single node system

Use the following API call to create a storage VM on a single node system.

```
POST /gcp/vsa/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{ "svmName": "NewSvmName"
  "svmPassword": "optional value, the API takes the cluster password if
not provided"
  "mgmtLif": "optional value, to create an additional management LIF, if
you want to use the storage VM for management purposes" }
```

### HA pair

Use the following API call to create a storage VM on an HA pair:

```
POST /gcp/ha/working-environments/{workingEnvironmentId}/svm/
```

Include the following parameters in the request body:

```
{ "svmName": "NewSvmName"
  "svmPassword": "optional value, the API takes the cluster password if
not provided"
}
```

### Manage storage VMs

Using the BlueXP APIs, you can rename and delete storage VMs in both single node and HA configurations.

### Before you begin

The Connector requires specific permissions to manage storage VMs for Cloud Volumes ONTAP HA pairs. The required permissions are included in [the policies provided by NetApp](#).

## Rename a storage VM

To rename a storage VM, you should provide the names of the existing storage VM and new storage VM as parameters.

### Steps

- Use the following API call to rename a storage VM on a single node system:

```
PUT /gcp/vsa/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{
  "svmNewName": "NewSvmName",
  "svmName": "OldSvmName"
}
```

- Use the following API call to rename a storage VM on an HA pair:

```
PUT /gcp/ha/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{
  "svmNewName": "NewSvmName",
  "svmName": "OldSvmName"
}
```

## Delete a storage VM

In a single node or HA configuration, you can remove a storage VM if it doesn't have any active volumes.

### Steps

- Use the following API call to delete a storage VM on a single node system:

```
DELETE /gcp/vsa/working-environments/{workingEnvironmentId}/svm/{svmName}
```

- Use the following API call to delete a storage VM on an HA pair:

```
DELETE /gcp/ha/working-environments/{workingEnvironmentId}/svm/{svmName}
```

### Related information

- [Prepare to use the API](#)
- [Cloud Volumes ONTAP workflows](#)
- [Get required identifiers](#)
- [Use the BlueXP REST APIs](#)

## Set up storage VM disaster recovery for Cloud Volumes ONTAP

BlueXP does not offer setup or orchestration support for storage VM (SVM) disaster recovery. To perform these tasks, use ONTAP System Manager or the ONTAP CLI.

If you set up SnapMirror SVM replication between two Cloud Volumes ONTAP systems, the replication must be between two HA pair systems or two single node systems. You can't set up SnapMirror SVM replication between an HA pair and a single node system.

Refer to the following documents for the ONTAP CLI instructions.

- [SVM Disaster Recovery Preparation Express Guide](#)
- [SVM Disaster Recovery Express Guide](#)

## Security and data encryption

### Encrypt volumes on Cloud Volumes ONTAP with NetApp encryption solutions

Cloud Volumes ONTAP supports NetApp Volume Encryption (NVE) and NetApp Aggregate Encryption (NAE). NVE and NAE are software-based solutions that enable FIPS 140-2–compliant data-at-rest encryption of volumes. [Learn more about these encryption solutions.](#)

Both NVE and NAE are supported with an external key manager.

If you use NVE, you have the option to use your cloud provider's key vault to protect ONTAP encryption keys:

- AWS Key Management Service (beginning in 9.12.0)
- Azure Key Vault (AKV)
- Google Cloud Key Management Service

New aggregates will have NAE enabled by default after you set up an external key manager. New volumes that aren't part of an NAE aggregate will have NVE enabled by default (for example, if you have existing aggregates that were created before setting up an external key manager).

Cloud Volumes ONTAP doesn't support onboard key management.

### Before you begin

Your Cloud Volumes ONTAP system should be registered with NetApp Support. A NetApp Volume Encryption license is automatically installed on each Cloud Volumes ONTAP system that is registered with NetApp Support.

- [Adding NetApp Support Site accounts to BlueXP](#)
- [Register pay-as-you-go systems](#)



BlueXP doesn't install the NVE license on systems that reside in the China region.

### Steps

1. Review the list of supported key managers in the [NetApp Interoperability Matrix Tool](#).



Search for the **Key Managers** solution.

2. [Connect to the Cloud Volumes ONTAP CLI](#).
3. Configure external key management.
  - AWS: [AWS Key Management Service](#)
  - Azure: [Azure Key Vault \(AKV\)](#)
  - Google Cloud: [Google Cloud Key Management Service](#)

## Manage Cloud Volumes ONTAP encryption keys with AWS Key Management Service

You can use [AWS's Key Management Service \(KMS\)](#) to protect your ONTAP encryption keys in an AWS-deployed application.

Key management with the AWS KMS can be enabled with the CLI or the ONTAP REST API.

When using the KMS, be aware that by default a data SVM's LIF is used to communicate with the cloud key management endpoint. A node management network is used to communicate with AWS's authentication services. If the cluster network is not configured correctly, the cluster will not properly utilize the key management service.

### Before you begin

- Cloud Volumes ONTAP must be running version 9.12.0 or later
- You must have installed the Volume Encryption (VE) license and
- You must have installed the Multi-tenant Encryption Key Management (MTEKM) license installed.
- You must be a cluster or SVM administrator
- You must have an active AWS subscription



You can only configure keys for a data SVM.

### Configuration

#### AWS

1. You must create a [grant](#) for the AWS KMS key that will be used by the IAM role managing encryption. The IAM role must include a policy that allows the following operations:
  - DescribeKey
  - Encrypt
  - DecryptTo create a grant, refer to [AWS documentation](#).
2. [Add a policy to the appropriate IAM role](#). The policy should support the DescribeKey, Encrypt, and Decrypt operations.

#### Cloud Volumes ONTAP

1. Switch to your Cloud Volumes ONTAP environment.
2. Switch to the advanced privilege level:

```
set -privilege advanced
```

3. Enable the AWS key manager:

```
security key-manager external aws enable -vserver data_svm_name -region  
AWS_region -key-id key_ID -encryption-context encryption_context
```

4. When prompted, enter the secret key.

5. Confirm the AWS KMS was configured correctly:

```
security key-manager external aws show -vserver svm_name
```

## Manage Cloud Volumes ONTAP encryption keys with Azure Key Vault

You can use Azure Key Vault (AKV) to protect your ONTAP encryption keys in an Azure-deployed application. Refer to the [Microsoft documentation](#).

AKV can be used to protect NetApp Volume Encryption (NVE) keys only for data SVMs. For more information, refer to the [ONTAP documentation](#).

Key management with AKV can be enabled with the CLI or the ONTAP REST API.

When using AKV, be aware that by default a data SVM LIF is used to communicate with the cloud key management endpoint. A node management network is used to communicate with the cloud provider's authentication services (login.microsoftonline.com). If the cluster network is not configured correctly, the cluster will not properly utilize the key management service.

### Before you begin

- Cloud Volumes ONTAP must be running version 9.10.1 or later
- Volume Encryption (VE) license installed (NetApp Volume Encryption license is automatically installed on each Cloud Volumes ONTAP system that is registered with NetApp Support)
- You must have a Multi-tenant Encryption Key Management (MT\_EK\_MGMT) license
- You must be a cluster or SVM administrator
- An Active Azure subscription

### Limitations

- AKV can only be configured on a data SVM
- NAE can't be used using AKV. NAE requires an external-supported KMIP server.
- Cloud Volumes ONTAP nodes poll AKV every 15 minutes to confirm accessibility and key availability. This polling period is non-configurable, and after four consecutive failures in the polling attempt (totaling 1 hour), the volumes are placed offline.

### Configuration process

The outlined steps capture how to register your Cloud Volumes ONTAP configuration with Azure and how to create an Azure Key Vault and keys. If you have already completed these steps, ensure you have the correct configuration settings, particularly in [Create an Azure Key Vault](#), and then proceed to [Cloud Volumes ONTAP configuration](#).

- [Azure Application Registration](#)
- [Create Azure client secret](#)
- [Create an Azure Key Vault](#)

- [Create encryption key](#)
- [Create an Azure Active Directory Endpoint \(HA only\)](#)
- [Cloud Volumes ONTAP configuration](#)

### Azure Application Registration

1. You must first register your application in the Azure subscription that you want the Cloud Volumes ONTAP to use for access the Azure Key Vault. Within the Azure portal, select **App registrations**.
2. Select **New registration**.
3. Provide a name for your application and select a supported application type. The default single tenant suffices for Azure Key Vault usage. Select **Register**.
4. In the Azure Overview window, select the application you have registered. Copy the **application (client) ID** and the **directory (tenant) ID** to a secure location. They will be required later in the registration process.

### Create Azure client secret

1. In the Azure portal for your Azure Key Vault app registration, select the **Certificates & secrets** pane.
2. Select **New client secret**. Enter a meaningful name for your client secret. NetApp recommends a 24-month expiration period; however, your specific cloud governance policies may require a different setting.
3. Click **Add** to create the client secret. Copy the secret string listed in the **Value** column and store it in a secure location for use later in [Cloud Volumes ONTAP configuration](#). The secret value will not be displayed again after you navigate away from the page.

### Create an Azure Key Vault

1. If you have an existing Azure Key Vault, you can connect it to your Cloud Volumes ONTAP configuration; however, you must adapt the access policies to the settings in this process.
2. In the Azure portal, navigate to the **Key Vaults** section.
3. Click **+Create** and enter the required information including resource group, region, and pricing tier. In addition, enter the number of days to retain deleted vaults and select **Enable purge protection** on the key vault.
4. Select **Next** to choose an access policy.
5. Select the following options:
  - a. Under **Access configuration**, select the **Vault access policy**.
  - b. Under **Resource access**, select **Azure Disk Encryption for volume encryption**.
6. Select **+Create** to add an access policy.
7. Under **Configure from a template**, click the drop-down menu and then select the **Key, Secret, and Certificate Management** template.
8. Choose each of the drop-down permissions menus (key, secret, certificate) and then **Select all** at the top of the menu list to select all the permissions available. You should have:
  - **Key permissions:** 20 selected
  - **Secret permissions:** 8 selected
  - **Certificate permissions:** 16 selected



# Create an access policy



- 1 **Permissions**   2 Principal   3 Application (optional)   4 Review + create

Configure from a template

Key, Secret, & Certificate Management

## Key permissions

### Key Management Operations

- Select all
- Get
- List
- Update
- Create
- Import
- Delete
- Recover
- Backup
- Restore

### Cryptographic Operations

- Select all
- Decrypt
- Encrypt
- Unwrap Key
- Wrap Key
- Verify
- Sign

### Privileged Key Operations

- Select all
- Purge
- Release

### Rotation Policy Operations

- Select all
- Rotate
- Get Rotation Policy
- Set Rotation Policy

## Secret permissions

### Secret Management Operations

- Select all
- Get
- List
- Set
- Delete
- Recover
- Backup
- Restore

### Privileged Secret Operations

- Select all
- Purge

## Certificate permissions

### Certificate Management Operations

- Select all
- Get
- List
- Update
- Create
- Import
- Delete
- Recover
- Backup
- Restore
- Manage Contacts
- Manage Certificate Authorities
- Get Certificate Authorities
- List Certificate Authorities
- Set Certificate Authorities
- Delete Certificate Authorities

### Privileged Certificate Operations

- Select all
- Purge

Previous

Next

9. Click **Next** to select the **Principal** Azure registered application you created in [Azure Application Registration](#). Select **Next**.



Only one principal can be assigned per policy.

**Create an access policy**

Permissions **Principal** Application (optional) Review + create

Only 1 principal can be assigned per access policy. Use the new embedded experience to select a principal. The previous popup experience can be accessed here. [Select a principal](#)

Search by object ID, name, or email address

**Selected item**

No item selected

Previous Next

10. Click **Next** two times until you arrive at **Review and create**. Then, click **Create**.
11. Select **Next** to advance to **Networking** options.
12. Choose the appropriate network access method or select **All networks** and **Review + Create** to create the key vault. (Network access method may be prescribed by a governance policy or your corporate cloud security team.)
13. Record the Key Vault URI: In the key vault you created, navigate to the Overview menu and copy the **Vault URI** from the right-hand column. You need this for a later step.

### Create encryption key

1. In the menu for the Key Vault you have created for Cloud Volumes ONTAP, navigate to the **Keys** option.
2. Select **Generate/import** to create a new key.
3. Leave the default option set to **Generate**.
4. Provide the following information:
  - Encryption key name

- Key type: RSA
  - RSA key size: 2048
  - Enabled: Yes
5. Select **Create** to create the encryption key.
  6. Return to the **Keys** menu and select the key you just created.
  7. Select the key ID under **Current version** to view the key properties.
  8. Locate the **Key Identifier** field. Copy the URI up to but not including the hexadecimal string.

#### **Create an Azure Active Directory Endpoint (HA only)**




1. This process is only required if you are configuring Azure Key Vault for an HA Cloud Volumes ONTAP Working Environment.
2. In the Azure portal navigate to **Virtual Networks**.
3. Select the Virtual Network where you deployed the Cloud Volumes ONTAP working environment and select the **Subnets** menu on the left side of the page.
4. Select the subnet name for you Cloud Volumes ONTAP deployment from the list.
5. Navigate to the **Service Endpoints** heading. In the drop-down menu, select the following:
  - **Microsoft.AzureActiveDirectory**
  - **Microsoft.KeyVault**
  - **Microsoft.Storage** (optional)

**SERVICE ENDPOINTS**

Create service endpoint policies to allow traffic to specific azure resources from your virtual network over service endpoints. [Learn more](#)

Services ⓘ

3 selected

Service	Status	
Microsoft.Storage	Succeeded	
Microsoft.AzureActiveDirectory	Succeeded	
Microsoft.KeyVault	Succeeded	

Service endpoint policies

0 selected

**SUBNET DELEGATION**

Delegate subnet to a service ⓘ

None

**NETWORK POLICY FOR PRIVATE ENDPOINTS**

The network policy affects all private endpoints in this subnet. To use network security groups, application security groups, or user defined routes to control traffic going to a private endpoint, set the private endpoint network policy to enabled. [Learn more](#)

Private endpoint network policy

Disabled

**Save** **Cancel**

6. Select **Save** to capture your settings.

### Cloud Volumes ONTAP configuration

1. Connect to the cluster management LIF with your preferred SSH client.
2. Enter the advanced privilege mode in ONTAP:

```
set advanced -con off
```

3. Identify the desired data SVM and verify its DNS configuration:

```
vserver services name-service dns show
```

- a. If a DNS entry for the desired data SVM exists and it contains an entry for the Azure DNS, then no action is required. If it does not, add a DNS server entry for the data SVM that points to the Azure DNS, private DNS, or on-premise server. This should match the entry for the cluster admin SVM:

```
vserver services name-service dns create -vserver SVM_name -domains domain  
-name-servers IP_address
```

- b. Verify the DNS service has been created for the data SVM:

```
vserver services name-service dns show
```

4. Enable Azure Key Vault using the client ID and tenant ID saved after the application registration:

```
security key-manager external azure enable -vserver SVM_name -client-id  
Azure_client_ID -tenant-id Azure_tenant_ID -name key_vault_URI -key-id  
full_key_URI
```



The `_full_key_URI` value must utilize the `<https:// <key vault host name>/keys/<key label>` format.

5. Upon successful enablement of the Azure Key Vault, enter the `client secret` value when prompted.

6. Check the status of the key manager:

```
security key-manager external azure check
```

The output will look like:

```
::*> security key-manager external azure check  
  
Vserver: data_svm_name  
Node: akvlab01-01  
  
Category: service_reachability  
Status: OK  
  
Category: ekvip_server  
Status: OK  
  
Category: kms_wrapped_key_status  
Status: UNKNOWN  
Details: No volumes created yet for the vserver. Wrapped KEK status  
will be available after creating encrypted volumes.  
  
3 entries were displayed.
```

If the `service_reachability` status is not `OK`, the SVM cannot reach the Azure Key Vault service with all the required connectivity and permissions. Ensure that your Azure network policies and routing don't block your private vNet from reaching the Azure Key Vault Public endpoint. If they do, consider using an Azure Private endpoint to access the Key vault from within the vNet. You may also need to add a static hosts entry on your SVM to resolve the private IP address for your endpoint.

The `kms_wrapped_key_status` will report UNKNOWN at initial configuration. Its status will change to OK after the first volume is encrypted.

7. OPTIONAL: Create a test volume to verify the functionality of NVE.

```
vol create -vserver SVM_name -volume volume_name -aggregate aggr -size size
-state online -policy default
```

If configured correctly, Cloud Volumes ONTAP will automatically create the volume and enable volume encryption.

8. Confirm the volume was created and encrypted correctly. If it is, the `-is-encrypted` parameter will display as `true`.

```
vol show -vserver SVM_name -fields is-encrypted
```

9. Optional: If you want to update the credentials on the Azure Key Vault authentication certificate, use the following command:

```
security key-manager external azure update-credentials -vserver v1
-authentication-method certificate
```

#### Related links

- [Set up Cloud Volumes ONTAP to use a customer-managed key in Azure](#)
- [Microsoft Azure documentation: About Azure Key Vault](#)
- [ONTAP command reference guide](#)

## Manage Cloud Volumes ONTAP encryption keys with Google Cloud KMS

You can use [Google Cloud Platform's Key Management Service \(Cloud KMS\)](#) to protect your Cloud Volumes ONTAP encryption keys in a Google Cloud Platform-deployed application.

Key management with Cloud KMS can be enabled with the ONTAP CLI or the ONTAP REST API.

When using Cloud KMS, be aware that by default a data SVM's LIF is used to communicate with the cloud key management endpoint. A node management network is used to communicate with the cloud provider's authentication services (`oauth2.googleapis.com`). If the cluster network is not configured correctly, the cluster will not properly utilize the key management service.

#### Before you begin

- Your system should be running Cloud Volumes ONTAP 9.10.1 or later
- You must use a data SVM. Cloud KMS can be configured only on a data SVM.
- You must be a cluster or SVM administrator
- Volume Encryption (VE) license should be installed on the SVM
- Beginning with Cloud Volumes ONTAP 9.12.1 GA, the multi-tenant Encryption Key Management (MTEKM) license should also be installed
- An active Google Cloud Platform subscription is required

#### Configuration

##### Google Cloud

1. In your Google Cloud environment, [create a symmetric GCP key ring and key](#).
2. Assign a custom role to the Cloud KMS key and Cloud Volumes ONTAP service account.
  - a. Create the custom role:

```
gcloud iam roles create kmsCustomRole
  --project=<project_id>
  --title=<kms_custom_role_name>
  --description=<custom_role_description>

  --permissions=cloudkms.cryptoKeyVersions.get,cloudkms.cryptoKeyVersions.list,cloudkms.cryptoKeyVersions.useToDecrypt,cloudkms.cryptoKeyVersions.useToEncrypt,cloudkms.cryptoKeys.get,cloudkms.keyRings.get,cloudkms.locations.get,cloudkms.locations.list,resourceManager.projects.get
  --stage=GA
```

- b. Assign the custom role you created:

```
gcloud kms keys add-iam-policy-binding key_name --keyring key_ring_name
  --location key_location --member serviceAccount:service_account_name
  --role projects/customer_project_id/roles/kmsCustomRole
```



If you are on Cloud Volumes ONTAP 9.13.0 or later, you don't need to create a custom role. You can assign the predefined `cloudkms.cryptoKeyEncrypterDecrypter` role.

3. Download service account JSON key:

```
gcloud iam service-accounts keys create key-file --iam-account=sa-name
  @project-id.iam.gserviceaccount.com
```

## Cloud Volumes ONTAP

1. Connect to the cluster management LIF with your preferred SSH client.
2. Switch to the advanced privilege level:
 

```
set -privilege advanced
```
3. Create a DNS for the data SVM.
 

```
dns create -domains c.<project>.internal -name-servers server_address -vserver SVM_name
```
4. Create CMEK entry:
 

```
security key-manager external gcp enable -vserver SVM_name -project-id project
  -key-ring-name key_ring_name -key-ring-location key_ring_location -key-name key_name
```
5. When prompted, enter the service account JSON key from your GCP account.
6. Confirm the enabled process succeeded:
 

```
security key-manager external gcp check -vserver svm_name
```
7. OPTIONAL: Create a volume to test encryption
 

```
vol create volume_name -aggregate aggregate
  -vserver vserver_name -size 10G
```

## Troubleshoot

If you need to troubleshoot, you can tail the raw REST API logs in the final two steps above:

1. `set d`
2. `systemshell -node node -command tail -f /mroot/etc/log/mlog/kmip2_client.log`

## Enable NetApp ransomware protection solutions for Cloud Volumes ONTAP

Ransomware attacks can cost a business time, resources, and reputation. BlueXP enables you to implement two NetApp solutions for ransomware: Protection from common ransomware file extensions and Autonomous Ransomware Protection (ARP). These solutions provide effective tools for visibility, detection, and remediation.









### Protection from common ransomware file extensions

Available through BlueXP, the Ransomware Protection setting allows you to utilize the ONTAP FPolicy functionality to guard against common ransomware file extension types.

### Steps

1. On the Canvas page, double-click the name of the system you configure to ransomware protection.
2. On the Overview tab, click the Features panel and then click the pencil icon next to **Ransomware Protection**.



Information		Features
Working Environment Tags		Tags 
Scheduled Downtime		Off 
S3 Storage Classes	Standard-Infrequent Access	
Instance Type	m5.xlarge	
Write Speed	Normal	
Ransomware Protection	Off	
Support Registration	Not Registered	
CIFs Setup		

3. Implement the NetApp solution for ransomware:

- a. Click **Activate Snapshot Policy**, if you have volumes that do not have a Snapshot policy enabled.

NetApp Snapshot technology provides the industry's best solution for ransomware remediation. The key to a successful recovery is restoring from uninfected backups. Snapshot copies are read-only, which prevents ransomware corruption. They can also provide the granularity to create images of a single file copy or a complete disaster recovery solution.

- b. Click **Activate FPolicy** to enable ONTAP's FPolicy solution, which can block file operations based on a file's extension.

This preventative solution improves protection from ransomware attacks by blocking common ransomware file types.

The default FPolicy scope blocks files that have the following extensions:

micro, encrypted, locked, crypto, crypt, crinf, r5a, XRNT, XTBL, R16M01D05, pzdc, good, LOL!, OMG!, RDM, RRK, encryptedRS, crjoker, EnCiPhErEd, LeChiffre



BlueXP creates this scope when you activate FPolicy on Cloud Volumes ONTAP. The list is based on common ransomware file types. You can customize the blocked file extensions by using the `vserver fpolicy policy scope` commands from the Cloud Volumes ONTAP CLI.

## Autonomous Ransomware Protection

Cloud Volumes ONTAP supports the Autonomous Ransomware Protection (ARP) feature, which performs analyses on workloads to proactively detect and warn about abnormal activity that might indicate a ransomware attack.

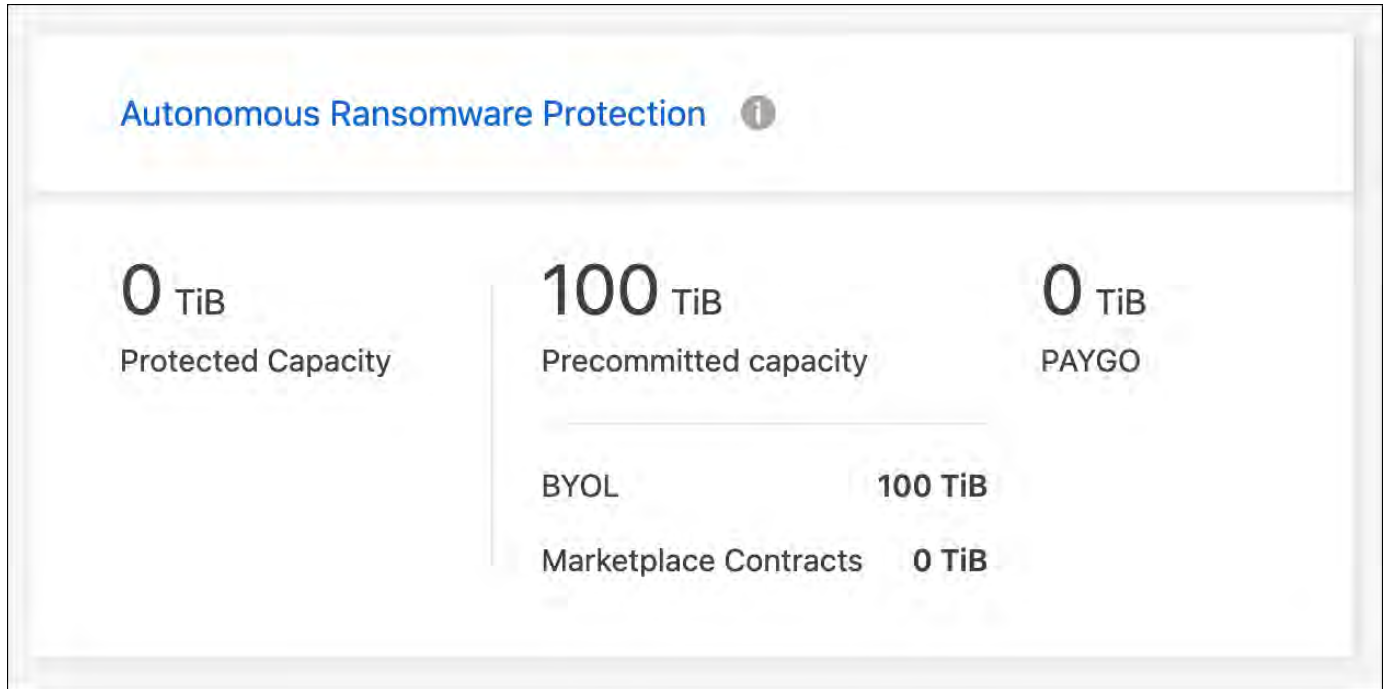
Separate from the file extension protections provided through the [ransomware protection setting](#), the ARP feature uses workload analysis to alert the user on potential attacks based on detected “abnormal activity”. Both the ransomware protection setting and the ARP feature can be used in conjunction for comprehensive ransomware protection.

The ARP feature is available for use with bring your own license (BYOL) and marketplace subscriptions for your licenses at no additional cost.

ARP-enabled volumes have a designated state of "Learning mode" or "Active".

Configuration of ARP for volumes is performed through ONTAP System Manager and ONTAP CLI.

For more information on how to enable ARP with ONTAP System Manager and the ONTAP CLI, refer to the [ONTAP documentation: Enable Autonomous Ransomware Protection](#).



## Create tamperproof Snapshot copies of WORM files on Cloud Volumes ONTAP

You can create tamperproof Snapshot copies of write once, read many (WORM) files on a Cloud Volumes ONTAP system and retain the snapshots in unmodified form for a specific retention period. This functionality is powered by the SnapLock technology, and provides an additional layer of data protection and compliance.

### Before you begin

Ensure that the volume that you use for creating Snapshot copies is a SnapLock volume. For information about enabling SnapLock protection on volumes, refer to the [ONTAP documentation: Configure SnapLock](#).

### Steps

1. Create Snapshot copies from the SnapLock volume. For information about creating Snapshot copies by using the CLI or System Manager, refer to the [ONTAP documentation: Manage local Snapshot copies overview](#).

The Snapshot copies inherit the WORM properties of the volume, making them tamperproof. The underlying SnapLock technology ensures that a snapshot remains protected from edit and deletion until the specified retention period has elapsed.

2. You can modify the retention period if there's a need to edit these snapshots. For information, refer to the [ONTAP documentation: Set the retention time](#).



Even though a Snapshot copy is protected for a specific retention period, the source volume can be deleted by a cluster administrator, as WORM storage in Cloud Volumes ONTAP operates under a "trusted storage administrator" model. Additionally, a trusted cloud administrator can delete the WORM data by operating on the cloud storage resources.

# System administration

## Upgrade Cloud Volumes ONTAP software

Upgrade Cloud Volumes ONTAP from BlueXP to gain access to the latest new features and enhancements. You should prepare Cloud Volumes ONTAP systems before you upgrade the software.

### Upgrade overview

You should be aware of the following before you start the Cloud Volumes ONTAP upgrade process.

#### Upgrade from BlueXP only

Upgrades of Cloud Volumes ONTAP must be completed from BlueXP. You should not upgrade Cloud Volumes ONTAP by using ONTAP System Manager or the ONTAP CLI. Doing so can impact system stability.

### How to upgrade

BlueXP provides two ways to upgrade Cloud Volumes ONTAP:

- By following upgrade notifications that appear in the working environment
- By placing the upgrade image at an HTTPS location and then providing BlueXP with the URL

### Supported upgrade paths

The version of Cloud Volumes ONTAP that you can upgrade to depends on the version of Cloud Volumes ONTAP that you're currently running.

Current version	Versions that you can directly upgrade to
9.15.1	9.16.1 (for Azure and Google Cloud only)
9.15.0	9.15.1
9.14.1	9.15.1
	9.15.0
9.14.0	9.14.1
9.13.1	9.14.1
	9.14.0
9.13.0	9.13.1
9.12.1	9.13.1
	9.13.0
9.12.0	9.12.1
9.11.1	9.12.1
	9.12.0

Current version	Versions that you can directly upgrade to
9.11.0	9.11.1
9.10.1	9.11.1
	9.11.0
9.10.0	9.10.1
9.9.1	9.10.1
	9.10.0
9.9.0	9.9.1
9.8	9.9.1
9.7	9.8
9.6	9.7
9.5	9.6
9.4	9.5
9.3	9.4
9.2	9.3
9.1	9.2
9.0	9.1
8.3	9.0

Note the following:

- The supported upgrade paths for Cloud Volumes ONTAP are different than they are for an on-premises ONTAP cluster.
- If you upgrade by following the upgrade notifications that appear in a working environment, BlueXP will prompt you to upgrade to a release that follows these supported upgrade paths.
- If you upgrade by placing an upgrade image at an HTTPS location, be sure to follow these supported upgrade paths.
- In some cases, you might need to upgrade a few times to reach your target release.

For example, if you're running version 9.8 and you want to upgrade to 9.10.1, you first need to upgrade to version 9.9.1 and then to 9.10.1.

### Patch releases

Starting in January 2024, patch upgrades are only available in BlueXP if there's a patch release for the three latest versions of Cloud Volumes ONTAP. Patch versions are occasionally available for deployment, when the RC or GA version isn't available for deployment.

We use the latest GA release to determine the three latest versions to display in BlueXP. For example, if the current GA release is 9.13.1, patches for 9.11.1-9.13.1 appear in BlueXP. If you want to upgrade to a patch release for versions 9.11.1 or below, you will need to use the manual upgrade procedure by [downloading the](#)

## ONTAP image.

As a general rule for patch (P) releases, you can upgrade from one version release to any P-release of the current version you're running or the next version.

Here are a couple of examples:

- 9.13.0 > 9.13.1P15
- 9.12.1 > 9.13.1P2

### Reverting or downgrading

Reverting or downgrading Cloud Volumes ONTAP to a previous release is not supported.

### Support registration

Cloud Volumes ONTAP must be registered with NetApp Support in order to upgrade the software using any of the methods described on this page. This applies to both pay-as-you-go (PAYGO) and bring your own license (BYOL). You'll need to [manually register PAYGO systems](#), while BYOL systems are registered by default.



A system that isn't registered for support will still receive the software update notifications that appear in BlueXP when a new version is available. But you will need to register the system before you can upgrade the software.

### Upgrades of the HA mediator

BlueXP also updates the mediator instance as needed during the Cloud Volumes ONTAP upgrade process.

### Upgrades in AWS with c4, m4, and r4 EC2 instance types

Cloud Volumes ONTAP no longer supports the c4, m4, and r4 EC2 instance types. You can upgrade existing deployments to Cloud Volumes ONTAP versions 9.8-9.12.1 with these instance types. Before you upgrade we recommend that you [change the instance type](#). If you can't change the instance type, you need to [enable enhanced networking](#) before you upgrade. Read the following sections to learn more about changing the instance type and enabling enhanced networking.

In Cloud Volumes ONTAP running versions 9.13.0 and above, you cannot upgrade with c4, m4, and r4 EC2 instance types. In this case, you need to reduce the number of disks and then [change the instance type](#) or deploy a new HA-pair configuration with the c5, m5, and r5 EC2 instance types and migrate the data.

### Change the instance type

c4, m4, and r4 EC2 instance types allow for more disks per node than the c5, m5, and r5 EC2 instance types. If the disk count per node for the c4, m4, or r4 EC2 instance you're running is below the max disk allowance per node for c5, m5, and r5 instances, you can change the EC2 instance type to c5, m5, or r5.

[Check disk and tiering limits by EC2 instance](#)  
[Change the EC2 instance type for Cloud Volumes ONTAP](#)

If you can't change the instance type, follow the steps in [Enable enhanced networking](#).

### Enable enhanced networking

To upgrade to Cloud Volumes ONTAP versions 9.8 and later, you must enable *enhanced networking* on the

cluster running the c4, m4, or r4 instance type. To enable ENA, refer to the Knowledge Base article "[How to enable Enhanced networking like SR-IOV or ENA on AWS Cloud Volumes ONTAP instances](#)".

## Prepare to upgrade

Before performing an upgrade, you must verify that your systems are ready and make any required configuration changes.

- [Plan for downtime](#)
- [Verify that automatic giveback is still enabled](#)
- [Suspend SnapMirror transfers](#)
- [Verify that aggregates are online](#)
- [Verify that all LIFs are on home ports](#)

### Plan for downtime

When you upgrade a single-node system, the upgrade process takes the system offline for up to 25 minutes, during which I/O is interrupted.

In many cases, upgrading an HA pair is nondisruptive and I/O is uninterrupted. During this nondisruptive upgrade process, each node is upgraded in tandem to continue serving I/O to clients.

Session-oriented protocols might cause adverse effects on clients and applications in certain areas during upgrades. For details, refer to the [ONTAP documentation](#)

### Verify that automatic giveback is still enabled

Automatic giveback must be enabled on a Cloud Volumes ONTAP HA pair (this is the default setting). If it isn't, then the operation will fail.

[ONTAP documentation: Commands for configuring automatic giveback](#)

### Suspend SnapMirror transfers

If a Cloud Volumes ONTAP system has active SnapMirror relationships, it is best to suspend transfers before you update the Cloud Volumes ONTAP software. Suspending the transfers prevents SnapMirror failures. You must suspend the transfers from the destination system.



Even though BlueXP backup and recovery uses an implementation of SnapMirror to create backup files (called SnapMirror Cloud), backups do not need to be suspended when a system is upgraded.

### About this task

These steps describe how to use ONTAP System Manager for version 9.3 and later.

### Steps

1. Log in to System Manager from the destination system.

You can log in to System Manager by pointing your web browser to the IP address of the cluster management LIF. You can find the IP address in the Cloud Volumes ONTAP working environment.



The computer from which you are accessing BlueXP must have a network connection to Cloud Volumes ONTAP. For example, you might need to log in to BlueXP from a jump host that's in your cloud provider network.

2. Click **Protection > Relationships**.
3. Select the relationship and click **Operations > Quiesce**.

#### Verify that aggregates are online

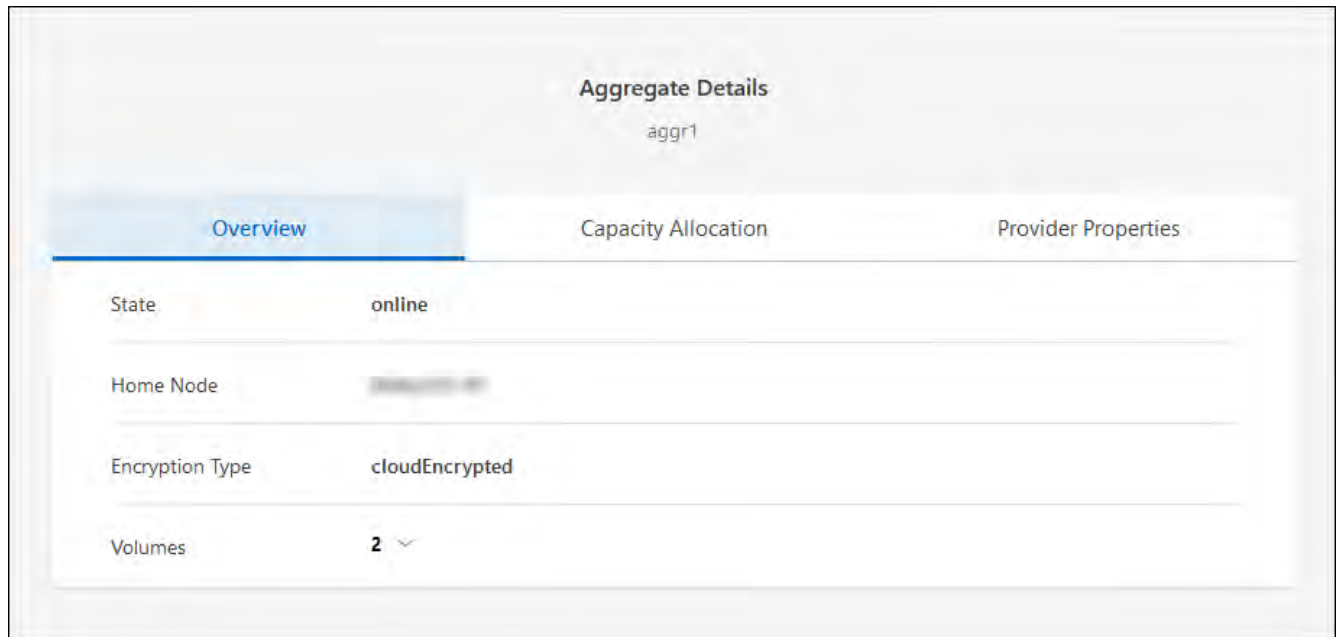
Aggregates for Cloud Volumes ONTAP must be online before you update the software. Aggregates should be online in most configurations, but if they are not, then you should bring them online.

#### About this task

These steps describe how to use ONTAP System Manager for version 9.3 and later.

#### Steps

1. In the working environment, click the **Aggregates** tab.
2. Under the aggregate title, click the ellipses button, and then select **View Aggregate details**.



3. If the aggregate is offline, use System Manager to bring the aggregate online:
  - a. Click **Storage > Aggregates & Disks > Aggregates**.
  - b. Select the aggregate, and then click **More Actions > Status > Online**.

#### Verify that all LIFs are on home ports

Before you upgrade, all LIFs must be on home ports. Refer to the ONTAP documentation to [verify that all LIFs are on home ports](#).

If an upgrade failure error occurs, consult the Knowledge Base (KB) article [Cloud Volumes ONTAP upgrade fails](#).



## Upgrade Cloud Volumes ONTAP

BlueXP notifies you when a new version is available for upgrade. You can start the upgrade process from this notification. For more information, see [Upgrade from BlueXP notifications](#).

Another way to perform software upgrades by using an image on an external URL. This option is helpful if BlueXP can't access the S3 bucket to upgrade the software or if you were provided with a patch. For more information, see [Upgrade from an image available at a URL](#).

### Upgrade from BlueXP notifications

BlueXP displays a notification in Cloud Volumes ONTAP working environments when a new version of Cloud Volumes ONTAP is available:



Before you can upgrade Cloud Volumes ONTAP through the BlueXP notification, you must have a NetApp Support Site account.

You can start the upgrade process from this notification, which automates the process by obtaining the software image from an S3 bucket, installing the image, and then restarting the system.

### Before you begin

BlueXP operations such as volume or aggregate creation must not be in progress on the Cloud Volumes ONTAP system.

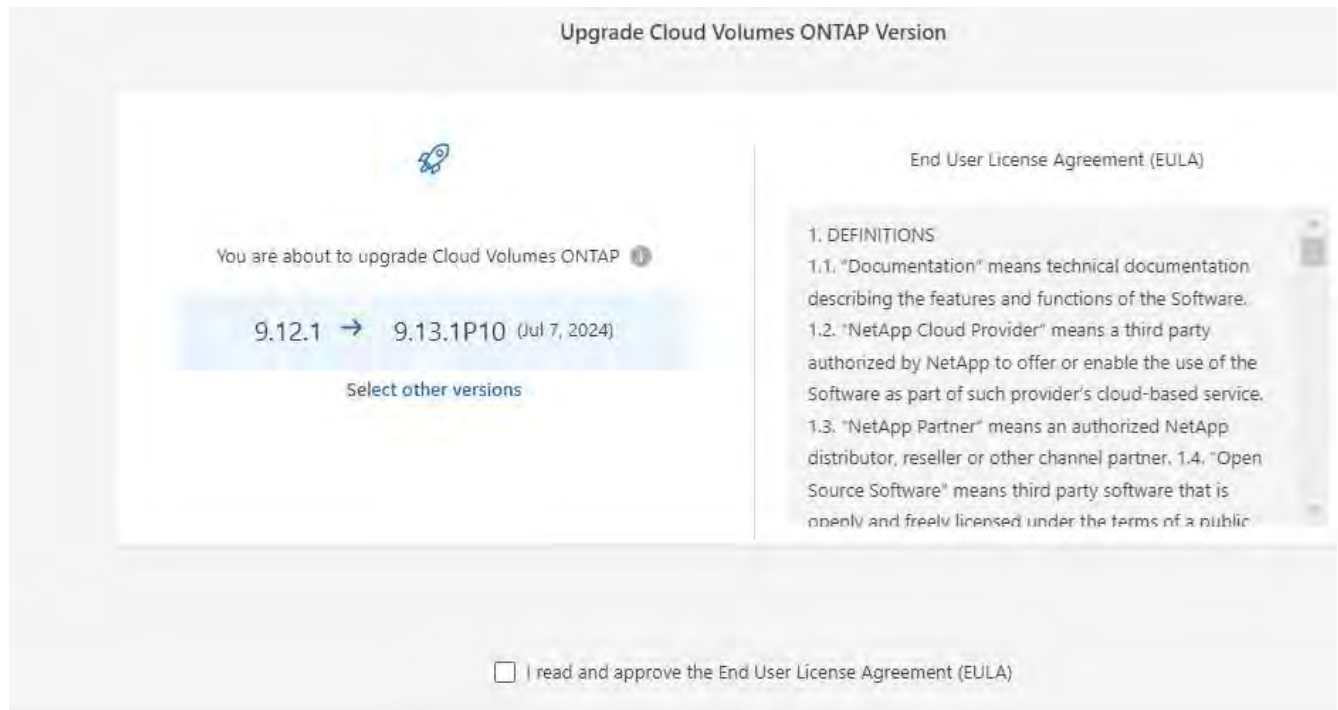
### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. Select a working environment.

A notification appears in the Overview tab if a new version is available:



3. If you want to upgrade the installed version of Cloud Volumes ONTAP, click **Upgrade Now!** By default, you see the latest, compatible version for upgrade.



If you want to upgrade to another version, click **Select other versions**. You see the latest Cloud Volumes ONTAP versions listed that are also compatible with the installed version on your system.

For example, the installed version on your system is 9.12.1P3, and the following compatible versions are available:

- 9.12.1P4 to 9.12.1P14
- 9.13.1 and 9.13.1P1

You see 9.13.1P1 as the default version for upgrade, and 9.12.1P13, 9.13.1P14, 9.13.1, and 9.13.1P1 as the other available versions.

4. Optionally, you can click **All versions** to enter another version that you want to upgrade to (say, the next patch of the installed version). For a compatible upgrade path of your current Cloud Volumes ONTAP version, refer to [Supported upgrade paths](#).

5. Click **Save**, and then **Apply**.

Select the ONTAP version you want to upgrade to:

Version	Date
<input type="radio"/> 9.12.1P14	Aug 22, 2024
<input type="radio"/> 9.12.1P13	Jul 7, 2024
<input type="radio"/> 9.13.1P10	Jul 7, 2024
<input type="radio"/> 9.13.1P9	May 9, 2024

All versions ⌵

Write the version you want to upgrade to:

6. In the Upgrade Cloud Volumes ONTAP page, read the EULA, and then select **I read and approve the EULA**.
7. Click **Upgrade**.
8. To check the status of the upgrade, click the Settings icon and select **Timeline**.

### Result

BlueXP starts the software upgrade. You can perform actions on the working environment when the software update is complete.

### After you finish

If you suspended SnapMirror transfers, use System Manager to resume the transfers.

### Upgrade from an image available at a URL

You can place the Cloud Volumes ONTAP software image on the Connector or on an HTTP server and then initiate the software upgrade from BlueXP. You might use this option if BlueXP can't access the S3 bucket to upgrade the software.

### Before you begin

- BlueXP operations such as volume or aggregate creation must not be in progress on the Cloud Volumes

ONTAP system.

- If you use HTTPS to host ONTAP images, the upgrade can fail due to SSL authentication issues, which are caused by missing certificates. The workaround is to generate and install a CA-signed certificate to be used for authentication between ONTAP and BlueXP.

Go to the NetApp Knowledge Base to view step-by-step instructions:

[NetApp KB: How to configure BlueXP as an HTTPS server to host upgrade images](#)

## Steps

1. Optional: Set up an HTTP server that can host the Cloud Volumes ONTAP software image.

If you have a VPN connection to the virtual network, you can place the Cloud Volumes ONTAP software image on an HTTP server in your own network. Otherwise, you must place the file on an HTTP server in the cloud.

2. If you use your own security group for Cloud Volumes ONTAP, ensure that the outbound rules allow HTTP connections so Cloud Volumes ONTAP can access the software image.



The predefined Cloud Volumes ONTAP security group allows outbound HTTP connections by default.

3. Obtain the software image from [the NetApp Support Site](#).
4. Copy the software image to a directory on the Connector or on an HTTP server from which the file will be served.

Two paths are available. The correct path depends on your Connector version.

- `/opt/application/netapp/cloudmanager/docker_occm/data/ontap/images/`

- `/opt/application/netapp/cloudmanager/ontap/images/`

5. From the working environment in BlueXP, click the ... (**ellipses icon**), and then click **Update Cloud Volumes ONTAP**.
6. On the Update Cloud Volumes ONTAP version page, enter the URL, and then click **Change Image**.

If you copied the software image to the Connector in the path shown above, you would enter the following URL:

`http://<Connector-private-IP-address>/ontap/images/<image-file-name>`



In the URL, **image-file-name** must follow the format "cot.image.9.13.1P2.tgz".

7. Click **Proceed** to confirm.

## Result

BlueXP starts the software update. You can perform actions on the working environment once the software update is complete.

## After you finish

If you suspended SnapMirror transfers, use System Manager to resume the transfers.

## Fix download failures when using a Google Cloud NAT gateway

The Connector automatically downloads software updates for Cloud Volumes ONTAP. The download can fail if your configuration uses a Google Cloud NAT gateway. You can correct this issue by limiting the number of parts that the software image is divided into. This step must be completed by using the BlueXP API.

### Step

1. Submit a PUT request to `/occm/config` with the following JSON as body:

```
{
  "maxDownloadSessions": 32
}
```

The value for `maxDownloadSessions` can be 1 or any integer greater than 1. If the value is 1, then the downloaded image will not be divided.

Note that 32 is an example value. The value that you should use depends on your NAT configuration and the number of sessions that you can have simultaneously.

[Learn more about the `/occm/config` API call.](#)

## Register Cloud Volumes ONTAP pay-as-you-go systems

Support from NetApp is included with Cloud Volumes ONTAP pay-as-you-go (PAYGO) systems, but you must first activate support by registering the systems with NetApp.

Registering a PAYGO system with NetApp is required to upgrade ONTAP software using any of the methods [described on this page](#).











A system that isn't registered for support will still receive the software update notifications that appear in BlueXP when a new version is available. But you will need to register the system before you can upgrade the software.

### Steps

1. If you have not yet added your NetApp Support Site account to BlueXP, go to **Account Settings** and add it now.

[Learn how to add NetApp Support Site accounts.](#)

2. On the Canvas page, double-click the name of the system you want to register..
3. On the Overview tab, click the Features panel and then click the pencil icon next to **Support Registration**.

Information		Features
Working Environment Tags	Tags	
Scheduled Downtime	Off	
S3 Storage Classes	Standard-Infrequent Access	
Instance Type	m5.xlarge	
Write Speed	Normal	
Ransomware Protection	Off	
Support Registration	Not Registered	
CI Fs Setup		

4. Select a NetApp Support Site account and click **Register**.

### Result

BlueXP registers the system with NetApp.

## Convert a Cloud Volumes ONTAP node-based license to a capacity-based license

After the end of availability (EOA) of your node-based licenses, you should transition to capacity-based licensing by using the BlueXP license conversion tool.

For annual or longer-term commitments, NetApp recommends that you contact your NetApp representative prior to the EOA date (11 November, 2024) or license expiration date to ensure that the prerequisites for the transition are in place. If you don't have a long-term contract for a Cloud Volumes ONTAP node and run your system against an on-demand pay-as-you-go (PAYGO) subscription, it is important to plan your conversion before the end of support (EOS) on 31 December, 2024. In both the cases, you should ensure that your system fulfills the requirements before you use the BlueXP license conversion tool for a seamless transition.

For information about the EOA and EOS, refer to [End of availability of node-based licenses](#).

### About this task

- When you use the license conversion tool, the transition from node-based to capacity-based licensing model is carried out in place and online that eliminates the need for any data migration or provisioning of additional cloud resources.
- It is a non-disruptive operation, and no service disruption or application downtime occurs.
- The account and application data in your Cloud Volumes ONTAP system remains intact.
- The underlying cloud resources remain unaffected post conversion.
- The license conversion tool supports all deployment types, such as single node, high availability (HA) in single availability zone (AZ), HA in multiple AZ, bring your own license (BYOL), and PAYGO.

- The tool supports all node-based licenses as the source and all capacity-based licenses as the destination. For example, if you have a PAYGO Standard node-based license, you can convert it to any capacity-based license purchased through the marketplace. NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAP](#).
- The conversion is supported for all cloud providers, AWS, Azure, and Google Cloud.
- Post conversion, the serial number of the node-based license will be replaced by a capacity-based format. This is done as a part of the conversion, and is reflected on your NetApp Support Site (NSS) account.
- When you transition to the capacity-based model, your data continues to be retained in the same location as the node-based licensing. This approach guarantees no disruption in data placement, and upholds data sovereignty principles throughout the transition.

### Before your begin

- You should have an NSS account with customer access or administrator access.
- Your NSS account should be registered with the BlueXP user credentials.
- The working environment should be linked to the NSS account with customer access or administrator access.
- You should have a valid capacity-based license in place, either a BYOL license or marketplace subscription.
- A capacity-based license should be available in the BlueXP account. This license can be a marketplace subscription or a BYOL/private offer package in BlueXP digital wallet.
- Understand the following criteria before selecting a destination package:
  - If the account has a capacity-based BYOL license, the destination package selected should align with the account's BYOL capacity-based licenses:
    - When `Professional` is selected as the destination package, the account should have a BYOL license with a Professional package:
    - When `Essentials` is selected as the destination package, the account should have a BYOL license with the Essentials package.
  - If the destination package does not align with the account's BYOL license availability, it implies that the capacity-based license might not include the selected package. In this case, you will be charged through your marketplace subscription.
  - If there is no capacity-based BYOL license but only a marketplace subscription, you should ensure that the selected package is included in your capacity-based marketplace subscription.
  - If there is not enough capacity in your existing capacity-based license, and if you have a marketplace subscription to charge for the additional capacity usage, you will be charged for the additional capacity through your marketplace subscription.
  - If there is not enough capacity in your existing capacity-based license, and you don't have a marketplace subscription to charge for the additional capacity usage, the conversion cannot occur. You should add a marketplace subscription to charge the additional capacity or extend the available capacity to your current license.
  - If the destination package does not align with the account's BYOL license availability and also if there is not enough capacity in your existing capacity-based license, then you will be charged through your marketplace subscription.



If any of these requirements is not fulfilled, the license conversion does not happen. In specific cases, the license might be converted, but cannot be used. Click the information icon to identify the issues and take corrective actions.

## Steps

1. On the Canvas page, double-click the name of the working environment for which you want to modify the license type.
2. On the Overview tab, click the Features panel.
3. Check the pencil icon next to **Charging method**. If the charging method for your system is `Node Based`, you can convert it to by-capacity charging.



The icon is disabled if your Cloud Volumes ONTAP system is already charged by capacity, or if any of the requirements is not fulfilled.

4. On the **Convert Node-based licenses to Capacity-based** screen, verify the working environment name and source license details.
5. Select the destination package for converting the existing license:
  - Essentials. The default value is `Essentials`.
  - Professional
6. If you have a BYOL license, you can select the checkbox to delete the node-based license from BlueXP digital wallet post conversion. If the conversion is not complete, then even on selecting this checkbox, the license will not be deleted from the digital wallet. If you have a marketplace subscription, this option is unavailable.
7. Select the check box to confirm that you understand the implications of the change, and then click **Proceed**.

## After you finish

View the new license serial number and verify the changes in BlueXP digital wallet.

## Pricing in different hyperscalars

For details on pricing, go to the [NetApp BlueXP website](#).

For information about private offers in specific hyperscalars, write to:

- AWS - [aws@netapp.com](mailto:aws@netapp.com)
- Azure - [azure@netapp.com](mailto:azure@netapp.com)
- Google Cloud - [gcp@netapp.com](mailto:gcp@netapp.com)

## Start and stop a Cloud Volumes ONTAP system

You can stop and start Cloud Volumes ONTAP from BlueXP to manage your cloud compute costs.

## Scheduling automatic shutdowns of Cloud Volumes ONTAP

You might want to shut down Cloud Volumes ONTAP during specific time intervals to lower your compute costs. Rather than do this manually, you can configure BlueXP to automatically shut down and then restart



systems at specific times.

### About this task

- When you schedule an automatic shutdown of your Cloud Volumes ONTAP system, BlueXP postpones the shutdown if an active data transfer is in progress.

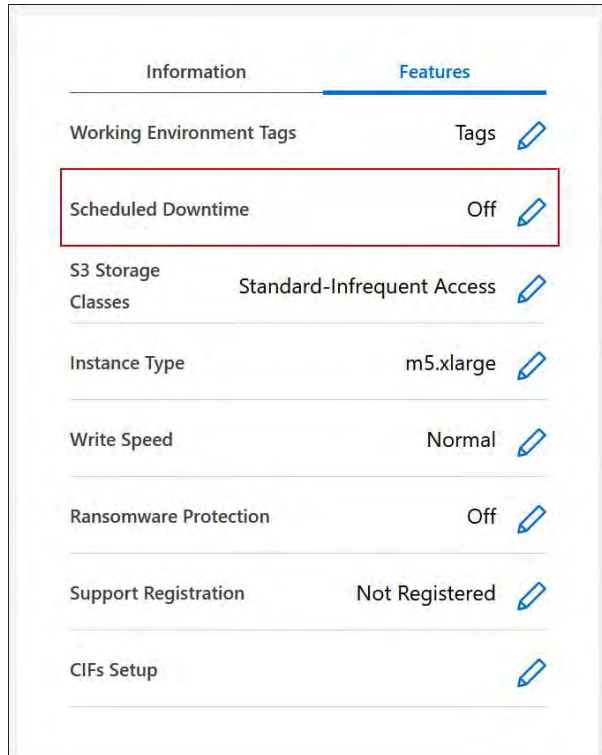
BlueXP shuts down the system after the transfer is complete.

- This task schedules automatic shutdowns of both nodes in an HA pair.
- Snapshots of boot and root disks are not created when turning off Cloud Volumes ONTAP through scheduled shutdowns.

Snapshots are automatically created only when performing a manual shutdown, as described in the next section.

### Steps

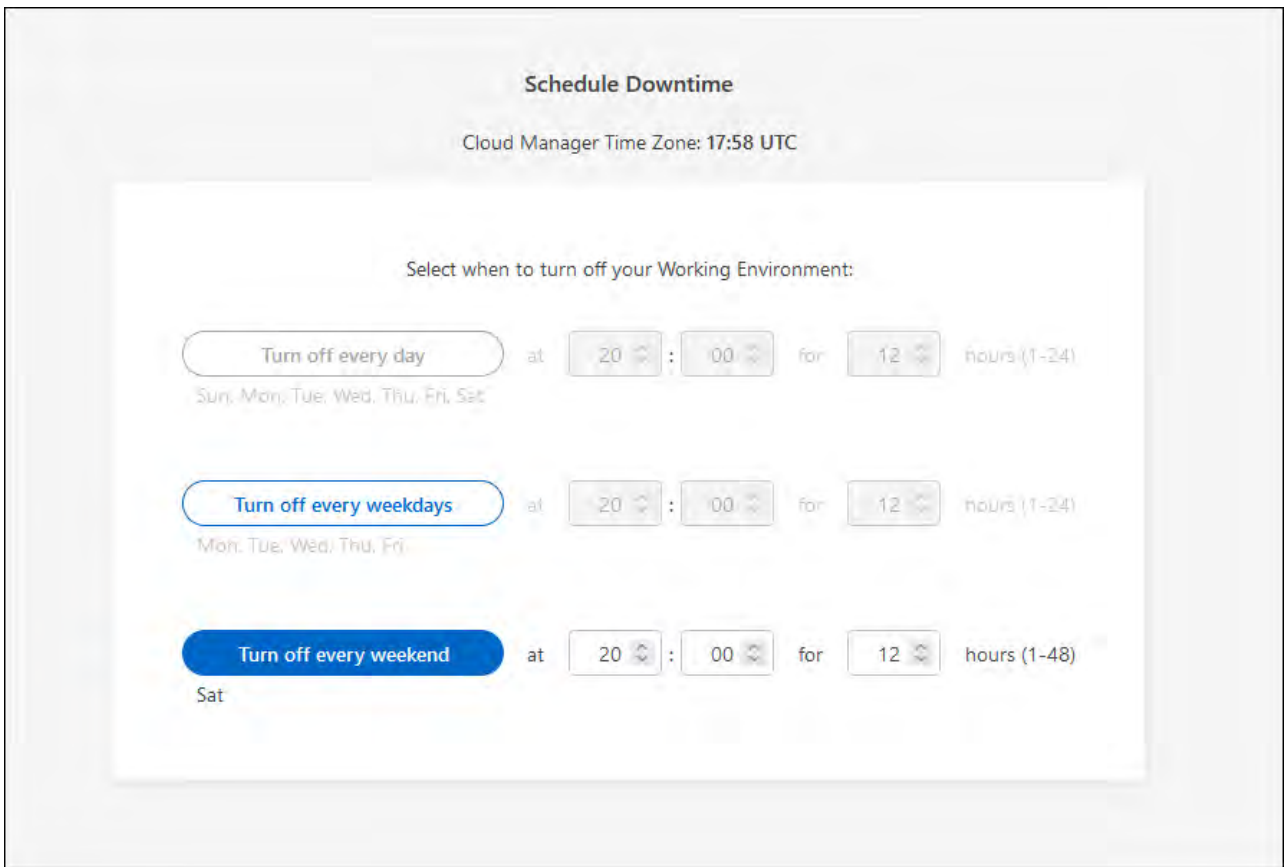
1. On the Canvas page, double-click the desired working environment.
2. On the Overview tab, click the Features panel and then click the pencil icon next to **Scheduled Downtime**.



3. Specify the shutdown schedule:
  - a. Choose whether you want to shut down the system every day, every weekday, every weekend, or any combination of the three options.
  - b. Specify when you want to turn off the system and for how long you want it turned off.

### Example

The following image shows a schedule that instructs BlueXP to shut down the system every Saturday at 20:00 P.M. (8:00 PM) for 12 hours. BlueXP restarts the system every Monday at 12:00 a.m.



4. Click **Save**.

### Result

BlueXP saves the schedule. The corresponding Scheduled Downtime line item under the Features panel displays 'On'.

### Stopping Cloud Volumes ONTAP

Stopping Cloud Volumes ONTAP saves you from accruing compute costs and creates snapshots of the root and boot disks, which can be helpful for troubleshooting.




To reduce costs, BlueXP periodically deletes older snapshots of root and boot disks. Only the two most recent snapshots are retained for both the root and boot disks.

### About this task

When you stop an HA pair, BlueXP shuts down both nodes.

### Steps

1. From the working environment, click the **Turn off** icon.



2. Keep the option to create snapshots enabled because the snapshots can enable system recovery.
3. Click **Turn Off**.

It can take up to a few minutes to stop the system. You can restart systems at a later time from the working environment page.



Snapshots are created automatically upon reboot.

## Synchronize Cloud Volumes ONTAP system time using the NTP server

Specifying an NTP server synchronizes the time between the systems in your network, which can help prevent issues due to time differences.

Specify an NTP server using the [BlueXP API](#) or from the user interface when you [create a CIFS server](#).

## Modify system write speed

BlueXP enables you to choose a normal or high write speed for Cloud Volumes ONTAP. The default write speed is normal. You can change to high write speed if fast write performance is required for your workload.

High write speed is supported with all types of single node systems and some HA pair configurations. View supported configurations in the [Cloud Volumes ONTAP Release Notes](#)









Before you change the write speed, you should [understand the differences between the normal and high settings](#).

### About this task

- Ensure that operations such as volume or aggregate creation are not in progress.
- Be aware that this change restarts the Cloud Volumes ONTAP system. This is disruptive process that requires downtime for the entire system.

### Steps

1. On the Canvas page, double-click the name of the system you configure to the write speed.
2. On the Overview tab, click the Features panel and then click the pencil icon next to **Write Speed**.

Information		Features
Working Environment Tags		Tags 
Scheduled Downtime		Off 
S3 Storage Classes	Standard-Infrequent Access	
Instance Type	m5.xlarge	
Write Speed	Normal	
Ransomware Protection		Off 
Support Registration	Not Registered	
CIFs Setup		

### 3. Select **Normal** or **High**.

If you choose High, then you'll need to read the "I understand..." statement and confirm by checking the box.



The **High** write speed option is supported with Cloud Volumes ONTAP HA pairs in Google Cloud starting with version 9.13.0.

### 4. Click **Save**, review the confirmation message, and then click **Approve**.

## Change the Cloud Volumes ONTAP cluster admin password

Cloud Volumes ONTAP includes a cluster admin account. You can change the password for this account from BlueXP, if needed.



You should not change the password for the admin account through ONTAP System Manager or the ONTAP CLI. The password will not be reflected in BlueXP. As a result, BlueXP cannot monitor the instance properly.

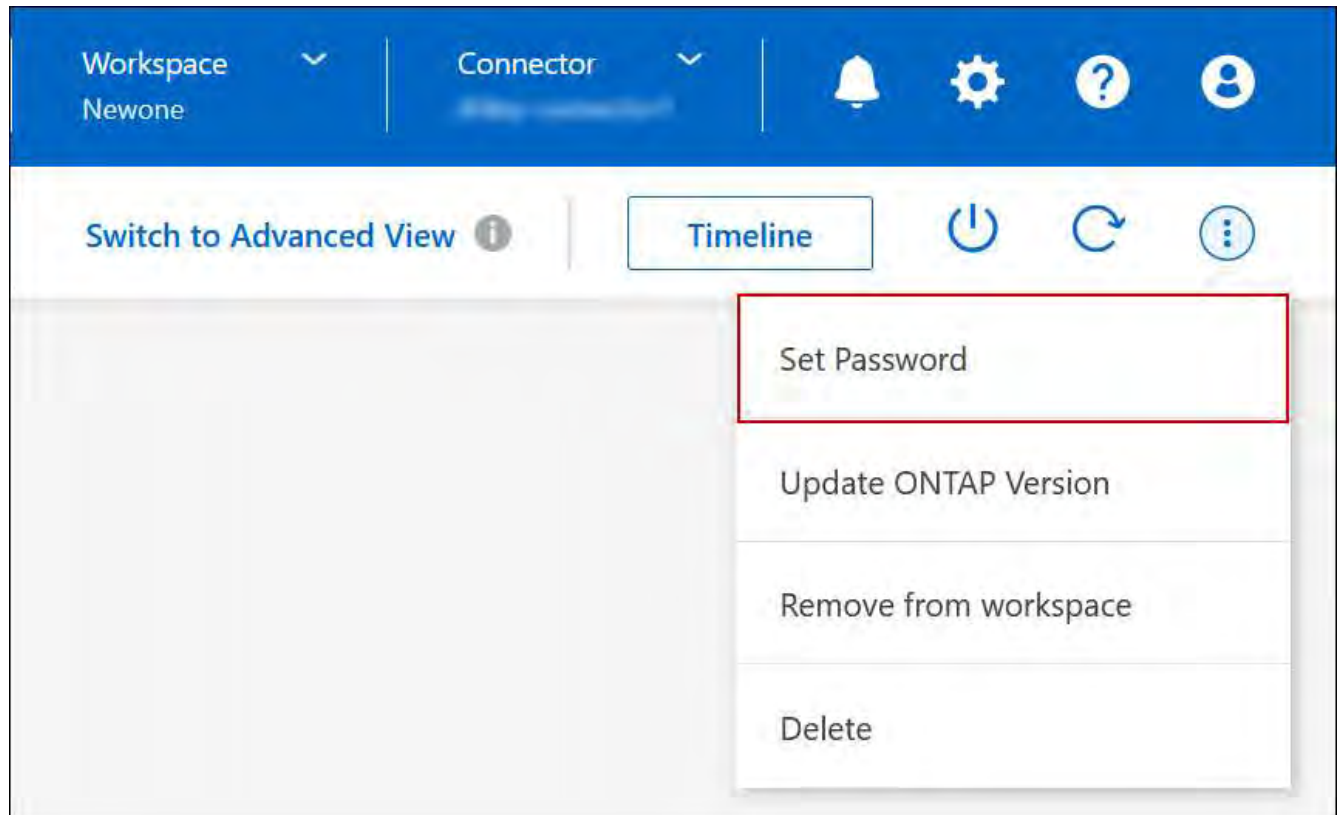
### About this task

The password must observe a few rules. The new password:

- Shouldn't contain the word `admin`
- Must be between eight and 50 characters in length
- Must contain at least one English letter and one digit
- Shouldn't contain these special characters: / ( ) { } [ ] # : % " ? \

### Steps

1. On the Canvas page, double-click the name of the Cloud Volumes ONTAP working environment.
2. On the upper right of the BlueXP console, click the ellipses icon, and select **Set password**.



## Add, remove, or delete systems

### Add an existing Cloud Volumes ONTAP system to BlueXP

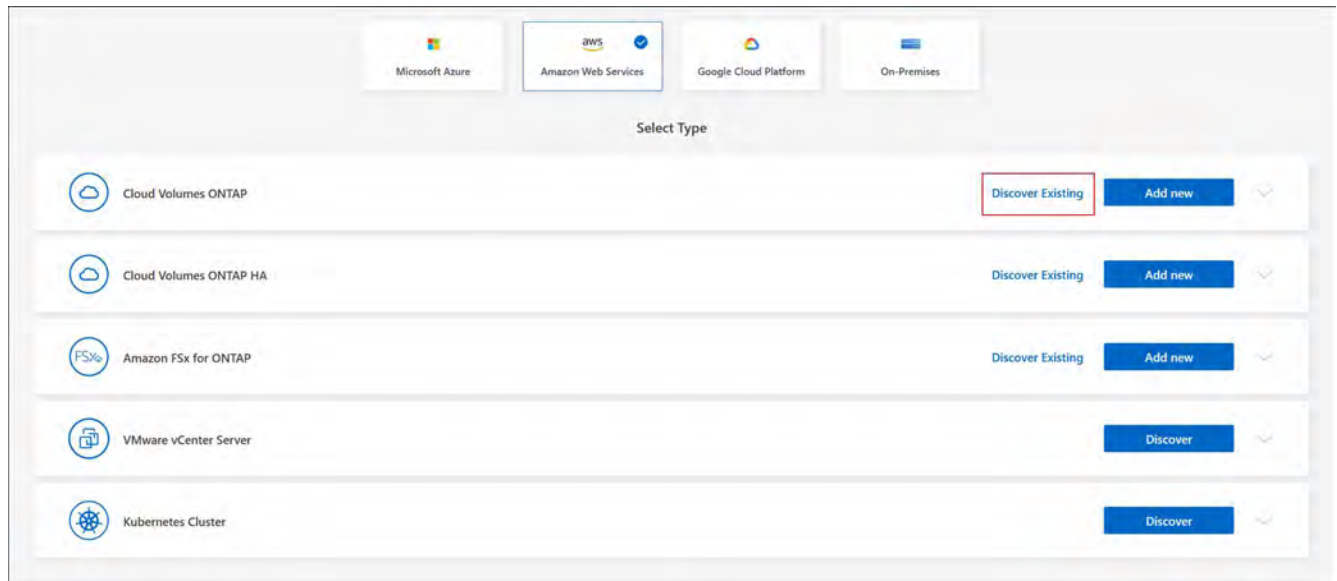
You can discover and add existing Cloud Volumes ONTAP systems to BlueXP. You might do this if you deployed a new BlueXP system.

#### Before you begin

You must know the password for the Cloud Volumes ONTAP admin user account.

#### Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment**.
3. Select the cloud provider in which the system resides.
4. Choose the type of Cloud Volumes ONTAP system to add.
5. Click the link to discover an existing system.



6. On the Region page, select a region. You can see the systems that are running in the selected region.



Cloud Volumes ONTAP systems are represented as instances on this page. From the list, you can select only those instances that are registered with the current account.

7. On the Credentials page, enter the password for the Cloud Volumes ONTAP admin user, and then select **Go**.

### Result

BlueXP adds the Cloud Volumes ONTAP instances to the project or workspace.

### Remove a Cloud Volumes ONTAP working environment from BlueXP

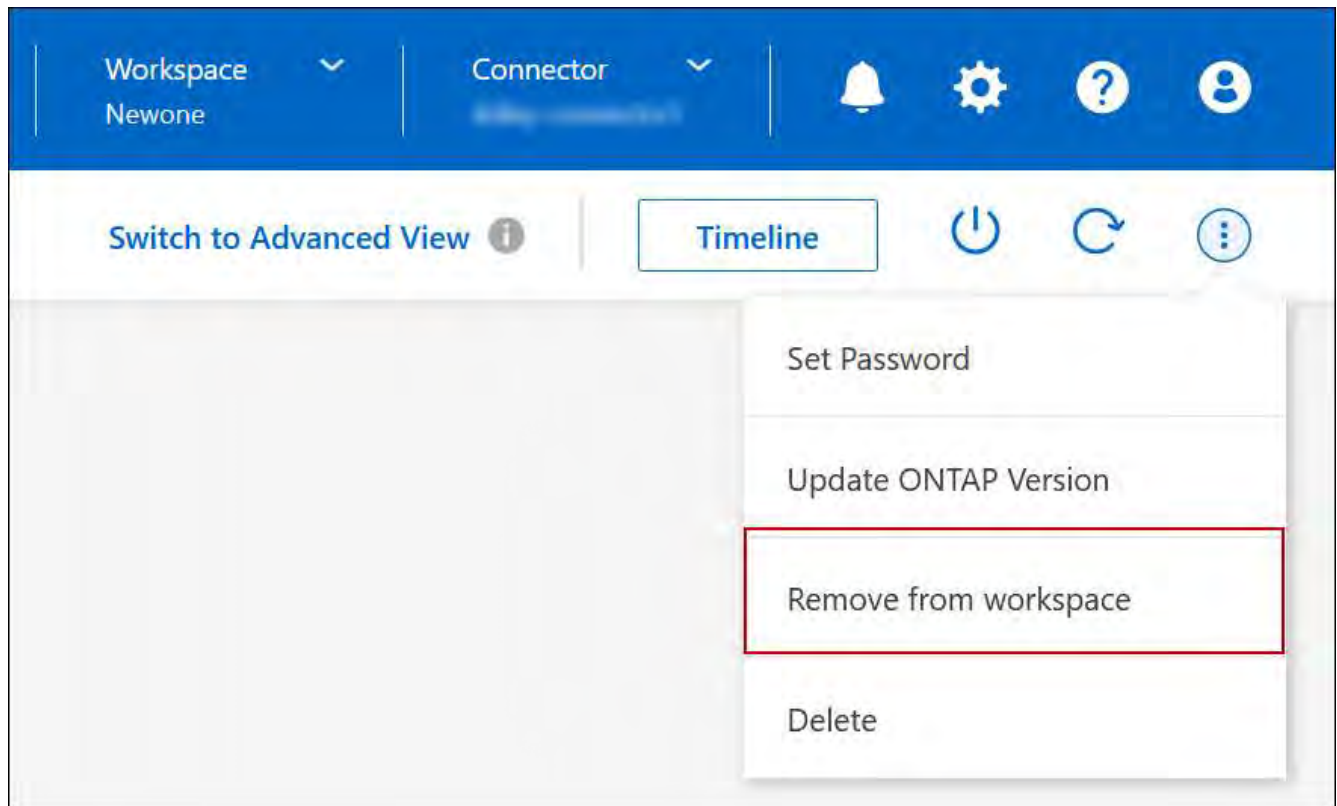
You can remove a Cloud Volumes ONTAP working environment to move it to another system or to troubleshoot discovery issues.

#### About this task

Removing a Cloud Volumes ONTAP working environment removes it from BlueXP. It does not delete the Cloud Volumes ONTAP system. You can later rediscover the working environment—for example, if you had problems during the initial discovery.

#### Steps

1. On the Canvas page, double-click on the working environment you want to remove.
2. On the upper right of the BlueXP console, click the ellipses icon, and select **Remove from workspace**.



3. In the Review from Workspace window, click **Remove**.

### Result

BlueXP removes the working environment. Users can rediscover this working environment from the Canvas page at any time.

### Delete a Cloud Volumes ONTAP system from BlueXP

You should always delete Cloud Volumes ONTAP systems from BlueXP, rather than from your cloud provider's console. For example, if you terminate a licensed Cloud Volumes ONTAP instance from your cloud provider, then you can't use the license key for another instance. You must delete the working environment from BlueXP to release the license.

When you delete a working environment, BlueXP terminates Cloud Volumes ONTAP instances and deletes disks and snapshots.



Other resources, such as backups managed by BlueXP backup and recovery, and instances for BlueXP classification, are not deleted when you delete a working environment. You'll need to manually delete them. If you don't, then you'll continue to incur charges for these resources.

When BlueXP deploys Cloud Volumes ONTAP in your cloud provider, it enables termination protection on the instances. This option helps prevent accidental termination.

### Steps

1. If you enabled BlueXP backup and recovery on the working environment, determine whether the backed up data is still required and then [delete the backups, if necessary](#).

BlueXP backup and recovery is independent from Cloud Volumes ONTAP by design. BlueXP backup and

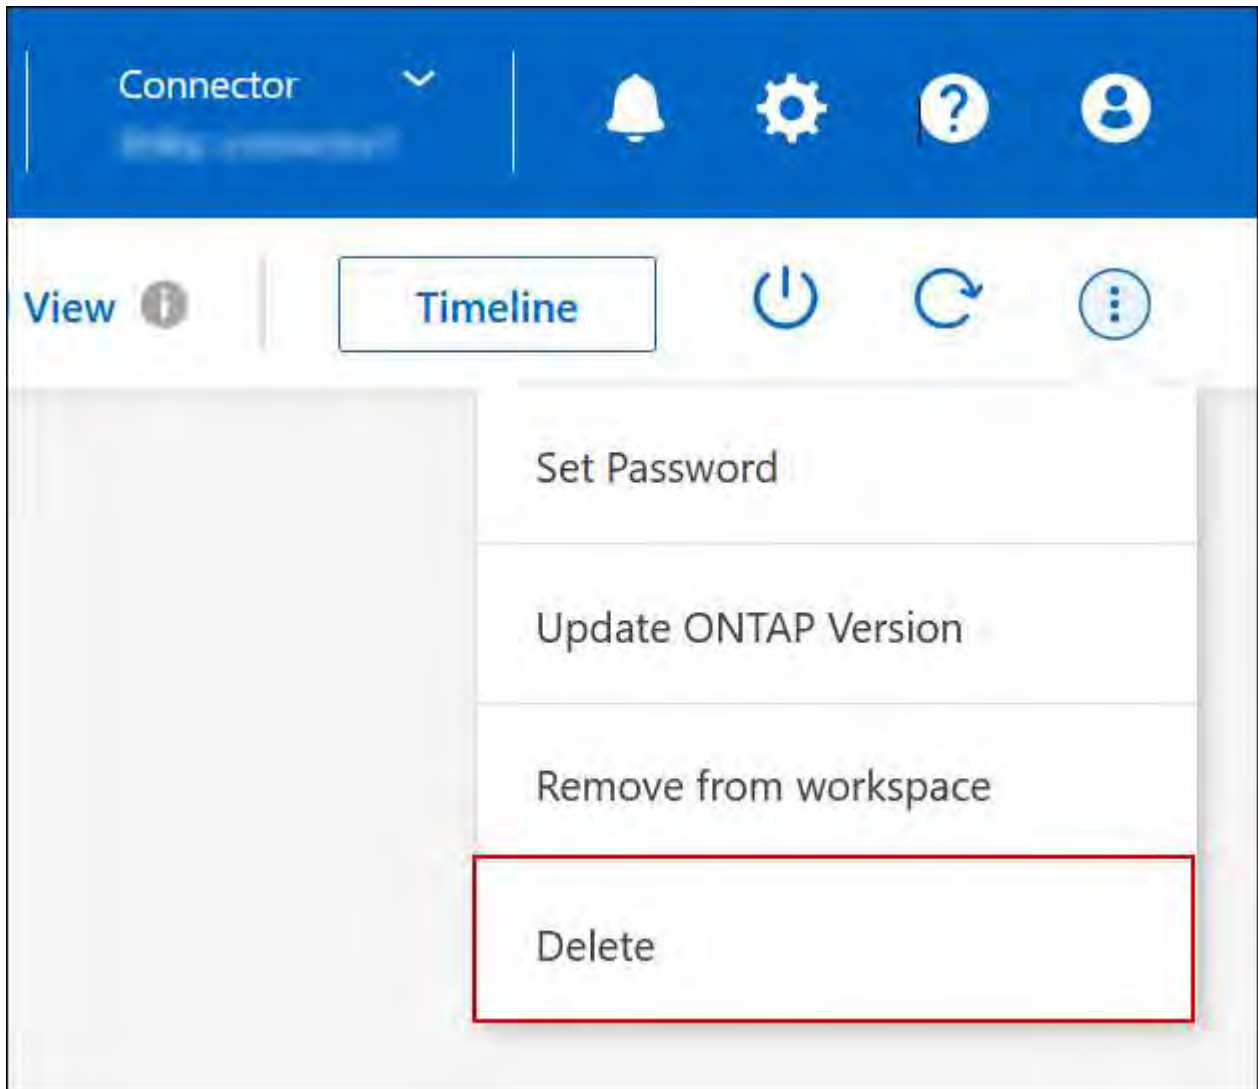


recovery doesn't automatically delete backups when you delete a Cloud Volumes ONTAP system, and there is no current support in the UI to delete the backups after the system has been deleted.

2. If you enabled BlueXP classification on this working environment and no other working environments use this service, then you'll need to delete the instance for the service.

[Learn more about the BlueXP classification instance.](#)

3. Delete the Cloud Volumes ONTAP working environment.
  - a. On the Canvas page, double-click the name of the Cloud Volumes ONTAP working environment that you want to delete.
  - b. On the upper right of the BlueXP console, click the ellipses icon, and select **Delete**.



- c. Under the Delete Working Environment window, type the name of the working environment and then click **Delete**. It can take up to five minutes to delete the working environment.



BlueXP backup and recovery is free only for Cloud Volumes ONTAP Professional licenses. This free benefit does not apply to deleted environments. If backed up copies of the Cloud Volumes ONTAP environment are retained in a BlueXP backup and recovery instance, you will be charged for the backed up copies until they are deleted.



## AWS administration

### Modify the EC2 instance type for a Cloud Volumes ONTAP system in AWS

You can choose from several instance or types when you launch Cloud Volumes ONTAP in AWS. You can change the instance type at any time if you determine that it is undersized or oversized for your needs.

#### About this task

- Automatic giveback must be enabled on a Cloud Volumes ONTAP HA pair (this is the default setting). If it isn't, then the operation will fail.

[ONTAP 9 Documentation: Commands for configuring automatic giveback](#)

- Changing the instance type can affect AWS service charges.
- The operation restarts Cloud Volumes ONTAP.

For single node systems, I/O is interrupted.

For HA pairs, the change is nondisruptive. HA pairs continue to serve data.



BlueXP changes one node at a time by initiating takeover and waiting for give back. NetApp's Quality Assurance team tested both writing and reading files during this process and didn't see any issues on the client side. As connections changed, some retries were observed on the I/O level, but the application layer overcame the rewiring of NFS/CIFS connections.

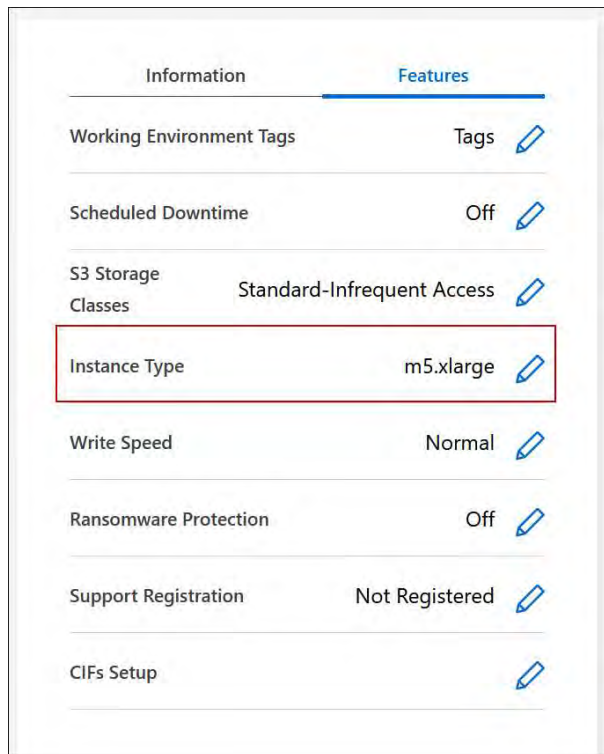
#### Reference

For a list of supported instance types in AWS, refer to [Supported EC2 instances](#).

If you can't change the instance type from c4, m4, or r4 instances, refer to KB article "[Converting an AWS Xen CVO instance to Nitro \(KVM\)](#)".

#### Steps

1. On the Canvas page, select the working environment.
2. On the Overview tab, click the Features panel and then click the pencil icon next to **Instance type**.



If you are using a node-based pay as you go (PAYGO) license, you can optionally choose a different license and instance type by clicking the pencil icon next to **License type**.

1. Choose an instance type, select the check box to confirm that you understand the implications of the change, and then click **Change**.

### Result

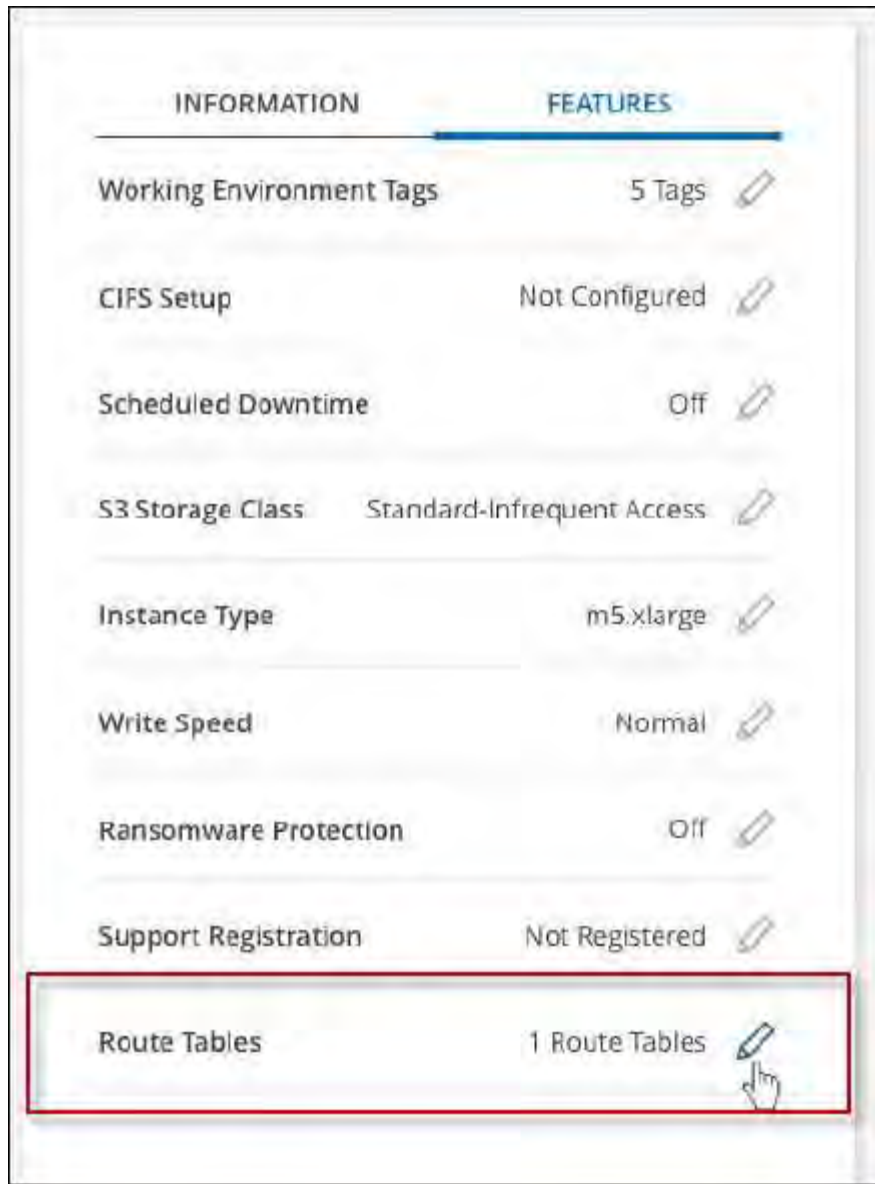
Cloud Volumes ONTAP reboots with the new configuration.

### Modify route tables for Cloud Volumes ONTAP HA pairs in multiple AWS AZs

You can modify the AWS route tables that include routes to the floating IP addresses for an HA pair that's deployed in multiple AWS Availability Zones (AZs). You might do this if new NFS or CIFS clients need to access an HA pair in AWS.

### Steps

1. On the Canvas page, select the working environment.
2. On the Overview tab, click the Features panel and then click the pencil icon next to **Route tables**.



3. Modify the list of selected route tables and then click **Save**.

### Result

BlueXP sends an AWS request to modify the route tables.

## Azure administration

### Change the Azure VM type for Cloud Volumes ONTAP

You can choose from several VM types when you launch Cloud Volumes ONTAP in Microsoft Azure. You can change the VM type at any time if you determine that it is undersized or oversized for your needs.

#### About this task

- Automatic giveback must be enabled on a Cloud Volumes ONTAP HA pair (this is the default setting). If it isn't, then the operation will fail.

[ONTAP 9 Documentation: Commands for configuring automatic giveback](#)

- Changing the VM type can affect Microsoft Azure service charges.
- The operation restarts Cloud Volumes ONTAP.

For single node systems, I/O is interrupted.

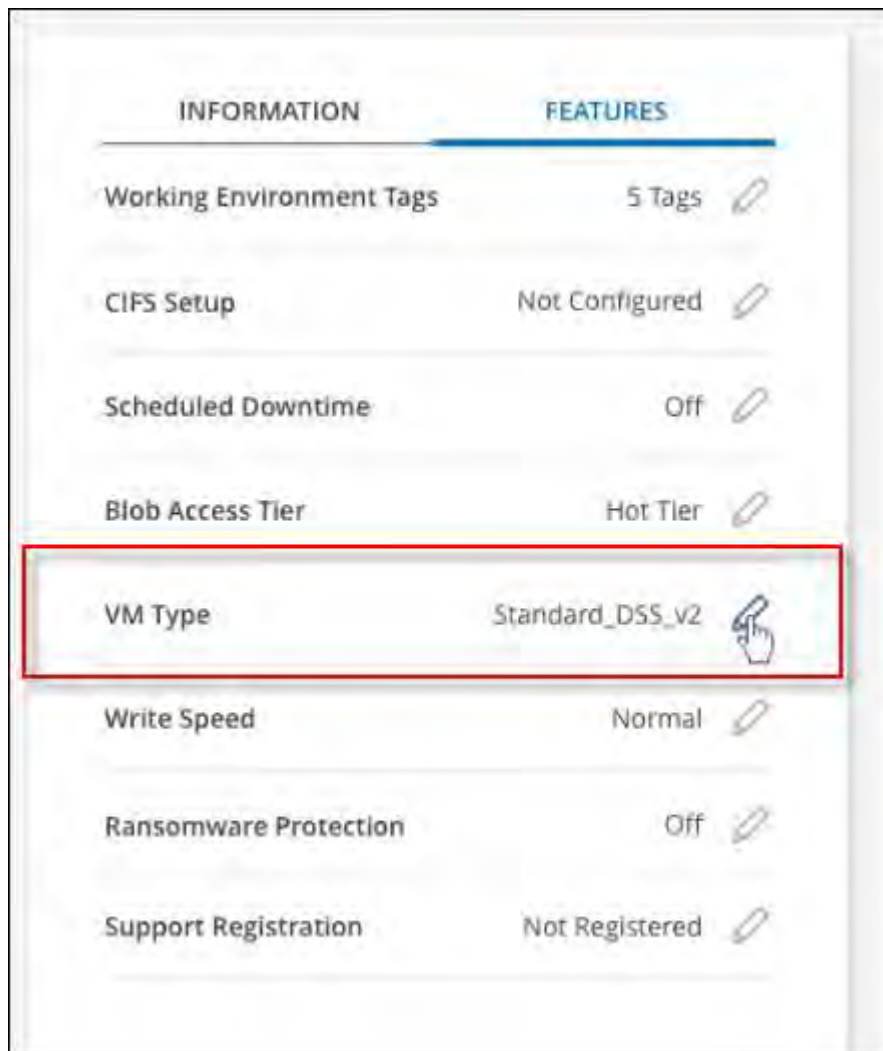
For HA pairs, the change is nondisruptive. HA pairs continue to serve data.



BlueXP changes one node at a time by initiating takeover and waiting for give back. NetApp's Quality Assurance team tested both writing and reading files during this process and didn't see any issues on the client side. As connections changed, some retries were observed on the I/O level, but the application layer overcame the rewiring of NFS/CIFS connections.

### Steps

1. On the Canvas page, select the working environment.
2. On the Overview tab, click the Features panel and then click the pencil icon next to **VM type**.



If you are using a node-based pay-as-you-go (PAYGO) license, you can optionally choose a different license and VM type by clicking the pencil icon next to **License type**.

1. Select a VM type, select the check box to confirm that you understand the implications of the change, and then click **Change**.

## Result

Cloud Volumes ONTAP reboots with the new configuration.

## Override CIFS locks for Cloud Volumes ONTAP HA pairs in Azure

The BlueXP Organization or Account admin can enable a setting in BlueXP that prevents issues with Cloud Volumes ONTAP storage giveback during Azure maintenance events. When you enable this setting, Cloud Volumes ONTAP vetoes CIFS locks and resets active CIFS sessions.

### About this task

Microsoft Azure schedules periodic maintenance events on its virtual machines. When a maintenance event occurs on a Cloud Volumes ONTAP HA pair, the HA pair initiates storage takeover. If there are active CIFS sessions during this maintenance event, the locks on CIFS files can prevent storage giveback.

If you enable this setting, Cloud Volumes ONTAP will veto the locks and reset the active CIFS sessions. As a result, the HA pair can complete storage giveback during these maintenance events.



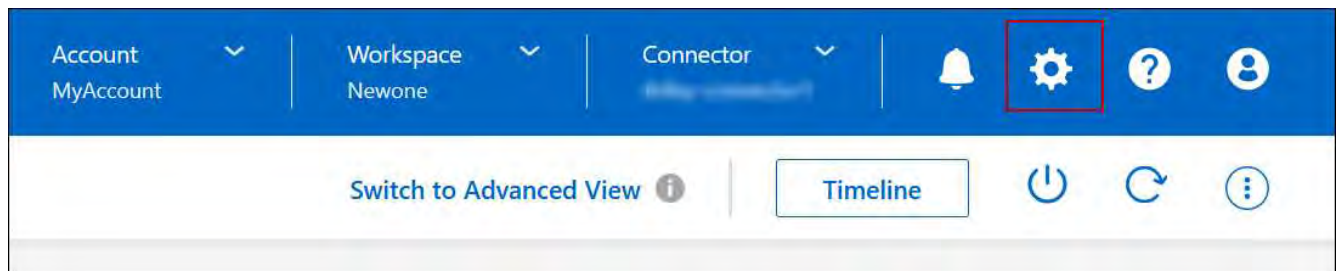
This process might be disruptive to CIFS clients. Data that is not committed from CIFS clients could be lost.

### Before you begin

You need to create a Connector before you can change BlueXP settings. [Learn how](#).

### Steps

1. In the upper right of the BlueXP console, click the Settings icon, and select **Cloud Volumes ONTAP Settings**.



2. Under **Azure**, click **Azure CIFS locks for Azure HA working environments**.
3. Click the checkbox to enable the feature and then click **Save**.

## Use an Azure Private Link or service endpoints for Cloud Volumes ONTAP systems

Cloud Volumes ONTAP uses an Azure Private Link for connections to its associated storage accounts. If needed, you can disable Azure Private Links and use service endpoints instead.

## Overview

By default, BlueXP enables an Azure Private Link for connections between Cloud Volumes ONTAP and its associated storage accounts. An Azure Private Link secures connections between endpoints in Azure and provides performance benefits.

If required, you can configure Cloud Volumes ONTAP to use service endpoints instead of an Azure Private Link.

With either configuration, BlueXP always limits network access for connections between Cloud Volumes ONTAP and storage accounts. Network access is limited to the VNet where Cloud Volumes ONTAP is deployed and the VNet where the Connector is deployed.

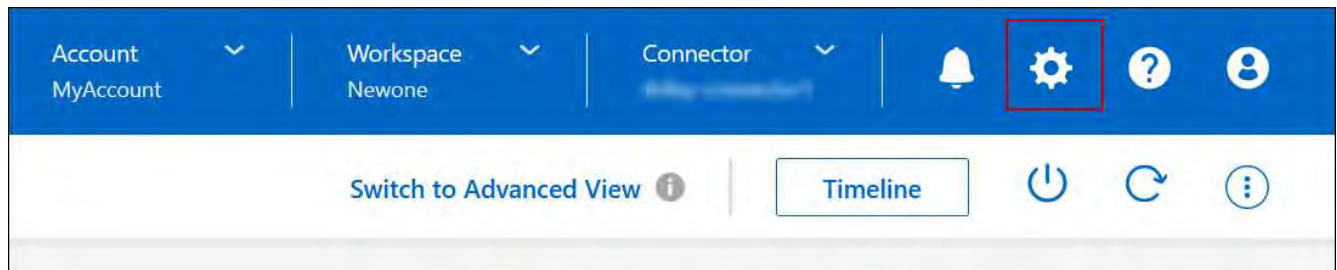
### Disable Azure Private Links and use service endpoints instead

If required by your business, you can change a setting in BlueXP so that it configures Cloud Volumes ONTAP to use service endpoints instead of an Azure Private Link. Changing this setting applies to new Cloud Volumes ONTAP systems that you create. Service endpoints are only supported in [Azure region pairs](#) between the Connector and Cloud Volumes ONTAP VNets.

The Connector should be deployed in the same Azure region as the Cloud Volumes ONTAP systems that it manages, or in the [Azure region pair](#) for the Cloud Volumes ONTAP systems.

## Steps

1. In the upper right of the BlueXP console, click the Settings icon, and select **Cloud Volumes ONTAP Settings**.



2. Under **Azure**, click **Use Azure Private Link**.
3. Deselect **Private Link connection between Cloud Volumes ONTAP and storage accounts**.
4. Click **Save**.

### After you finish

If you disabled Azure Private Links and the Connector uses a proxy server, you must enable direct API traffic.

[Learn how to enable direct API traffic on the Connector](#)

### Work with Azure Private Links

In most cases, there's nothing that you need to do to set up Azure Private links with Cloud Volumes ONTAP. BlueXP manages Azure Private Links for you. But if you use an existing Azure Private DNS zone, then you'll need to edit a configuration file.

### Requirement for custom DNS

Optionally, if you work with custom DNS, you need to create a conditional forwarder to the Azure private DNS

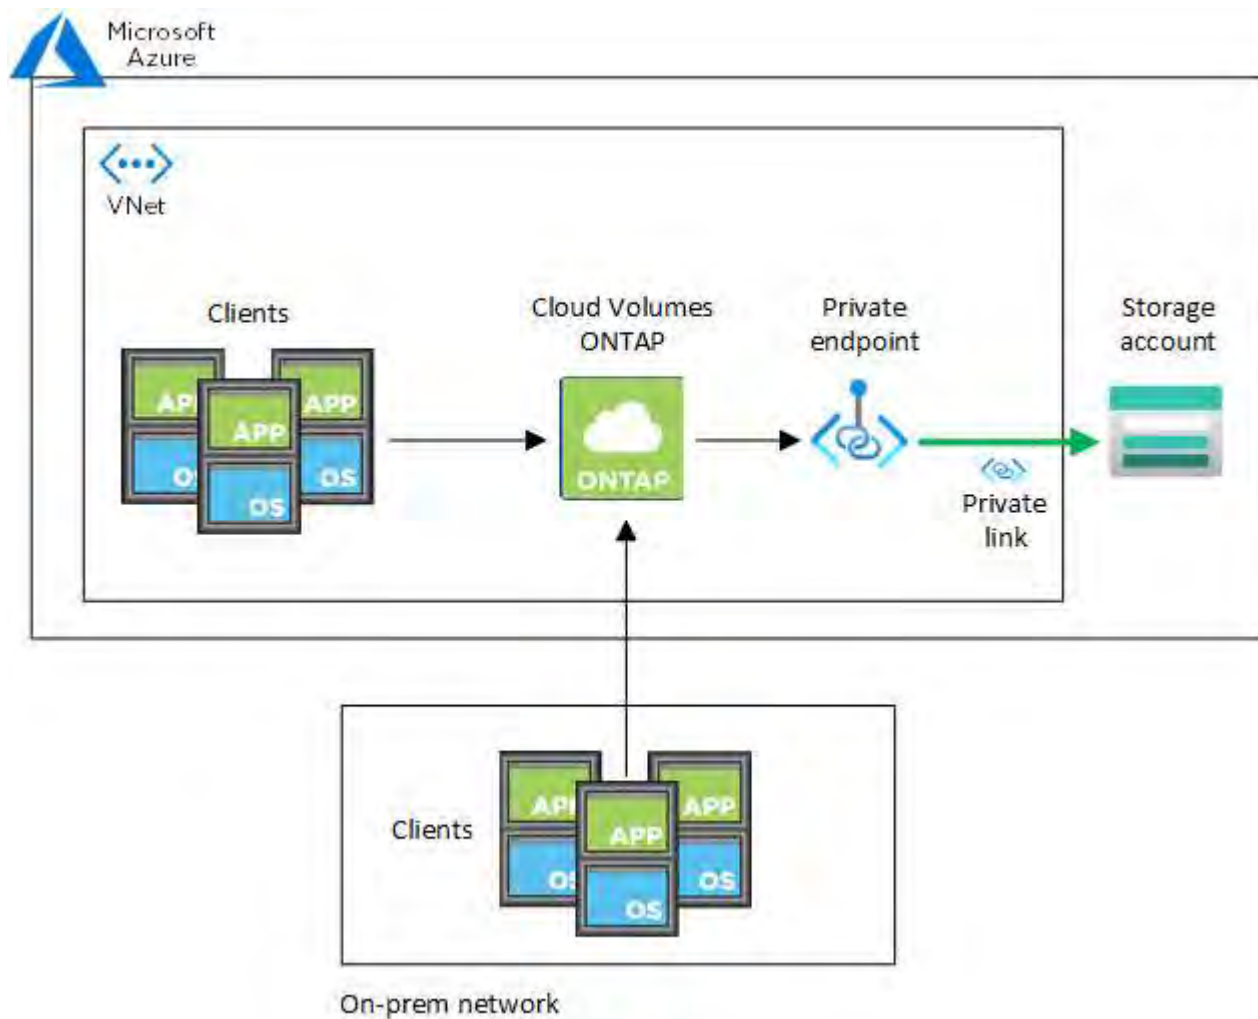
zone from your custom DNS servers. To learn more, refer to [Azure's documentation on using a DNS forwarder](#).

## How Private Link connections work

When BlueXP deploys Cloud Volumes ONTAP in Azure, it creates a private endpoint in the resource group. The private endpoint is associated with storage accounts for Cloud Volumes ONTAP. As a result, access to Cloud Volumes ONTAP storage travels through the Microsoft backbone network.

Client access goes through the private link when clients are within the same VNet as Cloud Volumes ONTAP, within peered VNets, or in your on-premises network when using a private VPN or ExpressRoute connection to the VNet.

Here's an example that shows client access over a private link from within the same VNet and from an on-premises network that has either a private VPN or ExpressRoute connection.



If the Connector and Cloud Volumes ONTAP systems are deployed in different VNets, then you must set up VNet peering between the VNet where the Connector is deployed and the VNet where the Cloud Volumes ONTAP systems are deployed.

## Provide BlueXP with details about your Azure Private DNS

If you use [Azure Private DNS](#), then you need to modify a configuration file on each Connector. Otherwise, BlueXP can't enable the Azure Private Link connection between Cloud Volumes ONTAP and its associated



storage accounts.

Note that the DNS name must match Azure DNS naming requirements [as shown in Azure documentation](#).

### Steps

1. SSH to the Connector host and log in.
2. Navigate to the `/opt/application/netapp/cloudmanager/docker_occm/data` directory.
3. Edit `app.conf` by adding the `user-private-dns-zone-settings` parameter with the following keyword-value pairs:

```
"user-private-dns-zone-settings" : {  
  "resource-group" : "<resource group name of the DNS zone>",  
  "subscription" : "<subscription ID>",  
  "use-existing" : true,  
  "create-private-dns-zone-link" : true  
}
```

The `subscription` keyword is required only if the private DNS zone is in a different subscription than that of the Connector.

4. Save the file and log off the Connector.

A reboot isn't required.

### Enable rollback on failures

If BlueXP fails to create an Azure Private Link as part of specific actions, it completes the action without the Azure Private Link connection. This can happen when creating a new working environment (single node or HA pair), or when the following actions occur on an HA pair: creating a new aggregate, adding disks to an existing aggregate, or creating a new storage account when going above 32 TiB.

You can change this default behavior by enabling rollback if BlueXP fails to create the Azure Private Link. This can help to ensure that you're fully compliant with your company's security regulations.

If you enable rollback, BlueXP stops the action and rolls back all resources that were created as part of the action.

You can enable rollback through the API or by updating the `app.conf` file.

### Enable rollback through the API

#### Step

1. Use the `PUT /occm/config` API call with the following request body:

```
{ "rollbackOnAzurePrivateLinkFailure": true }
```

### Enable rollback by updating `app.conf`



## Steps

1. SSH to the Connector host and log in.
2. Navigate to the following directory: `/opt/application/netapp/cloudmanager/docker_occm/data`
3. Edit `app.conf` by adding the following parameter and value:

```
"rollback-on-private-link-failure": true
```

4. Save the file and log off the Connector.

A reboot isn't required.

## Move an Azure resource group for Cloud Volumes ONTAP in Azure console

Cloud Volumes ONTAP supports Azure resource groups moves but the workflow happens in the Azure console only.

You can move a working environment from one resource group to a different resource group in Azure within the same Azure subscription. Moving resource groups between different Azure subscriptions is not supported.

## Steps

1. Remove the working environment from **Canvas**.

To learn how to remove a working environment, refer to [Removing Cloud Volumes ONTAP working environments](#).

2. Execute the resource group move in the Azure console.

To complete the move, refer to [Move resources to a new resource group or subscription in Microsoft Azure's documentation](#).

3. In **Canvas**, discover the working environment.
4. Look for the new resource group in the information for the working environment.

## Result

The working environment and its resources (VMs, disks, storage accounts, network interfaces, snapshots) are in the new resource group.

## Segregate SnapMirror traffic in Azure

With Cloud Volumes ONTAP in Azure, you can segregate SnapMirror replication traffic from data and management traffic. To segregate SnapMirror replication traffic from your data traffic, you'll add a new network interface card (NIC), an associated intercluster LIF and a non-routable subnet.

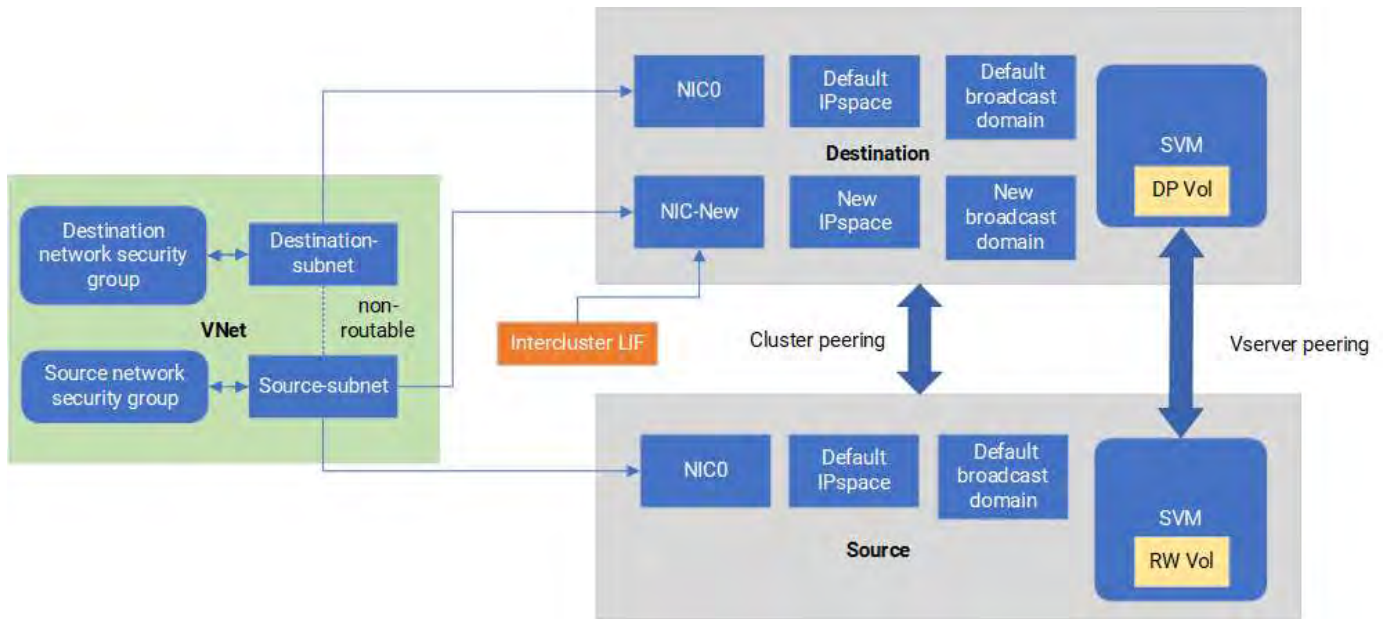
## About SnapMirror traffic segregation in Azure

By default, BlueXP configures all NICs and LIFs in a Cloud Volumes ONTAP deployment on the same subnet. In such configurations, SnapMirror replication traffic and data and management traffic use the same subnet. Segregating SnapMirror traffic leverages an additional subnet that isn't routable to the existing subnet used for

data and management traffic.

### Figure 1

The following diagrams show the segregation of SnapMirror replication traffic with an additional NIC, an associated intercluster LIF and a non-routable subnet in a single node deployment. An HA pair deployment differs slightly.



### Before you begin

Review the following considerations:

- You can only add a single NIC to a Cloud Volumes ONTAP single node or HA-pair deployment (VM instance) for SnapMirror traffic segregation.
- To add a new NIC, the VM instance type you deploy must have an unused NIC.
- The source and destination clusters should have access to the same Virtual Network (VNet). The destination cluster is a Cloud Volumes ONTAP system in Azure. The source cluster can be a Cloud Volumes ONTAP system in Azure or an ONTAP system.

### Step 1: Create an additional NIC and attach to the destination VM

This section provides instructions for how to create an additional NIC and attach it to the destination VM. The destination VM is the single node or HA-pair system in Cloud Volumes ONTAP in Azure where you want to set up your additional NIC.

### Steps

1. In the ONTAP CLI, stop the node.

```
dest::> halt -node <dest_node-vm>
```

2. In the Azure portal, check that the VM (node) status is stopped.

```
az vm get-instance-view --resource-group <dest-rg> --name <dest-vm>
--query instanceView.statuses[1].displayStatus
```

3. Use the Bash environment in Azure Cloud Shell to stop the node.

a. Stop the node.

```
az vm stop --resource-group <dest_node-rg> --name <dest_node-vm>
```

b. Deallocate the node.

```
az vm deallocate --resource-group <dest_node-rg> --name <dest_node-vm>
```

4. Configure network security group rules to make the two subnets (source cluster subnet and destination cluster subnet) non-routable to each other.

a. Create the new NIC on the destination VM.

b. Look up the subnet ID for the source cluster subnet.

```
az network vnet subnet show -g <src_vnet-rg> -n <src_subnet> --vnet
-name <vnet> --query id
```

c. Create the new NIC on the destination VM with the subnet ID for the source cluster subnet. Here you enter the name for the new NIC.

```
az network nic create -g <dest_node-rg> -n <dest_node-vm-nic-new>
--subnet <id_from_prev_command> --accelerated-networking true
```

d. Save the privateIPaddress. This IP address, <new\_added\_nic\_primary\_addr>, is used to create an intercluster LIF in [broadcast domain](#), [intercluster LIF for the new NIC](#).

5. Attach the new NIC to the VM.

```
az vm nic add -g <dest_node-rg> --vm-name <dest_node-vm> --nics
<dest_node-vm-nic-new>
```

6. Start the VM (node).

```
az vm start --resource-group <dest_node-rg> --name <dest_node-vm>
```

7. In the Azure portal, go to **Networking** and confirm that the new NIC, e.g. nic-new, exists and accelerated

networking is enabled.

```
az network nic list --resource-group azure-59806175-60147103-azure-rg
--query "[].{NIC: name, VM: virtualMachine.id}"
```

For HA-pair deployments, repeat the steps for the partner node.

## Step 2: Create a new IPspace, broadcast domain, and intercluster LIF for the new NIC

A separate IPspace for intercluster LIFs provides logical separation between networking functionality for replication between clusters.

Use the ONTAP CLI for the following steps.

### Steps

1. Create the new IPspace (`new_ipspace`).

```
dest::> network ipspace create -ipspace <new_ipspace>
```

2. Create a broadcast domain on the new IPspace (`new_ipspace`) and add the `nic-new` port.

```
dest::> network port show
```

3. For single node systems, the newly added port is `e0b`. For HA-pair deployments with managed disks, the newly added port is `e0d`. For HA-pair deployments with page blobs, the newly added port is `e0e`. Use the node name not the VM name. Find the node name by running `node show`.

```
dest::> broadcast-domain create -broadcast-domain <new_bd> -mtu 1500
-ipspace <new_ipspace> -ports <dest_node-cot-vm:e0b>
```

4. Create an intercluster LIF on the new broadcast-domain (`new_bd`) and on the new NIC (`nic-new`).

```
dest::> net int create -vserver <new_ipspace> -lif <new_dest_node-ic-
lif> -service-policy default-intercluster -address
<new_added_nic_primary_addr> -home-port <e0b> -home-node <node> -netmask
<new_netmask_ip> -broadcast-domain <new_bd>
```

5. Verify creation of the new intercluster LIF.

```
dest::> net int show
```

For HA-pair deployments, repeat the steps for the partner node.

### Step 3: Verify cluster peering between the source and destination systems

This section provides instructions for how to verify peering between the source and destination systems.

Use the ONTAP CLI for the following steps.

#### Steps

1. Verify that the intercluster LIF of the destination cluster can ping the intercluster LIF of the source cluster. Because the destination cluster executes this command, the destination IP address is the intercluster LIF IP address on the source.

```
dest::> ping -lif <new_dest_node-ic-lif> -vserver <new_ipspace>
-destination <10.161.189.6>
```

2. Verify that the intercluster LIF of the source cluster can ping the intercluster LIF of the destination cluster. The destination is the IP address of the new NIC created on the destination.

```
src::> ping -lif <src_node-ic-lif> -vserver <src_svm> -destination
<10.161.189.18>
```

For HA-pair deployments, repeat the steps for the partner node.

### Step 4: Create SVM peering between the source and destination system

This section provides instructions for how to create SVM peering between the source and destination system.

Use the ONTAP CLI for the following steps.

#### Steps

1. Create cluster peering on the destination using the source intercluster LIF IP address as the `-peer-addr`s. For HA pairs, list the source intercluster LIF IP address for both nodes as the `-peer-addr`s.

```
dest::> cluster peer create -peer-addr <10.161.189.6> -ipspace
<new_ipspace>
```

2. Enter and confirm the passphrase.
3. Create cluster peering on the source using the destination cluster LIF IP address as the `peer-addr`s. For HA pairs, list the destination intercluster LIF IP address for both nodes as the `-peer-addr`s.

```
src::> cluster peer create -peer-addr <10.161.189.18>
```

4. Enter and confirm the passphrase.
5. Check that the cluster peered.

```
src::> cluster peer show
```

Successful peering shows **Available** in the availability field.

6. Create SVM peering on the destination. Both source and destination SVMs should be data SVMs.

```
dest::> vserver peer create -vserver <dest_svm> -peer-vserver <src_svm>  
-peer-cluster <src_cluster> -applications snapmirror``
```

7. Accept SVM peering.

```
src::> vserver peer accept -vserver <src_svm> -peer-vserver <dest_svm>
```

8. Check that the SVM peered.

```
dest::> vserver peer show
```

Peer state shows **peered** and peering applications shows **snapmirror**.

#### Step 5: Create a SnapMirror replication relationship between the source and destination system

This section provides instructions for how to create a SnapMirror replication relationship between the source and destination system.

To move an existing SnapMirror replication relationship, you must first break the existing SnapMirror replication relationship before you create a new SnapMirror replication relationship.

Use the ONTAP CLI for the following steps.

#### Steps

1. Create a data protected volume on the destination SVM.

```
dest::> vol create -volume <new_dest_vol> -vserver <dest_svm> -type DP  
-size <10GB> -aggregate <aggr1>
```

2. Create the SnapMirror replication relationship on the destination which includes the SnapMirror policy and schedule for the replication.

```
dest::> snapmirror create -source-path src_svm:src_vol -destination  
-path dest_svm:new_dest_vol -vserver dest_svm -policy  
MirrorAllSnapshots -schedule 5min
```

3. Initialize the SnapMirror replication relationship on the destination.

```
dest::> snapmirror initialize -destination-path <dest_svm:new_dest_vol>
```

4. In the ONTAP CLI, validate the SnapMirror relationship status by running the following command:

```
dest::> snapmirror show
```

The relationship status is `Snapmirrored` and the health of the relationship is `true`.

5. Optional: In the ONTAP CLI, run the following command to view the actions history for the SnapMirror relationship.

```
dest::> snapmirror show-history
```

Optionally, you can mount the source and destination volumes, write a file to the source, and verify the volume is replicating to the destination.

## Google Cloud administration

### Change the Google Cloud machine type for Cloud Volumes ONTAP

You can choose from several machine types when you launch Cloud Volumes ONTAP in Google Cloud. You can change the instance or machine type at any time if you determine that it is undersized or oversized for your needs.

#### About this task

- Automatic giveback must be enabled on a Cloud Volumes ONTAP HA pair (this is the default setting). If it isn't, then the operation will fail.

[ONTAP 9 Documentation: Commands for configuring automatic giveback](#)

- Changing the machine type can affect Google Cloud service charges.
- The operation restarts Cloud Volumes ONTAP.

For single node systems, I/O is interrupted.

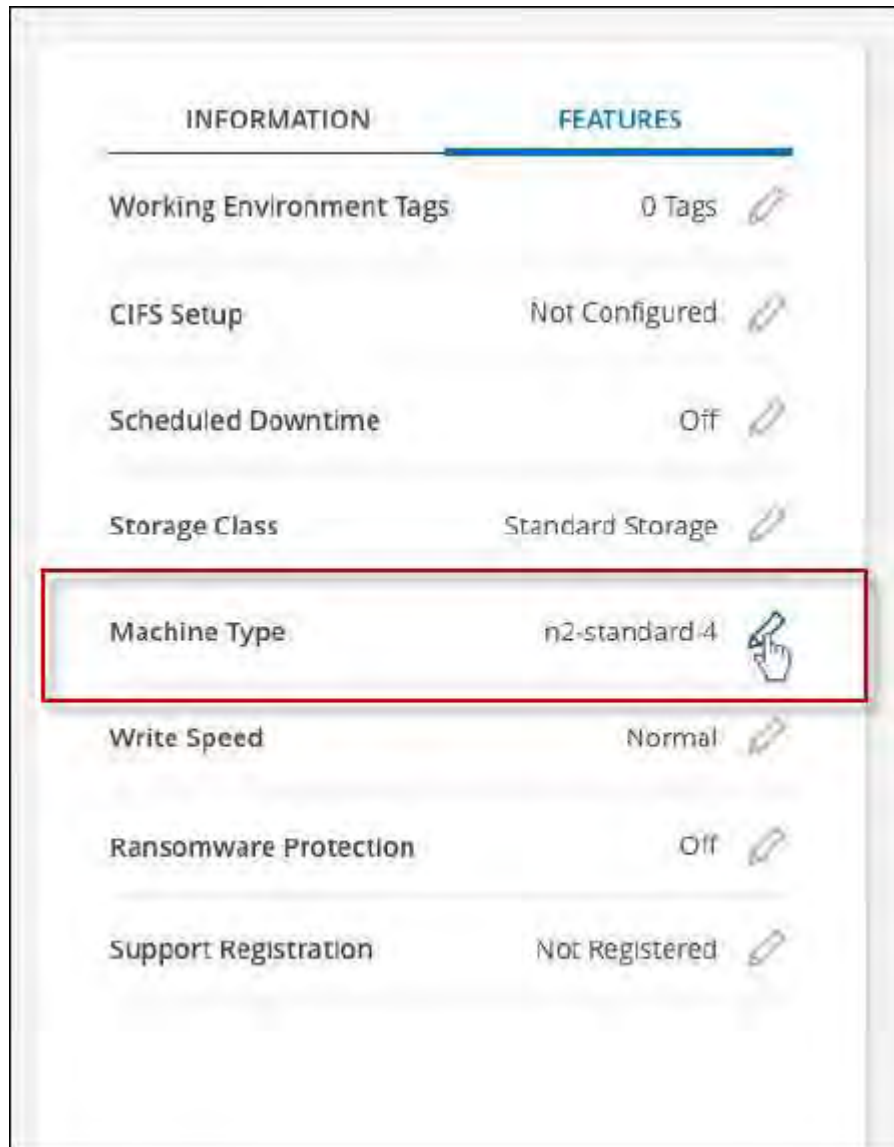
For HA pairs, the change is nondisruptive. HA pairs continue to serve data.



BlueXP changes one node at a time by initiating takeover and waiting for give back. NetApp's Quality Assurance team tested both writing and reading files during this process and didn't see any issues on the client side. As connections changed, some retries were observed on the I/O level, but the application layer overcame the rewiring of NFS/CIFS connections.

## Steps

1. On the Canvas page, select the working environment.
2. On the Overview tab, click the Features panel and then click the pencil icon next to **Machine type**.



If you are using a node-based pay-as-you-go (PAYGO) license, you can optionally choose a different license and machine type by clicking the pencil icon next to **License type**.

1. Choose an machine type, select the check box to confirm that you understand the implications of the change, and then click **Change**.

### Result

Cloud Volumes ONTAP reboots with the new configuration.

## Administer Cloud Volumes ONTAP using System Manager

Advanced storage management capabilities in Cloud Volumes ONTAP are available through ONTAP System Manager, a management interface provided with ONTAP systems. You can access System Manager directly from BlueXP.



## Features

You can perform various storage management functions using ONTAP System Manager in BlueXP. The following list includes some of those functionalities, though this list is not exhaustive:

- Advanced storage management: Manage consistency groups, shares, qtrees, quotas, and Storage VMs.
- Volume move: [Move a volume to a different aggregate](#).
- Networking management: Manage IPspaces, network interfaces, portsets, and ethernet ports.
- Manage FlexGroup volumes: You can create and manage FlexGroup volumes only through System Manager. BlueXP does not support FlexGroup volume creation.
- Events and jobs: View event logs, system alerts, jobs, and audit logs.
- Advanced data protection: Protect storage VMs, LUNs, and consistency groups.
- Host management: Set up SAN initiator groups and NFS clients.
- S3 object storage management: S3 storage management capabilities in Cloud Volumes ONTAP are available only in System Manager, and not in BlueXP.

## Supported configurations

- Advanced storage management through ONTAP System Manager is available in Cloud Volumes ONTAP 9.10.0 and later in standard cloud regions.
- System Manager integration is not supported in GovCloud regions or in regions that have no outbound internet access.

## Limitations

A few features that appear in the System Manager interface are not supported with Cloud Volumes ONTAP:

- BlueXP tiering: Cloud Volumes ONTAP does not support the BlueXP tiering service. You should set up tiering of data to object storage directly from BlueXP's Standard View when creating volumes.
- Tiers: Aggregate management (including local tiers and cloud tiers) is not supported from System Manager. You must manage aggregates directly from BlueXP's Standard View.
- Firmware upgrades: Cloud Volumes ONTAP does not support automatic firmware updates from the **Cluster > Settings** page.
- Role-based access control: Role-based access control from System Manager is not supported.
- SMB Continuous Availability (CA): Cloud Volumes ONTAP does not support [continuously available SMB shares](#) for nondisruptive operations.

## Configure authentication for accessing System Manager

As an administrator, you can activate authentication for users accessing ONTAP System Manager from BlueXP. You can determine the right level of access permissions based on the ONTAP user roles, and enable or disable authentication as needed. If you enable authentication, then users need to enter their ONTAP user credentials every time they access System Manager from BlueXP or when the page is reloaded, because BlueXP doesn't store the credentials internally. If you disable authentication, users can access System Manager using BlueXP admin credentials.



This setting is applicable per BlueXP Connector for the ONTAP users in your organization or account, irrespective of Cloud Volumes ONTAP working environments or BlueXP projects.

## Required permissions

You need to be assigned BlueXP Organization or Account admin privileges to modify the BlueXP Connector settings for Cloud Volumes ONTAP user authentication.

## Steps

1. Select the **Connector** drop-down from the BlueXP header.
2. Select **Manage Connectors**.
3. Click the action menu **...** for the required Connector and select **Edit Connector**.
4. Under **Force user credentials**, select the **Enable/Disable** check box. By default, authentication is disabled.



If you set this value to **Enable**, authentication is reset, and you have to modify any existing workflows to accommodate this change.

5. Click **Save**.

## Get started with System Manager

You can access ONTAP System Manager from a Cloud Volumes ONTAP working environment.

## Steps

1. From the left navigation menu, select **Storage > Canvas**.
2. On the Canvas page, select a Cloud Volumes ONTAP system.
3. From the right panel, select **Services > System Manager > Open**.
4. If prompted, enter your ONTAP user credentials and click **Login**.
5. If the confirmation message appears, read through it and click **Close**.
6. Use System Manager to manage Cloud Volumes ONTAP.
7. If needed, click **Switch to Standard View** to return to standard management through BlueXP.

## Help with using System Manager

If you need help using System Manager with Cloud Volumes ONTAP, you can refer to the [ONTAP documentation](#) for step-by-step instructions. Here are a few ONTAP documentation links that might help:

- [ONTAP roles, applications, and authentication](#)
- [Use System Manager to access a cluster](#)
- [Volume and LUN management](#)
- [Network management](#)
- [Data protection](#)
- [Create continuously available SMB shares](#)

## Administer Cloud Volumes ONTAP from the CLI

The Cloud Volumes ONTAP CLI enables you to run all administrative commands and is a good choice for advanced tasks or if you are more comfortable using the CLI. You can connect to the CLI using Secure Shell (SSH).

## Before you begin

The host from which you use SSH to connect to Cloud Volumes ONTAP must have a network connection to Cloud Volumes ONTAP. For example, you might need to SSH from a jump host that's in your cloud provider network.




When deployed in multiple AZs, Cloud Volumes ONTAP HA configurations use a floating IP address for the cluster management interface, which means external routing is not available. You must connect from a host that is part of the same routing domain.

## Steps

1. In BlueXP, identify the IP address of the cluster management interface:
  - a. From the left navigation menu, select **Storage > Canvas**.
  - b. On the Canvas page, select the Cloud Volumes ONTAP system.
  - c. Copy the cluster management IP address that appears in the right pane.
2. Use SSH to connect to the cluster management interface IP address using the admin account.

## Example

The following image shows an example using PuTTY:



3. At the login prompt, enter the password for the admin account.

## Example

```
Password: *****  
COT2:::>
```

# System health and events

## Verify AutoSupport setup for Cloud Volumes ONTAP

AutoSupport proactively monitors the health of your system and sends messages to NetApp technical support. By default, AutoSupport is enabled on each node to send messages to technical support using the HTTPS transport protocol. It's best to verify that AutoSupport can send these messages.

The only required configuration step is to ensure that Cloud Volumes ONTAP has outbound internet connectivity. For details, refer to the networking requirements for your cloud provider.

## AutoSupport requirements

Cloud Volumes ONTAP nodes require outbound internet access for NetApp AutoSupport, which proactively monitors the health of your system and sends messages to NetApp technical support.

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

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

If an outbound internet connection isn't available to send AutoSupport messages, BlueXP automatically configures your Cloud Volumes ONTAP systems to use the Connector as a proxy server. The only requirement is to ensure that the Connector's security group allows *inbound* connections over port 3128. You'll need to open this port after you deploy the Connector.

If you defined strict outbound rules for Cloud Volumes ONTAP, then you'll also need to ensure that the Cloud Volumes ONTAP security group allows *outbound* connections over port 3128.



If you're using an HA pair, the HA mediator doesn't require outbound internet access.

After you've verified that outbound internet access is available, you can test AutoSupport to ensure that it can send messages. For instructions, refer to the [ONTAP documentation: Set up AutoSupport](#).

## Troubleshoot your AutoSupport configuration

If an outbound connection isn't available and BlueXP can't configure your Cloud Volumes ONTAP system to use the Connector as a proxy server, you'll receive a notification from BlueXP titled "<working environment name> is unable to send AutoSupport messages."

You're most likely receiving this message because of networking issues.

Follow these steps to address this problem.

### Steps

1. SSH to the Cloud Volumes ONTAP system so that you can administer the system from the ONTAP CLI.

[Learn how to SSH to Cloud Volumes ONTAP.](#)

2. Display the detailed status of the AutoSupport subsystem:

```
autosupport check show-details
```

The response should be similar to the following:

```

Category: smtp
  Component: mail-server
  Status: failed
  Detail: SMTP connectivity check failed for destination:
         mailhost. Error: Could not resolve host -
'mailhost'
  Corrective Action: Check the hostname of the SMTP server

Category: http-https
  Component: http-put-destination
  Status: ok
  Detail: Successfully connected to:
         <https://support.netapp.com/put/AsupPut/>.

  Component: http-post-destination
  Status: ok
  Detail: Successfully connected to:

https://support.netapp.com/asupprod/post/1.0/postAsup.

Category: on-demand
  Component: ondemand-server
  Status: ok
  Detail: Successfully connected to:
         https://support.netapp.com/aods/asupmessage.

Category: configuration
  Component: configuration
  Status: ok
  Detail: No configuration issues found.
5 entries were displayed.

```

If the status of the http-https category is "ok" then it means AutoSupport is configured properly and messages can be sent.

3. If the status is not ok, verify the proxy URL for each Cloud Volumes ONTAP node:

```
autosupport show -fields proxy-url
```

4. If the proxy URL parameter is empty, configure Cloud Volumes ONTAP to use the Connector as a proxy:

```
autosupport modify -proxy-url http://<connector private ip>:3128
```

5. Verify AutoSupport status again:

```
autosupport check show-details
```

6. If the status is still failed, validate that there is connectivity between Cloud Volumes ONTAP and the Connector over port 3128.
7. If the status ID is still failed after verifying that there is connectivity, SSH to the Connector.

[Learn more about Connecting to the Linux VM for the Connector](#)

8. Go to `/opt/application/netapp/cloudmanager/docker_occm/data/`
9. Open the proxy configuration file `squid.conf`

The basic structure of the file is as follows:

```
http_port 3128
acl localnet src 172.31.0.0/16
acl azure_aws_metadata dst 169.254.169.254

http_access allow localnet
http_access deny azure_aws_metadata
http_access allow localhost
http_access deny all
```

The `localnet` `src` value is the CIDR of the Cloud Volumes ONTAP system.

10. If the CIDR block of the Cloud Volumes ONTAP system isn't in the range that's specified in the file, either update the value or add a new entry as follows:

```
acl cvonet src <cidr>
```

If you add this new entry, don't forget to also add an allow entry:

```
http_access allow cvonet
```

Here's an example:

```
http_port 3128
acl localnet src 172.31.0.0/16
acl cvonet src 172.33.0.0/16
acl azure_aws_metadata dst 169.254.169.254

http_access allow localnet
http_access allow cvonet
http_access deny azure_aws_metadata
http_access allow localhost
http_access deny all
```

11. After editing the config file, restart the proxy container as `sudo`:

```
docker restart squid
```

12. Go back to the Cloud Volumes ONTAP CLI and verify that Cloud Volumes ONTAP can send AutoSupport messages:

```
autosupport check show-details
```

## Configure EMS for Cloud Volumes ONTAP systems

### Related links

The Event Management System (EMS) collects and displays information about events that occur on ONTAP systems. To receive event notifications, you can set event destinations (email addresses, SNMP trap hosts, or syslog servers) and event routes for a particular event severity.

You can configure EMS using the CLI. For instructions, refer to the [ONTAP documentation: EMS configuration overview](#).

# Concepts

## Licensing

### Licensing for Cloud Volumes ONTAP

Several licensing options are available for Cloud Volumes ONTAP. Each option enables you to choose a consumption model that meets your needs.

#### Licensing overview

The following licensing options are available for new customers.

#### Capacity-based licensing

Pay for multiple Cloud Volumes ONTAP systems in your NetApp account by provisioned capacity. Includes the ability to purchase add-on cloud data services. For more information about consumption models in capacity-based licenses, refer to [Learn more about capacity-based licenses](#).

#### Keystone Subscription

A pay-as-you-grow subscription-based service that delivers a seamless hybrid cloud experience for High Availability (HA) pairs.

The following sections provide more details about each of these options.



Support is not available for the use of licensed features without a license.

#### Capacity-based licensing

Capacity-based licensing packages enable you to pay for Cloud Volumes ONTAP per TiB of capacity. The license is associated with your NetApp account and enables you to charge multiple systems against the license, as long as enough capacity is available through the license.

For example, you could purchase a single 20 TiB license, deploy four Cloud Volumes ONTAP systems, and then allocate a 5 TiB volume to each system, for a total of 20 TiB. The capacity is available to the volumes on each Cloud Volumes ONTAP system deployed in that account.

Capacity-based licensing is available in the form of a *package*. When you deploy a Cloud Volumes ONTAP system, you can choose from several licensing packages based on your business needs.



While the actual usage and metering for the products and services managed in BlueXP are always calculated in GiB and TiB, the terms GB/GiB and TB/TiB are used interchangeably. This is reflected in the Cloud marketplace listings, price quotes, listing descriptions, and in other supporting documentation.

#### Packages

The following capacity-based packages are available for Cloud Volumes ONTAP. For more information about capacity-based license packages, refer to [Learn more about capacity-based licenses](#).

For a list of supported VM types with the following capacity-based packages, refer to:



- [Supported configurations in Azure](#)
- [Supported configurations in Google Cloud](#)

## Freemium

Provides all Cloud Volumes ONTAP features free of charge from NetApp (cloud provider charges still apply). A Freemium package has these characteristics:

- No license or contract is needed.
- Support from NetApp is not included.
- You're limited to 500 GiB of provisioned capacity per Cloud Volumes ONTAP system.
- You can use up to 10 Cloud Volumes ONTAP systems with the Freemium offering per NetApp account, for any cloud provider.
- If the provisioned capacity for a Cloud Volumes ONTAP system exceeds 500 GiB, BlueXP converts the system to an Essentials package.

As soon as a system is converted to the Essentials package, [minimum charging](#) applies to it.

A Cloud Volumes ONTAP system that has been converted into an Essentials package cannot be switched back to Freemium even if the provisioned capacity is reduced to less than 500 GiB. Other systems with less than 500 GiB of provisioned capacity stay on Freemium (as long as they were deployed using the Freemium offering).

## Essentials

You can pay by capacity in a number of different configurations:

- Choose your Cloud Volumes ONTAP configuration:
  - A single node or HA system
  - File and block storage or secondary data for disaster recovery (DR)
- Add on any of NetApp's cloud data services at extra cost

## Professional

Pay by capacity for any type of Cloud Volumes ONTAP configuration with unlimited backups.

- Provides licensing for any Cloud Volumes ONTAP configuration

Single node or HA with capacity charging for primary and secondary volumes at the same rate

- Includes unlimited volume backups using BlueXP backup and recovery, but only for Cloud Volumes ONTAP systems that use the Professional package.



A pay-as-you-go (PAYGO) subscription is required for BlueXP backup and recovery, however no charges will be incurred for using this service. For more information on setting up licensing for BlueXP backup and recovery, refer to [Set up licensing for BlueXP backup and recovery](#).

- Add on any of NetApp's cloud data services at extra cost

## Availability of capacity-based licenses

The availability of the PAYGO and BYOL licenses for Cloud Volumes ONTAP systems requires the BlueXP Connector to be up and running. For more information, refer to [Learn about Connectors](#).



NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAPP](#).

## How to get started

Learn how to get started with capacity-based licensing:

- [Set up licensing for Cloud Volumes ONTAP in AWS](#)
- [Set up licensing for Cloud Volumes ONTAP in Azure](#)
- [Set up licensing for Cloud Volumes ONTAP in Google Cloud](#)

## Keystone Subscription

A pay-as-you-grow subscription-based service that delivers a seamless hybrid cloud experience for those preferring OpEx consumption models to upfront CapEx or leasing.

Charging is based on the size of your committed capacity for one or more Cloud Volumes ONTAP HA pairs in your Keystone Subscription.

The provisioned capacity for each volume is aggregated and compared to the committed capacity on your Keystone Subscription periodically, and any overages are charged as burst on your Keystone Subscription.

[Learn more about NetApp Keystone](#).

## Supported configurations

Keystone Subscriptions are supported with HA pairs. This licensing option isn't supported with single node systems at this time.

## Capacity limit

In the capacity-based licensing model, each Cloud Volumes ONTAP system supports tiering to object storage, and the total tiered capacity can scale up to the cloud provider's bucket limit. Although the license does not impose capacity restrictions, follow the [FabricPool Best Practices](#) to ensure optimal performance, reliability, and cost efficiency when configuring and managing tiering.

For information about the capacity limits of each cloud provider, refer to their documentation:

- [AWS documentation](#)
- [Azure documentation for managed disks](#) and [Azure documentation for blob storage](#)
- [Google Cloud documentation](#)

## How to get started

Learn how to get started with a Keystone Subscription:

- [Set up licensing for Cloud Volumes ONTAP in AWS](#)
- [Set up licensing for Cloud Volumes ONTAP in Azure](#)

- [Set up licensing for Cloud Volumes ONTAP in Google Cloud](#)

## Node-based licensing

Node-based licensing is the previous generation licensing model that enabled you to license Cloud Volumes ONTAP by node. This licensing model is not available for new customers. By-node charging has been replaced with the by-capacity charging methods described above.

NetApp has planned the end of availability (EOA) and support (EOS) of node-based licensing. After the EOA and EOS, node-based licenses will need to be converted to capacity-based licenses.

For information, refer to [Customer communique: CPC-00589](#).

### End of availability of node-based licenses

Beginning on 11 November, 2024, the limited availability of node-based licenses has been terminated. The support for node-based licensing ends on 31 December, 2024.

If you have a valid node-based contract that extends beyond the EOA date, you can continue to use the license until the contract expires. Once the contract expires, it will be necessary to transition to the capacity-based licensing model. If you don't have a long-term contract for a Cloud Volumes ONTAP node, it is important to plan your conversion before the EOS date.

Learn more about each license type and the impact of EOA on it from this table:

License type	Impact after EOA
Valid node-based license purchased through bring your own license (BYOL)	License remains valid till expiration. Existing unused node-based licenses can be used for deploying new Cloud Volumes ONTAP systems.
Expired node-based license purchased through BYOL	You won't be entitled to deploy new Cloud Volumes ONTAP systems using this license. The existing systems might continue to work, but you won't receive any support or updates for your systems post the EOS date.
Valid node-based license with PAYGO subscription	Will cease to receive NetApp support post the EOS date, until you transition to a capacity-based license.

### Exclusions

NetApp recognizes that certain situations require special consideration, and EOA and EOS of node-based licensing will not apply to the following cases:

- U.S. Public Sector customers
- Deployments in private mode
- China region deployments of Cloud Volumes ONTAP in AWS

For these particular scenarios, NetApp will offer support to address the unique licensing requirements in compliance with contractual obligations and operational needs.



Even in these scenarios, new node-based licenses and license renewals are valid for a maximum of one year from the date of approval.

## License conversion

BlueXP enables a seamless conversion of node-based licenses to capacity based through the license conversion tool. For information about EOA of node-based licensing, refer to [End of availability of node-based licenses](#).

Before transitioning, it is good to familiarize yourself with the difference between the two licensing models. Node-based licensing includes fixed capacity for each ONTAP instance, which can restrict flexibility. Capacity-based licensing, on the other hand, allows for a shared pool of storage across multiple instances, offering enhanced flexibility, optimizing resource utilization, and reducing the potential for financial penalties when redistributing workloads. Capacity-based charging seamlessly adjusts to changing storage requirements.

To know how you can perform this conversion, refer to [Convert a Cloud Volumes ONTAP node-based license to capacity-based license](#).



Conversion of a system from capacity-based to node-based licensing is not supported.

## Learn more about capacity-based licenses for Cloud Volumes ONTAP

You should be familiar with the charging and capacity usage for capacity-based licenses

### Consumption models

Capacity-based licensing packages are available with the following consumption models:

- **BYOL**: Bring your own license (BYOL). A license purchased from NetApp that can be used to deploy Cloud Volumes ONTAP in any cloud provider.



NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAP](#).

- **PAYGO**: A pay-as-you-go (PAYGO) subscription is an hourly subscription from your cloud provider's marketplace.
- **Annual**: An annual contract from your cloud provider's marketplace.

Note the following:

- If you purchase a license from NetApp (BYOL), you also need to subscribe to the PAYGO offering from your cloud provider's marketplace. NetApp has restricted BYOL licensing. When your BYOL licenses expire, you are required to replace them with cloud marketplace subscriptions.

Your license is always charged first, but you'll be charged from the hourly rate in the marketplace in these cases:

- If you exceed your licensed capacity
- If the term of your license expires
- If you have an annual contract from a marketplace, *all* Cloud Volumes ONTAP systems that you deploy are charged against that contract. You can't mix and match an annual marketplace contract with BYOL.

- Only single node systems with BYOL are supported in China regions. China region deployments are exempt from BYOL licensing restrictions.

## Changing packages

After deployment, you can change the package for a Cloud Volumes ONTAP system that uses capacity-based licensing. For example, if you deployed a Cloud Volumes ONTAP system with the Essentials package, you can change it to the Professional package if your business needs changed.

[Learn how to change charging methods.](#)

For information about converting node-based licenses to capacity-based, see

## Pricing and supported configurations

For details about pricing, go to the [NetApp BlueXP website](#).

Capacity-based licensing packages are available with Cloud Volumes ONTAP 9.7 and later.

## Storage VMs

- There are no extra licensing costs for additional data-serving storage VMs (SVMs), but there is a 4 TiB minimum capacity charge per data-serving SVM.
- Disaster recovery SVMs are charged according to the provisioned capacity.

## HA pairs

For HA pairs, you're only charged for the provisioned capacity on a node. You aren't charged for data that is synchronously mirrored to the partner node.

## FlexClone and FlexCache volumes

- You won't be charged for the capacity used by FlexClone volumes.
- Source and destination FlexCache volumes are considered primary data and charged according to the provisioned space.

## Capacity limit

In the capacity-based licensing model, each Cloud Volumes ONTAP system supports tiering to object storage, and the total tiered capacity can scale up to the cloud provider's bucket limit. Although the license does not impose capacity restrictions, follow the [FabricPool Best Practices](#) to ensure optimal performance, reliability, and cost efficiency when configuring and managing tiering.

For information about the capacity limits of each cloud provider, refer to their documentation:

- [AWS documentation](#)
- [Azure documentation for managed disks](#) and [Azure documentation for blob storage](#)
- [Google Cloud documentation](#)

## Max number of systems

With capacity-based licensing, the maximum number of Cloud Volumes ONTAP systems is limited to 24 per BlueXP account. A *system* is a Cloud Volumes ONTAP HA pair, a Cloud Volumes ONTAP single node system,

or any additional storage VMs that you create. The default storage VM does not count against the limit. This limit applies to all licensing models.

For example, let's say you have three working environments:

- A single node Cloud Volumes ONTAP system with one storage VM (this is the default storage VM that's created when you deploy Cloud Volumes ONTAP)

This working environment counts as one system.

- A single node Cloud Volumes ONTAP system with two storage VMs (the default storage VM, plus one additional storage VM that you created)

This working environment counts as two systems: one for the single node system and one for the additional storage VM.

- A Cloud Volumes ONTAP HA pair with three storage VMs (the default storage VM, plus two additional storage VMs that you created)

This working environment counts as three systems: one for the HA pair and two for the additional storage VMs.

That's six systems in total. You would then have room for an additional 14 systems in your account.

If you have a large deployment that requires more than 24 systems, contact your account rep or sales team.

[Learn more about BlueXP accounts.](#)

[Learn about storage limits for AWS, Azure, and Google Cloud.](#)

## Notes about charging

The following details can help you understand how charging works with capacity-based licensing.

### Minimum charge

There is a 4 TiB minimum charge for each data-serving storage VM that has at least one primary (read-write) volume. If the sum of the primary volumes is less than 4 TiB, then BlueXP applies the 4 TiB minimum charge to that storage VM.

If you haven't provisioned any volumes yet, then the minimum charge doesn't apply.

For the Essentials package, the 4 TiB minimum capacity charge doesn't apply to storage VMs that contain secondary (data protection) volumes only. For example, if you have a storage VM with 1 TiB of secondary data, then you're charged just for that 1 TiB of data. With the Professional package type, the minimum capacity charging of 4 TiB applies regardless of the volume type.

### Overages

If you exceed your BYOL capacity, you'll be charged for overages at hourly rates based on your marketplace subscription. Overages are charged at marketplace rates, with a preference for using available capacity from other licenses first. If your BYOL license expires, you need to transition to a capacity-based licensing model through cloud marketplaces.

## Essentials package

With the Essentials package, you're billed by the deployment type (HA or single node) and the volume type (primary or secondary). Pricing from high to low is in the following order: *Essentials Primary HA*, *Essentials Primary Single Node*, *Essentials Secondary HA*, and *Essentials Secondary Single Node*. Alternately, when you purchase a marketplace contract or accept a private offer, capacity charges are the same for any deployment or volume type.

Licensing is based entirely on the volume type created within Cloud Volumes ONTAP systems:

- Essentials Single Node: Read/write volumes created on a Cloud Volumes ONTAP system using one ONTAP node only.
- Essentials HA: Read/write volumes using two ONTAP nodes that can fail over to each other for non-disruptive data access.
- Essentials Secondary Single Node: Data Protection (DP) type volumes (typically SnapMirror or SnapVault destination volumes that are read-only) created on a Cloud Volumes ONTAP system using one ONTAP node only.



If a read-only/DP volume becomes a primary volume, BlueXP considers it as primary data and the charging costs are calculated based on the time the volume was in read/write mode. When the volume is again made read-only/DP, BlueXP considers it as secondary data again and charges accordingly using the best matching license in the digital wallet.

- Essentials Secondary HA: Data Protection (DP) type volumes (typically SnapMirror or SnapVault destination volumes that are read-only) created on a Cloud Volumes ONTAP system using two ONTAP nodes that can fail over to each other for non-disruptive data access.

## BYOL

If you purchased an Essentials license from NetApp (BYOL) and you exceed the licensed capacity for that deployment and volume type, the BlueXP digital wallet charges overages against a higher priced Essentials license (if you have one and there is available capacity). This happens because we first use the available capacity that you've already purchased as prepaid capacity before charging against the marketplace. If there is no available capacity with your BYOL license, the exceeded capacity will be charged at marketplace on-demand hourly rates (PAYGO) and will add costs to your monthly bill.

Here's an example. Let's say you have the following licenses for the Essentials package:

- A 500 TiB *Essentials Secondary HA* license that has 500 TiB of committed capacity
- A 500 TiB *Essentials Single Node* license that only has 100 TiB of committed capacity

Another 50 TiB is provisioned on an HA pair with secondary volumes. Instead of charging that 50 TiB to PAYGO, the BlueXP digital wallet charges the 50 TiB overage against the *Essentials Single Node* license. That license is priced higher than *Essentials Secondary HA*, but it's making use of a license you have already purchased, and it will not add costs to your monthly bill.

In the BlueXP digital wallet, that 50 TiB will be shown as charged against the *Essentials Single Node* license.

Here's another example. Let's say you have the following licenses for the Essentials package:

- A 500 TiB *Essentials Secondary HA* license that has 500 TiB of committed capacity
- A 500 TiB *Essentials Single Node* license that only has 100 TiB of committed capacity

Another 100 TiB is provisioned on an HA pair with primary volumes. The license you purchased doesn't have

*Essentials Primary HA* committed capacity. The *Essentials Primary HA* license is priced higher than both the *Essentials Primary Single Node* and *Essentials Secondary HA* licenses.

In this example, the BlueXP digital wallet charges overages at the marketplace rate for the additional 100 TiB. The overage charges will appear on your monthly bill.

### **Marketplace contracts or private offers**

If you purchased an Essentials license as part of a marketplace contract or a private offer, the BYOL logic does not apply, and you must have the exact license type for the usage. License type includes volume type (primary or secondary) and the deployment type (HA or single node).

For example, let's say you deploy a Cloud Volumes ONTAP instance with the Essentials license. You then provision read-write volumes (primary single node) and read-only (secondary single node) volumes. Your marketplace contract or private offer must include capacity for *Essentials Single Node* and *Essentials Secondary Single Node* to cover the provisioned capacity. Any provisioned capacity that isn't part of your marketplace contract or private offer will be charged at the on-demand hourly rates (PAYGO) and will add costs to your monthly bill.

## **Storage**

### **Supported client protocols for Cloud Volumes ONTAP**

Cloud Volumes ONTAP supports the iSCSI, NFS, SMB, NVMe-TCP, and S3 client protocols.

#### **iSCSI**

iSCSI is a block protocol that can run on standard Ethernet networks. Most client operating systems offer a software initiator that runs over a standard Ethernet port.

#### **NFS**

NFS is the traditional file access protocol for UNIX and LINUX systems. Clients can access files in ONTAP volumes using the NFSv3, NFSv4, and NFSv4.1 protocols. You can control file access using UNIX-style permissions, NTFS-style permissions, or a mix of both.

Clients can access the same files using both NFS and SMB protocols.

#### **SMB**

SMB is the traditional file access protocol for Windows systems. Clients can access files in ONTAP volumes using the SMB 2.0, SMB 2.1, SMB 3.0, and SMB 3.1.1 protocols. Just like with NFS, a mix of permission styles are supported.

#### **S3**

Cloud Volumes ONTAP supports S3 as an option for scale-out storage. S3 protocol support enables you to configure S3 client access to objects contained in a bucket in a storage VM (SVM).

[ONTAP documentation: Learn how S3 multiprotocol works.](#)

[ONTAP documentation: Learn how to configure and manage S3 object storage services in ONTAP.](#)



## NVMe-TCP

Beginning with ONTAP version 9.12.1, NVMe-TCP is supported for cloud providers. BlueXP does not provide any management capabilities for NVMe-TCP.

For more information on configuring NVMe through ONTAP, refer to the [ONTAP documentation: Configure a storage VM for NVMe](#).

## Disks and aggregates used for Cloud Volumes ONTAP clusters

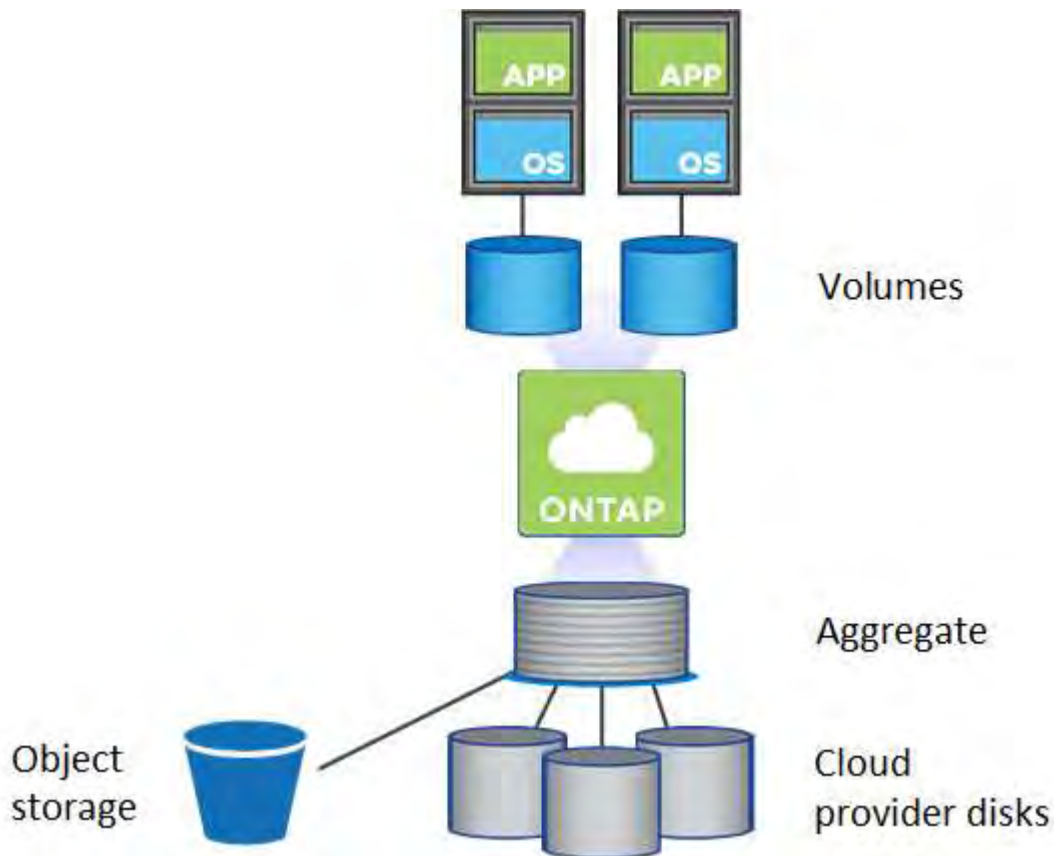
Understanding how Cloud Volumes ONTAP uses cloud storage can help you understand your storage costs.



All disks and aggregates must be created and deleted directly from BlueXP. You should not perform these actions from another management tool. Doing so can impact system stability, hamper the ability to add disks in the future, and potentially generate redundant cloud provider fees.

### Overview

Cloud Volumes ONTAP uses cloud provider storage as disks and groups them into one or more aggregates. Aggregates provide storage to one or more volumes.



Several types of cloud disks are supported. You choose the disk type when you create a volume and the default disk size when you deploy Cloud Volumes ONTAP.



The total amount of storage purchased from a cloud provider is the *raw capacity*. The *usable capacity* is less because approximately 12 to 14 percent is overhead that is reserved for Cloud Volumes ONTAP use. For example, if BlueXP creates a 500 GiB aggregate, the usable capacity is 442.94 GiB.

## AWS storage

In AWS, Cloud Volumes ONTAP uses EBS storage for user data and local NVMe storage as Flash Cache on some EC2 instance types.

### EBS storage

In AWS, an aggregate can contain up to 6 disks that are all the same size. But if you have a configuration that supports the Amazon EBS Elastic Volumes feature, then an aggregate can contain up to 8 disks. [Learn more about support for Elastic Volumes](#).

The maximum disk size is 16 TiB.

The underlying EBS disk type can be either General Purpose SSDs (gp3 or gp2), Provisioned IOPS SSD (io1), or Throughput Optimized HDD (st1). You can pair an EBS disk with Amazon S3 to [low-cost object storage](#).



Tiering data to object storage is not recommended when using Throughput Optimized HDDs (st1).

## Local NVMe storage

Some EC2 instance types include local NVMe storage, which Cloud Volumes ONTAP uses as [Flash Cache](#).

## Related links

- [AWS documentation: EBS Volume Types](#)
- [Learn how to choose disk types and disk sizes for your systems in AWS](#)
- [Review storage limits for Cloud Volumes ONTAP in AWS](#)
- [Review supported configurations for Cloud Volumes ONTAP in AWS](#)

## Azure storage

In Azure, an aggregate can contain up to 12 disks that are all the same size. The disk type and maximum disk size depends on whether you use a single node system or an HA pair:

### Single node systems

Single node systems can use these types of Azure Managed Disks:

- *Premium SSD Managed Disks* provide high performance for I/O-intensive workloads at a higher cost.
- *Premium SSD v2 Managed Disks* provide higher performance with lower latency at a lower cost for both single node and HA pairs, compared to Premium SSD Managed Disks.
- *Standard SSD Managed Disks* provide consistent performance for workloads that require low IOPS.
- *Standard HDD Managed Disks* are a good choice if you don't need high IOPS and want to reduce your costs.

Each managed disk type has a maximum disk size of 32 TiB.

You can pair a managed disk with Azure Blob storage to [low-cost object storage](#).

## HA pairs

HA pairs use two types of disks which provide high performance for I/O-intensive workloads at a higher cost:

- *Premium page blobs* with a maximum disk size of 8 TiB
- *Managed disks* with a maximum disk size of 32 TiB

## Related links

- [Learn how to choose disk types and disk sizes for your systems in Azure](#)
- [Launch a Cloud Volumes ONTAP HA pair in Azure](#)
- [Microsoft Azure documentation: Azure managed disk types](#)
- [Microsoft Azure documentation: Overview of Azure page blobs](#)
- [Review storage limits for Cloud Volumes ONTAP in Azure](#)

## Google Cloud storage

In Google Cloud, an aggregate can contain up to 6 disks that are all the same size. The maximum disk size is 64 TiB.

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

## Related links

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

## RAID type

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

## Hot spares

RAID0 doesn't support the use of hot spares for redundancy.

Creating unused disks (hot spares) attached to a Cloud Volumes ONTAP instance is an unnecessary expense and may prevent provisioning additional space as needed. Therefore, it's not recommended.

## Learn about support for AWS Elastic Volumes with Cloud Volumes ONTAP

Support for the Amazon EBS Elastic Volumes feature with a Cloud Volumes ONTAP aggregate provides better performance and additional capacity, while enabling BlueXP to automatically increase the underlying disk capacity as needed.

## Benefits

- Dynamic disk growth

BlueXP can dynamically increase the size of disks while Cloud Volumes ONTAP is running and while disks are still attached.

- Better performance

Aggregates that are enabled with Elastic Volumes can have up to eight disks that are equally utilized across two RAID groups. This configuration provides more throughput and consistent performance.

- Larger aggregates

Support for eight disks provides a maximum aggregate capacity of 128 TiB. These limits are higher than the six disk limit and 96 TiB limit for aggregates that aren't enabled with the Elastic Volumes feature.

Note that total system capacity limits remain the same.

[AWS Documentation: Learn more about Elastic Volumes from AWS](#)

## Supported configurations

The Amazon EBS Elastic Volumes feature is supported with specific Cloud Volumes ONTAP versions and specific EBS disk types.

### Cloud Volumes ONTAP version

The Elastic Volumes feature is supported with *new* Cloud Volumes ONTAP systems created from version 9.11.0 or later. The feature is *not* supported with existing Cloud Volumes ONTAP systems that were deployed prior to 9.11.0.

For example, the Elastic Volumes feature is not supported if you created a Cloud Volumes ONTAP 9.9.0 system and then later upgraded that system to version 9.11.0. It must be a new system deployed using version 9.11.0 or later.

### EBS disk types

The Elastic Volumes feature is automatically enabled at the aggregate level when using General Purpose SSDs (gp3) or Provisioned IOPS SSDs (io1). The Elastic Volumes feature is not supported with aggregates that use any other disk type.

## Required AWS permissions


Starting with the 3.9.19 release, the Connector requires the following permissions to enable and manage the Elastic Volumes feature on a Cloud Volumes ONTAP aggregate:

- ec2:DescribeVolumesModifications
- ec2:ModifyVolume

These permissions are included in [the policies provided by NetApp](#)

## How support for Elastic Volumes works

An aggregate that has the Elastic Volumes feature enabled is comprised of one or two RAID groups. Each RAID group has four identical disks that have the same capacity. Here's an example of a 10 TiB aggregate that has four disks that are 2.5 TiB each:




When BlueXP creates an aggregate, it starts with one RAID group. If additional capacity is needed, BlueXP grows the aggregate by increasing the capacity of all disks in the RAID group by the same amount. The capacity increase is either a minimum of 256 GiB or 10% of the aggregate's size.

For example, if you have a 1 TiB aggregate, each disk is 250 GiB. 10% of the aggregate's capacity is 100 GiB. That's lower than 256 GiB, so the size of the aggregate is increased by the 256 GiB minimum (or 64 GiB for each disk).

BlueXP increases the size of the disks while the Cloud Volumes ONTAP system is running and while the disks are still attached. The change is non-disruptive.

If an aggregate reaches 64 TiB (or 16 TiB on each disk), BlueXP creates a second RAID group for additional capacity. This second RAID group works just like the first one: it has four disks that have the exact same capacity and it can grow up to 64 TiB. That means an aggregate can have a maximum capacity of 128 TiB.

Here's an example of an aggregate with two RAID groups. The capacity limit has been reached on the first RAID group, while the disks in the second RAID group have plenty of free space.



### What happens when you create a volume

If you create a volume that uses gp3 or io1 disks, BlueXP creates the volume on an aggregate as follows:

- If there is an existing gp3 or io1 aggregate that has Elastic Volumes enabled, BlueXP creates the volume on that aggregate.

- If there are multiple gp3 or io1 aggregates that have Elastic Volumes enabled, BlueXP creates the volume on the aggregate that requires the least amount of resources.
- If the system only has gp3 or io1 aggregates that aren't enabled for Elastic Volumes, then the volume is created on that aggregate.



While this scenario is unlikely, it's possible in two cases:

- You explicitly disabled the Elastic Volumes feature when creating an aggregate from the API.
- You created a new Cloud Volumes ONTAP system from the user interface, in which case the Elastic Volumes feature is disabled on the initial aggregate. Review [Limitations](#) below to learn more.

- If no existing aggregates have enough capacity, BlueXP creates the aggregate with Elastic Volumes enabled and then creates the volume on that new aggregate.

The size of the aggregate is based on the requested volume size plus an additional 10% capacity.

### Capacity Management Mode

The Capacity Management Mode for a Connector works with Elastic Volumes similar to how it works with other types of aggregates:

- When Automatic mode is enabled (this is the default setting), BlueXP automatically increases the size of aggregates if additional capacity is needed.
- If you change the capacity management mode to Manual, BlueXP asks for your approval to purchase additional capacity.

[Learn more about the Capacity Management Mode.](#)

### Limitations

Increasing the size of an aggregate can take up to 6 hours. During that time, BlueXP can't request any additional capacity for that aggregate.

### How to work with Elastic Volumes

You can work with Elastic Volumes in BlueXP as follows:

- Create a new system that has Elastic Volumes enabled on the initial aggregate when using gp3 or io1 disks

[Learn how to create Cloud Volumes ONTAP system](#)

- Create a new volume on an aggregate that has Elastic Volumes enabled

If you create a volume that uses gp3 or io1 disks, BlueXP automatically creates the volume on an aggregate that has Elastic Volumes enabled. For more details, refer to [What happens when you create a volume.](#)

[Learn how to create volumes.](#)

- Create a new aggregate that has Elastic Volumes enabled

Elastic Volumes is automatically enabled on new aggregates that use gp3 or io1 disks, as long as the Cloud Volumes ONTAP system was created from version 9.11.0 or later.

When you create the aggregate, BlueXP will prompt you for the aggregate's capacity size. This is different than other configurations where you choose a disk size and number of disks.


The following screenshot shows an example of a new aggregate comprised of gp3 disks.

1 Disk Type    2 Aggregate details    3 Tiering Data    4 Review

### Select Disk Type

Disk Type

GP3 - General Purpose SSD Dynamic Performance

 **General Purpose SSD (gp3) Disk Properties**

**Description:** General purpose SSD volume that balances price and performance (performance level is independent of storage capacity)

IOPS Value *i*    Throughput MB/s *i*

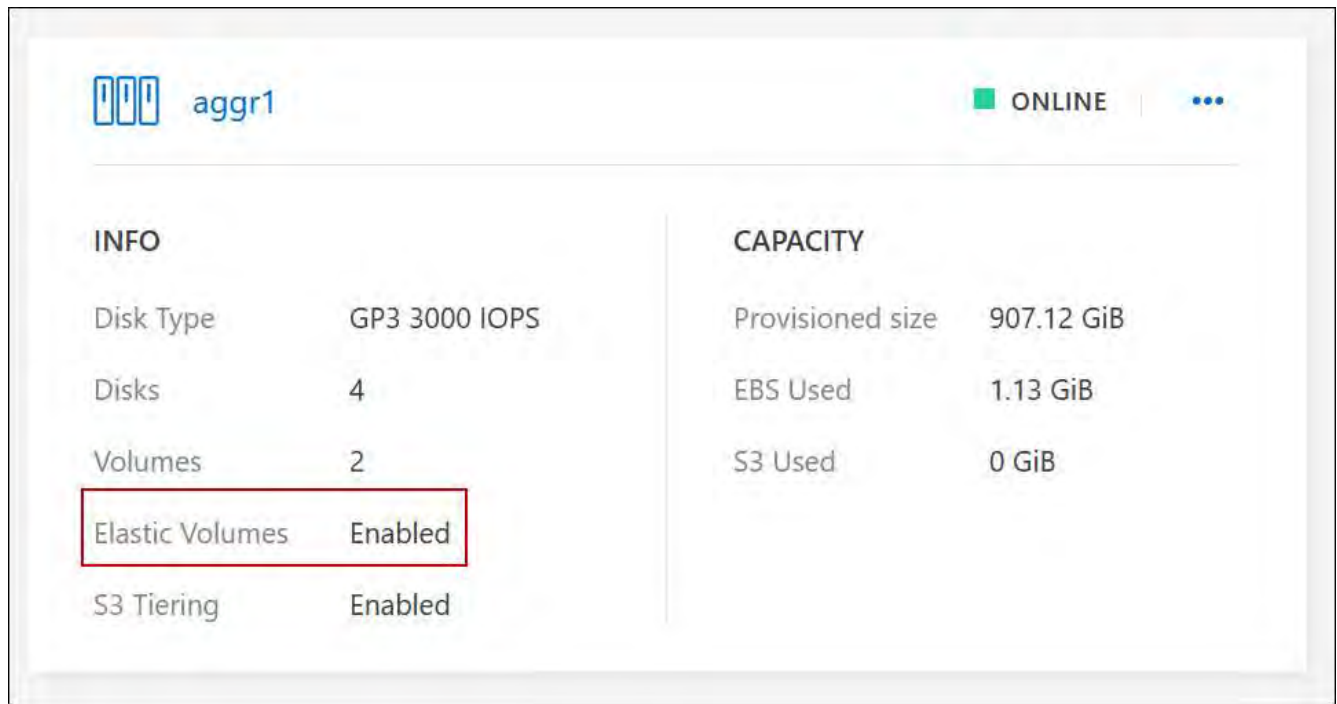
12000    250

[Learn how to create aggregates.](#)

- Identify aggregates that have Elastic Volumes enabled

When you go to the Advanced Allocation page, you can identify whether the Elastic Volumes feature is enabled on an aggregate. In the following example, aggr1 has Elastic Volumes enabled.





- Add capacity to an aggregate

While BlueXP automatically adds capacity to aggregates as needed, you can manually increase the capacity yourself.

[Learn how to increase aggregate capacity.](#)


- Replicate data to an aggregate that has Elastic Volumes enabled

If the destination Cloud Volumes ONTAP system supports Elastic Volumes, a destination volume will be placed on an aggregate that has Elastic Volumes enabled (as long as you choose a gp3 or io1 disk).

[Learn how to set up data replication](#)

## Learn about data tiering with Cloud Volumes ONTAP in AWS, Azure, or Google Cloud

Reduce your storage costs by enabling automated tiering of inactive data to low-cost object storage. Active data remains in high-performance SSDs or HDDs, while inactive data is tiered to low-cost object storage. This enables you to reclaim space on your primary storage and shrink secondary storage.



Data tiering is powered by FabricPool technology. Cloud Volumes ONTAP provides data tiering for all Cloud Volumes ONTAP clusters without an additional license. When you enable data tiering, data tiered to object storage incurs charges. Refer to your cloud provider’s documentation for details about object storage costs.

### Data tiering in AWS

When you enable data tiering in AWS, Cloud Volumes ONTAP uses EBS as a performance tier for hot data and AWS S3 as a capacity tier for inactive data.

#### Performance tier

The performance tier can be General Purpose SSDs (gp3 or gp2) or Provisioned IOPS SSDs (io1).

Tiering data to object storage is not recommended when using Throughput Optimized HDDs (st1).

#### Capacity tier

A Cloud Volumes ONTAP system tiers inactive data to a single S3 bucket.

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

When BlueXP creates the S3 bucket, it uses the following default settings:

- Storage class: Standard
- Default encryption: Disabled
- Block public access: Block all public access
- Object ownership: ACLs enabled
- Bucket versioning: Disabled
- Object lock: Disabled

#### Storage classes

The default storage class for tiered data in AWS is *Standard*. Standard is ideal for frequently accessed data stored across multiple Availability Zones.

If you don’t plan to access the inactive data, you can reduce your storage costs by changing the storage class to one of the following: *Intelligent Tiering*, *One-Zone Infrequent Access*, *Standard-Infrequent Access*, or *S3 Glacier Instant Retrieval*. When you change the storage class, inactive data starts in the Standard

storage class and transitions to the storage class that you selected, if the data is not accessed after 30 days.

Access costs are higher if you access the data, so consider this before changing the storage class. [Amazon S3 documentation: Learn more about Amazon S3 storage classes](#).

You can select a storage class when you create the working environment and you can change it any time afterwards. For instructions on changing the storage class, refer to [Tier inactive data to low-cost object storage](#).

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

## Data tiering in Azure

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

### Performance tier

The performance tier can be either SSDs or HDDs.

### Capacity tier

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

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

BlueXP creates the storage account with the following settings:

- Access tier: Hot
- Performance: Standard
- Redundancy: Accordingly to Cloud Volume ONTAP Deployment
  - Single availability zone: Locally-redundant storage (LRS)
  - Multiple availability zone: Zone-redundant storage (ZRS)
- Account: StorageV2 (general purpose v2)
- Require secure transfer for REST API operations: Enabled
- Storage account key access: Enabled
- Minimum TLS version: Version 1.2
- Infrastructure encryption: Disabled

### Storage access tiers

The default storage access tier for tiered data in Azure is the *hot* tier. The hot tier is ideal for frequently accessed data in the capacity tier.

If you don't plan to access the inactive data in the capacity tier, you can choose the *cool* storage tier, where the inactive data is retained for a minimum of 30 days. You can also opt for the *cold* tier, where the inactive data is stored for a minimum of 90 days. Based on your storage requirements and cost considerations, you can select the tier that best suits your needs. When you change the storage tier to *cool* or *cold*, the inactive capacity tier data moves directly to the cool or cold storage tier. The cool and cold tiers offer lower storage costs compared to the hot tier, but they come with higher access costs, so take that into consideration

before you change the storage tier. Refer to [Microsoft Azure documentation: Learn more about Azure Blob storage access tiers](#).

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

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

## Data tiering in Google Cloud

When you enable data tiering in Google Cloud, Cloud Volumes ONTAP uses persistent disks as a performance tier for hot data and a Google Cloud Storage bucket as a capacity tier for inactive data.

### Performance tier

The performance tier can be either SSD persistent disks, balanced persistent disks, or standard persistent disks.

### Capacity tier

A Cloud Volumes ONTAP system tiers inactive data to a single Google Cloud Storage bucket.

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

When BlueXP creates the bucket, it uses the following default settings:

- Location type: Region
- Storage class: Standard
- Public access: Subject to object ACLs
- Access control: Fine-grained
- Protection: None
- Data encryption: Google-managed key

### Storage classes

The default storage class for tiered data is the *Standard Storage* class. If the data is infrequently accessed, you can reduce your storage costs by changing to *Nearline Storage* or *Coldline Storage*. When you change the storage class, subsequent inactive data moves directly to the class that you selected.



Any existing inactive data will maintain the default storage class when you change the storage class. To change the storage class for existing inactive data, you must perform the designation manually.

The access costs are higher if you do access the data, so take that into consideration before you change the storage class. To learn more, refer to [Google Cloud documentation: Storage classes](#).

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

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

## Data tiering and capacity limits

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

## Volume tiering policies

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

Some tiering policies have an associated minimum cooling period, which sets the time that user data in a volume must remain inactive for the data to be considered "cold" and moved to the capacity tier. The cooling period starts when data is written to the aggregate.



You can change the minimum cooling period and default aggregate threshold of 50% (more on that below). [Learn how to change the cooling period](#) and [learn how to change the threshold](#).

BlueXP enables you to choose from the following volume tiering policies when you create or modify a volume:

### Snapshot Only

After an aggregate has reached 50% capacity, Cloud Volumes ONTAP tiers cold user data of Snapshot copies that are not associated with the active file system to the capacity tier. The cooling period is approximately 2 days.

If read, cold data blocks on the capacity tier become hot and are moved to the performance tier.

### All

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

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

### Auto

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

This policy is supported starting with Cloud Volumes ONTAP 9.4.

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

### None

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

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

## Turning off Cloud Volumes ONTAP impacts the cooling period

Data blocks are cooled by cooling scans. During this process, blocks that haven't been used have their block temperature moved (cooled) to the next lower value. The default cooling time depends on the volume tiering policy:

- Auto: 31 days
- Snapshot Only: 2 days

Cloud Volumes ONTAP must be running for the cooling scan to work. If Cloud Volumes ONTAP is turned off, cooling will stop, as well. As a result, you can experience longer cooling times.



When Cloud Volumes ONTAP is turned off, the temperature of each block is preserved until you restart the system. For example, if the temperature of a block is 5 when you turn the system off, the temp is still 5 when you turn the system back on.

## Setting up data tiering

For instructions and a list of supported configurations, refer to [Tier inactive data to low-cost object storage](#).

## Cloud Volumes ONTAP storage management

BlueXP provides simplified and advanced management of Cloud Volumes ONTAP storage.



All disks and aggregates must be created and deleted directly from BlueXP. You should not perform these actions from another management tool. Doing so can impact system stability, hamper the ability to add disks in the future, and potentially generate redundant cloud provider fees.

## Storage provisioning

BlueXP makes storage provisioning for Cloud Volumes ONTAP easy by purchasing disks and managing aggregates for you. You simply need to create volumes. You can use an advanced allocation option to provision aggregates yourself, if desired.

### Simplified provisioning

Aggregates provide cloud storage to volumes. BlueXP creates aggregates for you when you launch an instance, and when you provision additional volumes.

When you create a volume, BlueXP does one of three things:

- It places the volume on an existing aggregate that has sufficient free space.
- It places the volume on an existing aggregate by purchasing more disks for that aggregate.

In the case of an aggregate in AWS that supports Elastic Volumes, BlueXP also increases the size of the disks in a RAID group. [Learn more about support for Elastic Volumes](#).

- It purchases disks for a new aggregate and places the volume on that aggregate.

BlueXP determines where to place a new volume by looking at several factors: an aggregate's maximum size, whether thin provisioning is enabled, and free space thresholds for aggregates.

## Disk size selection for aggregates in AWS

When BlueXP creates new aggregates for Cloud Volumes ONTAP in AWS, it gradually increases the disk size in an aggregate, as the number of aggregates in the system increases. BlueXP does this to ensure that you can utilize the system's maximum capacity before it reaches the maximum number of data disks allowed by AWS.

For example, BlueXP might choose the following disk sizes:

Aggregate number	Disk size	Max aggregate capacity
1	500 GiB	3 TiB
4	1 TiB	6 TiB
6	2 TiB	12 TiB



This behavior does not apply to aggregates that support the Amazon EBS Elastic Volumes feature. Aggregates that have Elastic Volumes enabled are comprised of one or two RAID groups. Each RAID group has four identical disks that have the same capacity. [Learn more about support for Elastic Volumes.](#)

You can choose the disk size yourself by using the advanced allocation option.

### Advanced allocation

Rather than let BlueXP manage aggregates for you, you can do it yourself. [From the Advanced allocation page](#), you can create new aggregates that include a specific number of disks, add disks to an existing aggregate, and create volumes in specific aggregates.

### Capacity management

The BlueXP Organization or Account admin can choose whether BlueXP notifies you of storage capacity decisions or whether BlueXP automatically manages capacity requirements for you.

This behavior is determined by the *Capacity Management Mode* on a Connector. The Capacity Management Mode affects all Cloud Volumes ONTAP systems managed by that Connector. If you have another Connector, it can be configured differently.

### Automatic capacity management

The Capacity Management Mode is set to automatic by default. In this mode, BlueXP checks the free space ratio every 15 minutes to determine if the free space ratio falls below the specified threshold. If more capacity is needed, BlueXP automatically initiates purchase of new disks, deletes unused collections of disks (aggregates), moves volumes between aggregates as required, and attempts to prevent disk failure.

The following examples illustrate how this mode works:

- If an aggregate reaches the capacity threshold and it has room for more disks, BlueXP automatically purchases new disks for that aggregate so volumes can continue to grow.

In the case of an aggregate in AWS that supports Elastic Volumes, BlueXP also increases the size of the disks in a RAID group. [Learn more about support for Elastic Volumes.](#)

- If an aggregate reaches the capacity threshold and it can't support any additional disks, BlueXP automatically moves a volume from that aggregate to an aggregate with available capacity or to a new aggregate.

If BlueXP creates a new aggregate for the volume, it chooses a disk size that accommodates the size of that volume.

Note that free space is now available on the original aggregate. Existing volumes or new volumes can use that space. The space can't be returned to the cloud provider in this scenario.

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

## Management of LUNs with automatic capacity management

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

### Manual capacity management

If the BlueXP Organization or Account admin set the Capacity Management Mode to manual, BlueXP displays Action Required messages when capacity decisions must be made. The same examples described in the automatic mode apply to the manual mode, but it is up to you to accept the actions.

#### Learn more

[Learn how to modify the capacity management mode.](#)

## Write speed

BlueXP enables you to choose normal or high write speed for most Cloud Volumes ONTAP configurations. Before you choose a write speed, you should understand the differences between the normal and high settings and risks and recommendations when using high write speed.

### Normal write speed

When you choose normal write speed, data is written directly to disk. When data is written directly to disk, reduces the likelihood of data loss in the event of an unplanned system outage, or a cascading failure involving an unplanned system outage (HA pairs only).

Normal write speed is the default option.

### High write speed

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

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



## When to use high write speed

High write speed is a good choice if fast write performance is required for your workload and you can withstand the risk of data loss in the event of an unplanned system outage, or a cascading failure involving an unplanned system outage (HA pairs only).

## Recommendations when using high write speed

If you enable high write speed, you should ensure write protection at the application layer, or that the applications can tolerate data loss, if it occurs.

## High write speed with an HA pair in AWS

If you plan to enable high write speed on an HA pair in AWS, you should understand the difference in protection levels between a multiple Availability Zone (AZ) deployment and a single AZ deployment. Deploying an HA pair across multiple AZs provides more resiliency and can help to mitigate the chance of data loss.

[Learn more about HA pairs in AWS.](#)

## Configurations that support high write speed

Not all Cloud Volumes ONTAP configurations support high write speed. Those configurations use normal write speed by default.

## AWS

If you use a single node system, Cloud Volumes ONTAP supports high write speed with all instance types.

Starting with the 9.8 release, Cloud Volumes ONTAP supports high write speed with HA pairs when using almost all supported EC2 instance types, except for m5.xlarge and r5.xlarge.

[Learn more about the Amazon EC2 instances that Cloud Volumes ONTAP supports.](#)

## Azure

If you use a single node system, Cloud Volumes ONTAP supports high write speed with all VM types.

If you use an HA pair, Cloud Volumes ONTAP supports high write speed with several VM types, starting with the 9.8 release. Go to the [Cloud Volumes ONTAP Release Notes](#) to view the VM types that support high write speed.

## Google Cloud

If you use a single node system, Cloud Volumes ONTAP supports high write speed with all machine types.

If you use an HA pair, Cloud Volumes ONTAP supports high write speed with several VM types, starting with the 9.13.0 release. Go to the [Cloud Volumes ONTAP Release Notes](#) to view the VM types that support high write speed.

[Learn more about the Google Cloud machine types that Cloud Volumes ONTAP supports.](#)

## How to select a write speed

You can choose a write speed when you create a new working environment and you can [change the write speed for an existing system](#).

## What to expect if data loss occurs

If data loss occurs due to high write speed, the Event Management System (EMS) reports the following two events:

- Cloud Volumes ONTAP 9.12.1 or later

```
NOTICE nv.data.loss.possible: An unexpected shutdown occurred while in high write speed mode, which possibly caused a loss of data.
```

- Cloud Volumes ONTAP 9.11.0 to 9.11.1

```
DEBUG nv.check.failed: NVRAM check failed with error "NVRAM disabled due to dirty shutdown with High Write Speed mode"
```

```
ERROR wafl.root.content.changed: Contents of the root volume '' might have changed. Verify that all recent configuration changes are still in effect..
```

- Cloud Volumes ONTAP 9.8 to 9.10.1

```
DEBUG nv.check.failed: NVRAM check failed with error "NVRAM disabled due to dirty shutdown"
```

```
ERROR wafl.root.content.changed: Contents of the root volume '' might have changed. Verify that all recent configuration changes are still in effect.
```

When this happens, Cloud Volumes ONTAP should be able to boot up and continue to serve data without user intervention.

## How to stop data access if data loss occurs

If you are concerned about data loss, want the applications to stop running upon data loss, and the data access to be resumed after the data loss issue is properly addressed, you can use the NVFAIL option from the CLI to achieve that goal.

### To enable the NVFAIL option

```
vol modify -volume <vol-name> -nvfail on
```

### To check NVFAIL settings

```
vol show -volume <vol-name> -fields nvfail
```

## To disable the NVFAIL option

```
vol modify -volume <vol-name> -nvfail off
```

When data loss occurs, an NFS or iSCSI volume with NVFAIL enabled should stop serving data (there's no impact to CIFS which is a stateless protocol). For more details, refer to [How NVFAIL impacts access to NFS volumes or LUNs](#).

## To check the NVFAIL state

```
vol show -fields in-nvfailed-state
```

After the data loss issue is properly addressed, you can clear the NVFAIL state and the volume will be available for data access.

## To clear the NVFAIL state

```
vol modify -volume <vol-name> -in-nvfailed-state false
```

## Flash Cache

Some Cloud Volumes ONTAP configurations include local NVMe storage, which Cloud Volumes ONTAP uses as *Flash Cache* for better performance.

### What's Flash Cache?

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

### Supported configurations

Flash Cache is supported with specific Cloud Volumes ONTAP configurations. View supported configurations in the [Cloud Volumes ONTAP Release Notes](#)

### Limitations

- When configuring Flash Cache for Cloud Volumes ONTAP 9.12.0 or earlier in AWS, compression must be disabled on all volumes to take advantage of the Flash Cache performance improvements. When you deploy or upgrade to Cloud Volumes ONTAP 9.12.1 or later, you don't need to disable compression.

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

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

## Learn about WORM storage on Cloud Volumes ONTAP

You can activate write once, read many (WORM) storage on a Cloud Volumes ONTAP system to retain files in unmodified form for a specified retention period. Cloud WORM storage is powered by SnapLock technology, which means WORM files are protected at the file level.

The WORM feature is available for use with bring your own license (BYOL) and marketplace subscriptions for your licenses at no additional cost. Contact your NetApp sales representative to add WORM to your current license.

## How WORM storage works

Once a file has been committed to WORM storage, it can't be modified, even after the retention period has expired. A tamper-proof clock determines when the retention period for a WORM file has elapsed.

After the retention period has elapsed, you are responsible for deleting any files that you no longer need.

## Activating WORM storage

How you activate WORM storage depends on the Cloud Volumes ONTAP version that you're using.

### Version 9.10.1 and later

Beginning with Cloud Volumes ONTAP 9.10.1, you have the option to enable or disable WORM at the volume level.

When you create a new Cloud Volumes ONTAP working environment, you're prompted to enable or disable WORM storage:

- If you enable WORM storage when creating a working environment, every volume that you create from BlueXP has WORM enabled. But you can use ONTAP System Manager or the ONTAP CLI to create volumes that have WORM disabled.
- If you disable WORM storage when creating a working environment, every volume that you create from BlueXP, ONTAP System Manager, or the ONTAP CLI has WORM disabled.

### Version 9.10.0 and earlier

You can activate WORM storage on a Cloud Volumes ONTAP system when you create a new working environment. Every volume that you create from BlueXP has WORM enabled. You can't disable WORM storage on individual volumes.

## Committing files to WORM

You can use an application to commit files to WORM over NFS or CIFS, or use the ONTAP CLI to autocommit files to WORM automatically. You can also use a WORM appendable file to retain data that is written incrementally, like log information.

After you activate WORM storage on a Cloud Volumes ONTAP system, you must use the ONTAP CLI for all management of WORM storage. For instructions, refer to the [ONTAP documentation on SnapLock](#).

## Enabling WORM on a Cloud Volumes ONTAP working environment

You can enable WORM storage when creating a Cloud Volumes ONTAP working environment on BlueXP. You can also enable WORM on a working environment if WORM is not enabled on it during creation. After you enable it, you cannot disable WORM.

### About this task

- WORM is supported on ONTAP 9.10.1 and later.
- WORM with backup is supported on ONTAP 9.11.1 and later.

### Steps

1. On the Canvas page, double-click the name of the working environment on which you want to enable WORM.

2. On the Overview tab, click the Features panel and then click the pencil icon next to **WORM**.

If WORM is already enabled on the system, the pencil icon is disabled.

3. On the **WORM** page, set the retention period for the cluster Compliance Clock.

For more information, refer to the [ONTAP documentation: Initialize the Compliance Clock](#).

4. Click **Set**.

### After you finish

You can verify the status of **WORM** on the Features panel.

After WORM is enabled, the SnapLock license is automatically installed on the cluster. You can view the SnapLock license on ONTAP System Manager.

### Deleting WORM files

You can delete WORM files during the retention period using the privileged delete feature.

For instructions, refer to the [ONTAP documentation](#).

### WORM and data tiering

When you create a new Cloud Volumes ONTAP 9.8 system or later, you can enable both data tiering and WORM storage together. Enabling data tiering with WORM storage allows you to tier the data to an object store in the cloud.

You should understand the following about enabling both data tiering and WORM storage:

- Data that is tiered to object storage doesn't include the ONTAP WORM functionality. To ensure end-to-end WORM capability, you'll need to set up the bucket permissions correctly.
- The data that is tiered to object storage doesn't carry the WORM functionality, which means technically anyone with full access to buckets and containers can go and delete the objects tiered by ONTAP.
- Reverting or downgrading to Cloud Volumes ONTAP 9.8 is blocked after enabling WORM and tiering.

### Limitations

- WORM storage in Cloud Volumes ONTAP operates under a "trusted storage administrator" model. While WORM files are protected from alteration or modification, volumes can be deleted by a cluster administrator even if those volumes contain unexpired WORM data.
- In addition to the trusted storage administrator model, WORM storage in Cloud Volumes ONTAP also implicitly operates under a "trusted cloud administrator" model. A cloud administrator could delete WORM data before its expiration date by removing or editing cloud storage directly from the cloud provider.

### Related link

- [Create tamperproof Snapshot copies for WORM storage](#)

## High-availability pairs

### Learn about Cloud Volumes ONTAP HA pairs in AWS

A Cloud Volumes ONTAP high-availability (HA) configuration provides nondisruptive

operations and fault tolerance. In AWS, data is synchronously mirrored between the two nodes.

## HA components

In AWS, Cloud Volumes ONTAP HA configurations include the following components:

- Two Cloud Volumes ONTAP nodes whose data is synchronously mirrored between each other.
- A mediator instance that provides a communication channel between the nodes to assist in storage takeover and giveback processes.

### Mediator

Here are some key details about the mediator instance in AWS:

#### Instance type

t3-micro

#### Disks

Two st1 disks of 8 GiB and 4 GiB

#### Operating system

Debian 11



For Cloud Volumes ONTAP 9.10.0 and earlier, Debian 10 was installed on the mediator.

### Upgrades

When you upgrade Cloud Volumes ONTAP, BlueXP also updates the mediator instance as needed.

### Access to the instance

When you create a Cloud Volumes ONTAP HA pair from BlueXP, you're prompted to provide a key pair for the mediator instance. You can use that key pair for SSH access using the `admin` user.

### Third-party agents

Third-party agents or VM extensions are not supported on the mediator instance.

### Storage takeover and giveback

If a node goes down, the other node can serve data for its partner to provide continued data service. Clients can access the same data from the partner node because the data was synchronously mirrored to the partner.

After the node reboots, the partner must resync data before it can return the storage. The time that it takes to resync data depends on how much data was changed while the node was down.

Storage takeover, resync, and giveback are all automatic by default. No user action is required.

### RPO and RTO

An HA configuration maintains high-availability of your data as follows:

- The recovery point objective (RPO) is 0 seconds.  
Your data is transactionally consistent with no data loss.

- The recovery time objective (RTO) is 120 seconds.  
In the event of an outage, data should be available in 120 seconds or less.

## HA deployment models

You can ensure the high availability of your data by deploying an HA configuration across multiple availability zones (AZs) or in a single availability zone (AZ). You should review more details about each configuration to choose which best fits your needs.

### Multiple availability zones

Deploying an HA configuration in multiple availability zones (AZs) ensures high availability of your data if a failure occurs with an AZ or an instance that runs a Cloud Volumes ONTAP node. You should understand how NAS IP addresses impact data access and storage failover.

### NFS and CIFS data access

When an HA configuration is spread across multiple Availability Zones, *floating IP addresses* enable NAS client access. The floating IP addresses, which must be outside of the CIDR blocks for all VPCs in the region, can migrate between nodes when failures occur. They aren't natively accessible to clients that are outside of the VPC, unless you [set up an AWS transit gateway](#).

If you can't set up a transit gateway, private IP addresses are available for NAS clients that are outside the VPC. However, these IP addresses are static—they can't failover between nodes.

You should review requirements for floating IP addresses and route tables before you deploy an HA configuration across multiple availability zones. You must specify the floating IP addresses when you deploy the configuration. The private IP addresses are automatically created by BlueXP.

For more information, refer to [AWS networking requirements for Cloud Volumes ONTAP HA in multiple AZs](#).

### iSCSI data access

Cross-VPC data communication is not an issue since iSCSI does not use floating IP addresses.

### Takeover and giveback for iSCSI


For iSCSI, Cloud Volumes ONTAP uses multipath I/O (MPIO) and Asymmetric Logical Unit Access (ALUA) to manage path failover between the active-optimized and non-optimized paths.



For information about which specific host configurations support ALUA, refer to the [NetApp Interoperability Matrix Tool](#) and the [SAN hosts and cloud clients guide](#) for your host operating system.

### Takeover and giveback for NAS

When takeover occurs in a NAS configuration using floating IPs, the node's floating IP address that clients use to access data moves to the other node. The following image depicts storage takeover in a NAS configuration using floating IPs. If node 2 goes down, the floating IP address for node 2 moves to node 1.



NAS data IPs used for external VPC access cannot migrate between nodes if failures occur. If a node goes offline, you must manually remount volumes to clients outside the VPC by using the IP address on the other node.

After the failed node comes back online, remount clients to volumes using the original IP address. This step is needed to avoid transferring unnecessary data between two HA nodes, which can cause significant performance and stability impact.

You can easily identify the correct IP address from BlueXP by selecting the volume and clicking **Mount Command**.

### Single availability zone

Deploying an HA configuration in a single availability zone (AZ) can ensure high availability of your data if an instance that runs a Cloud Volumes ONTAP node fails. All data is natively accessible from outside of the VPC.



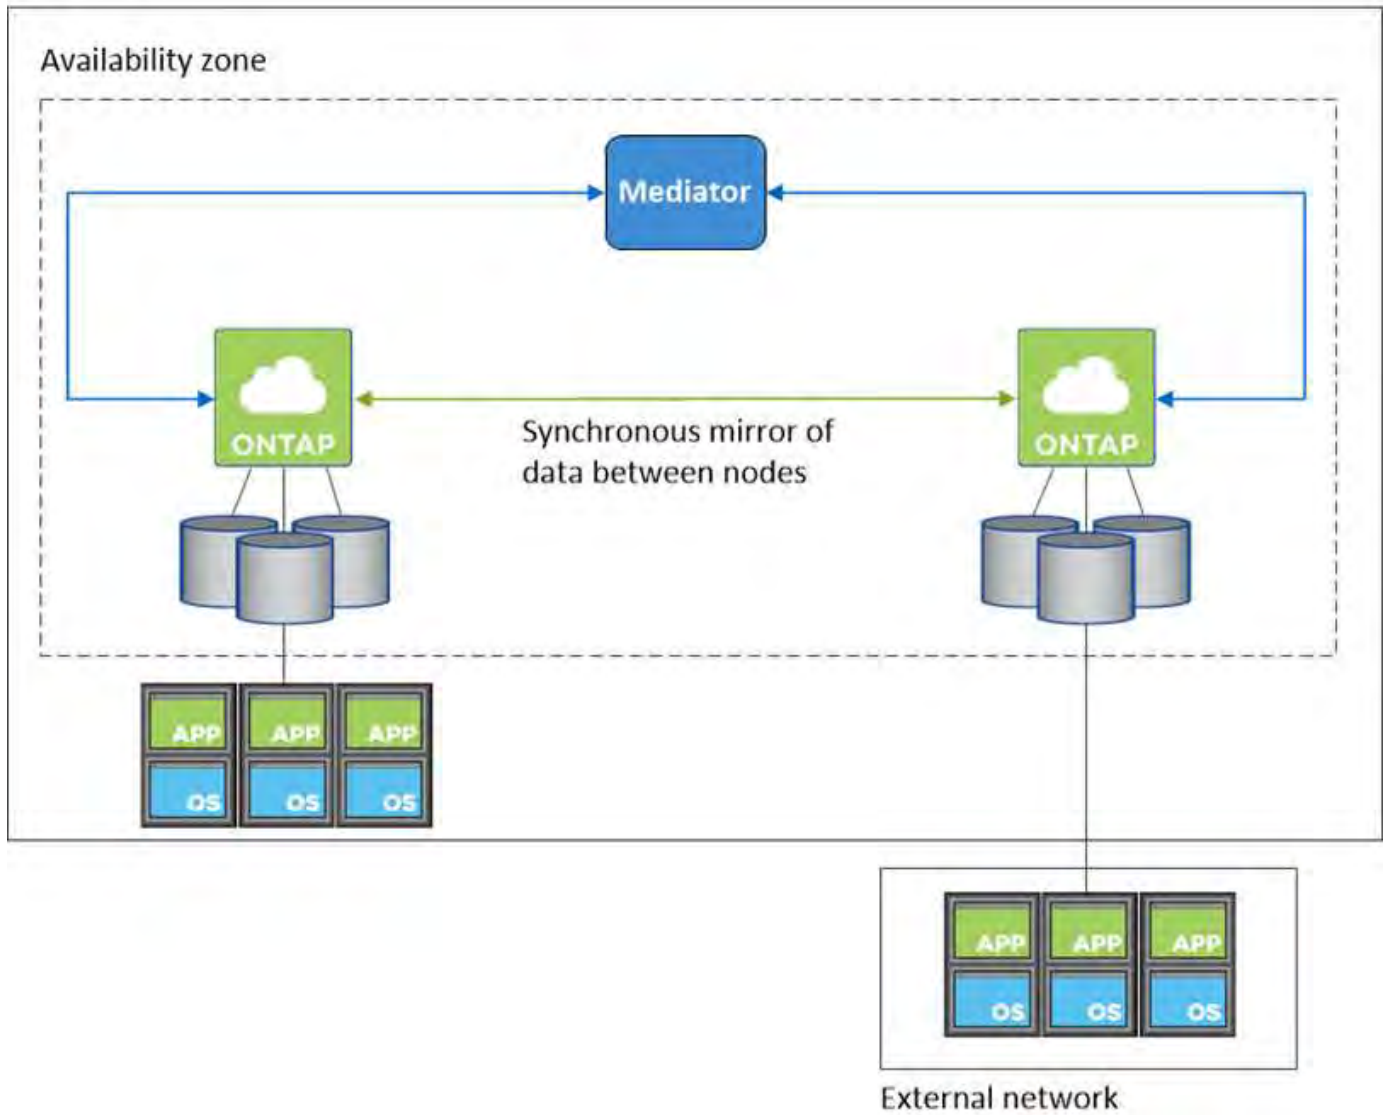
BlueXP creates an [AWS Documentation: AWS spread placement group](#) and launches the two HA nodes in that placement group. The placement group reduces the risk of simultaneous failures by spreading the instances across distinct underlying hardware. This feature improves redundancy from a compute perspective and not from disk failure perspective.

### Data access

Because this configuration is in a single AZ, it does not require floating IP addresses. You can use the same IP address for data access from within the VPC and from outside the VPC.

The following image shows an HA configuration in a single AZ. Data is accessible from within the VPC and from outside the VPC.





### Takeover and giveback

For iSCSI, Cloud Volumes ONTAP uses multipath I/O (MPIO) and Asymmetric Logical Unit Access (ALUA) to manage path failover between the active-optimized and non-optimized paths.



For information about which specific host configurations support ALUA, refer to the [NetApp Interoperability Matrix Tool](#) and the [SAN hosts and cloud clients guide](#) for your host operating system.

For NAS configurations, the data IP addresses can migrate between HA nodes if failures occur. This ensures client access to storage.

### AWS Local Zones

AWS Local Zones are an infrastructure deployment where storage, compute, database, and other select AWS services are located close to large cities and industry areas. With AWS Local Zones, you can bring AWS services closer to you which improves latency for your workloads and maintain databases locally. On Cloud Volumes ONTAP,

You can deploy a single AZ or multiple AZ configuration in AWS Local Zones.



AWS Local Zones are supported when using BlueXP in standard and private modes. At this time, AWS Local Zones are not supported when using BlueXP in restricted mode.

### Example AWS Local Zone configurations

Cloud Volumes ONTAP in AWS supports only high availability (HA) mode in a single availability zone. Single node deployments are not supported.

Cloud Volumes ONTAP does not support data tiering, cloud tiering, and unqualified instances in AWS Local Zones.

The following are example configurations:

- Single availability zone: Both cluster nodes and the mediator are in the same Local Zone.
- Multiple availability zones  
In multiple availability zone configurations, there are three instances, two nodes and one mediator. One instance out of the three instances must be in a separate zone. You can choose how you set this up.

Here are three example configurations:

- Each cluster node is in a different Local Zone and the mediator in a public availability zone.
- One cluster node in a Local Zone, the mediator in a Local Zone, and the second cluster node is in an availability zone.
- Each cluster node and the mediator are in separate Local Zones.

### Supported disk and instance types

The only supported disk type is GP2. The following EC2 instance type families with sizes xlarge to 4xlarge are currently supported:

- M5
- C5
- C5d
- R5
- R5d



Cloud Volumes ONTAP supports only these configurations. Selecting unsupported disk types or unqualified instances in AWS Local Zone configuration might result in deployment failure. Data tiering to AWS S3 is not available in AWS Local Zones due to lack of connectivity.

Refer to AWS documentation for the latest and complete details of the [AWS Documentation: EC2 instance types in Local Zones](#).

### How storage works in an HA pair

Unlike an ONTAP cluster, storage in a Cloud Volumes ONTAP HA pair is not shared between nodes. Instead, data is synchronously mirrored between the nodes so that the data is available in the event of failure.

## Storage allocation

When you create a new volume and additional disks are required, BlueXP allocates the same number of disks to both nodes, creates a mirrored aggregate, and then creates the new volume. For example, if two disks are required for the volume, BlueXP allocates two disks per node for a total of four disks.

## Storage configurations

You can use an HA pair as an active-active configuration, in which both nodes serve data to clients, or as an active-passive configuration, in which the passive node responds to data requests only if it has taken over storage for the active node.



You can set up an active-active configuration only when using BlueXP in the Storage System View.

## Performance expectations

A Cloud Volumes ONTAP HA configuration synchronously replicates data between nodes, which consumes network bandwidth. As a result, you can expect the following performance in comparison to a single-node Cloud Volumes ONTAP configuration:

- For HA configurations that serve data from only one node, read performance is comparable to the read performance of a single-node configuration, whereas write performance is lower.
- For HA configurations that serve data from both nodes, read performance is higher than the read performance of a single-node configuration, and write performance is the same or higher.

For more details about Cloud Volumes ONTAP performance, refer to [Performance](#).

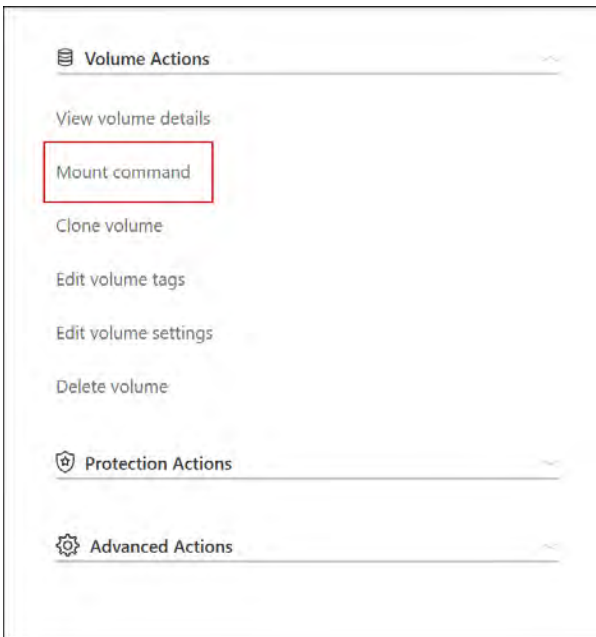
## Client access to storage

Clients should access NFS and CIFS volumes by using the data IP address of the node on which the volume resides. If NAS clients access a volume by using the IP address of the partner node, traffic goes between both nodes, which reduces performance.



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

You can easily identify the correct IP address through the *Mount Command* option under the manage volumes panel in BlueXP.



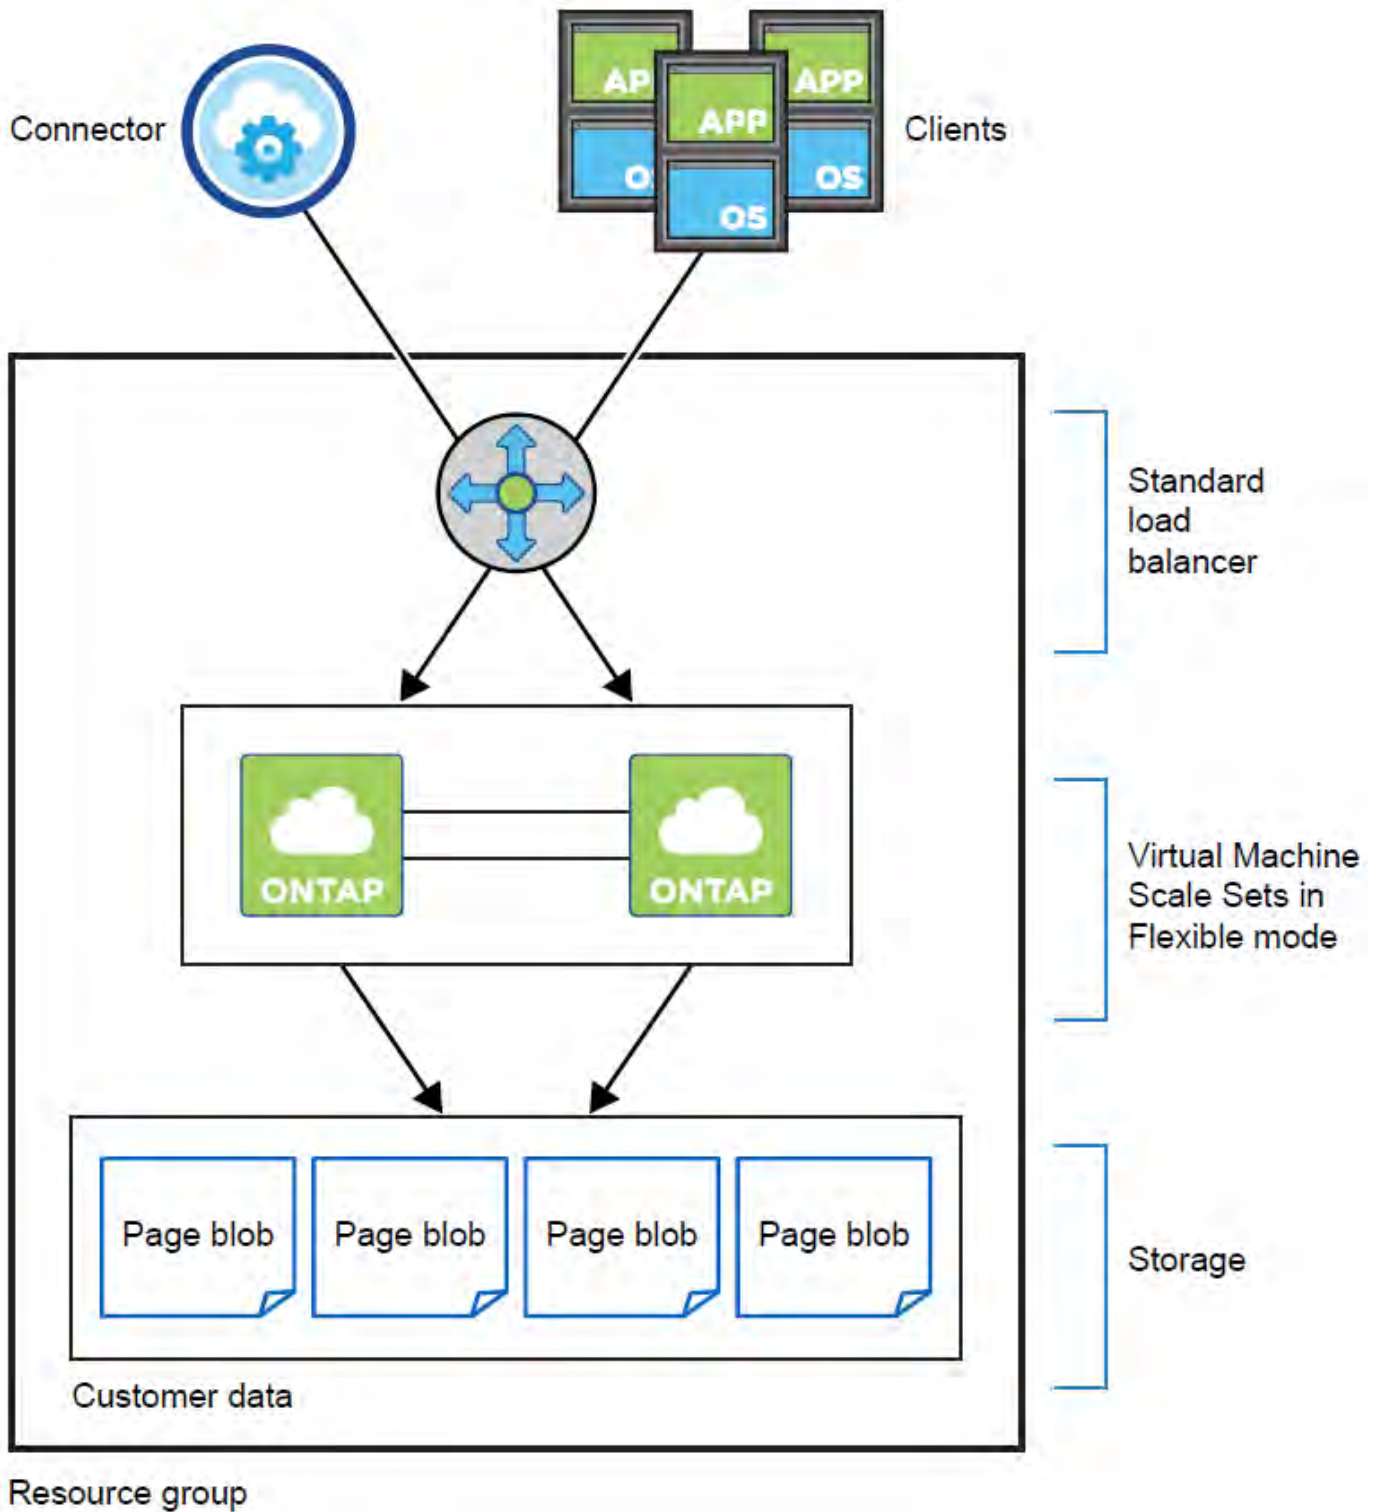
## Learn about Cloud Volumes ONTAP HA pairs in Azure

A Cloud Volumes ONTAP high-availability (HA) pair provides enterprise reliability and continuous operations in case of failures in your cloud environment. In Azure, storage is shared between the two nodes.

### HA components

#### HA single availability zone configuration with page blobs

A Cloud Volumes ONTAP HA page blob configuration in Azure includes the following components:



Note the following about the Azure components that BlueXP deploys for you:

#### Azure Standard Load Balancer

The load balancer manages incoming traffic to the Cloud Volumes ONTAP HA pair.

#### VMs in single availability zones

Beginning with Cloud Volumes ONTAP 9.15.1, you can create and manage heterogeneous virtual machines (VMs) in a single availability zone (AZ). You can deploy high-availability (HA) nodes in separate fault domains within the same AZ, guaranteeing optimal availability. To learn more about the flexible

orchestration mode that enables this capability, refer to [Microsoft Azure documentation: Virtual Machine Scale Sets](#).

## Disks

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

## Storage accounts

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

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

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

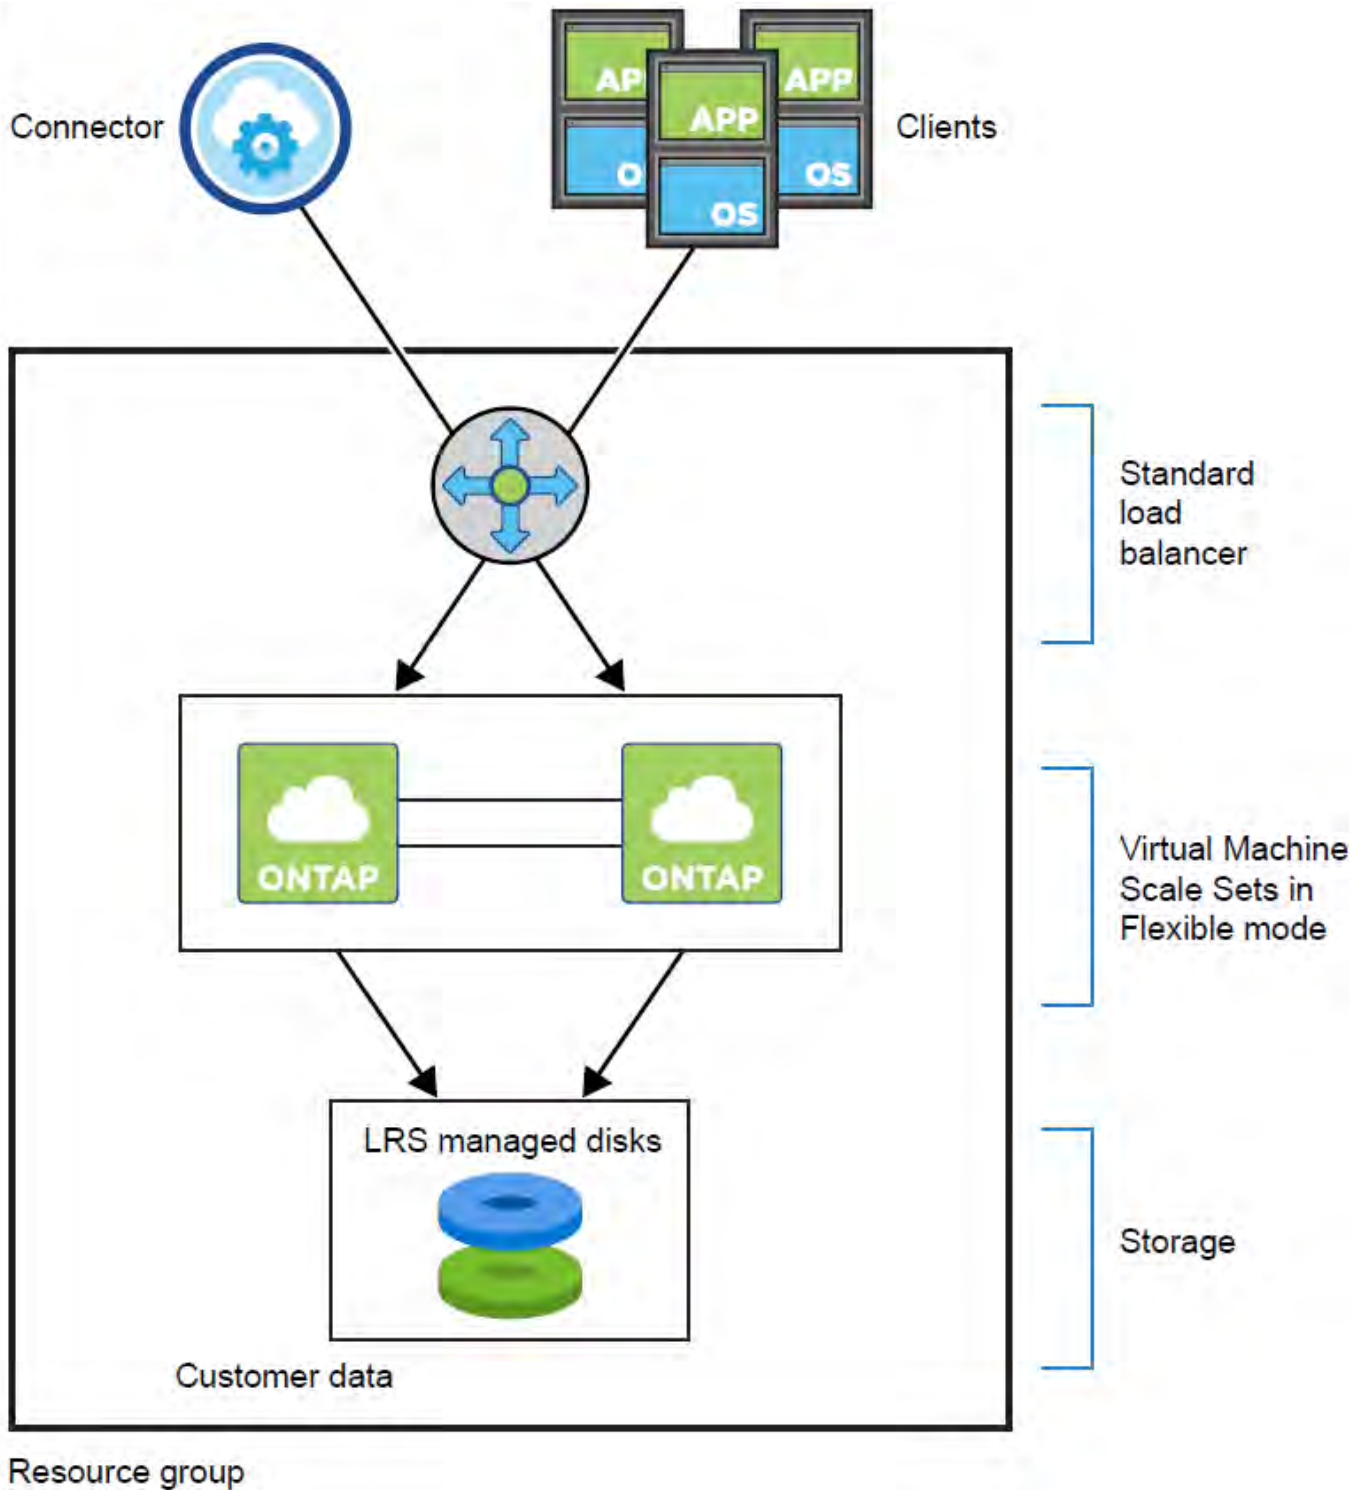


Starting with Cloud Volumes ONTAP 9.15.0P1, Azure page blobs are no longer supported for new high-availability pair deployments. If you currently use Azure page blobs in existing high-availability pair deployments, you can migrate to newer VM instance types in the Edsv4-series VMs and Edsv5-series VMs.

[Learn more about supported configurations in Azure.](#)

## HA single availability zone configuration with shared managed disks

A Cloud Volumes ONTAP HA single availability zone configuration running on top of shared managed disk includes the following components:



Note the following about the Azure components that BlueXP deploys for you:

### Azure Standard Load Balancer

The load balancer manages incoming traffic to the Cloud Volumes ONTAP HA pair.

### VMs in single availability zones

Beginning with Cloud Volumes ONTAP 9.15.1, you can create and manage heterogeneous virtual machines (VMs) in a single availability zone (AZ). You can deploy high-availability (HA) nodes in separate fault domains within the same AZ, guaranteeing optimal availability. To learn more about the flexible orchestration mode that enables this capability, refer to [Microsoft Azure documentation: Virtual Machine](#)

## Scale Sets.

The zonal deployment uses Premium SSD v2 Managed Disks when the following conditions are fulfilled:

- The version of Cloud Volumes ONTAP is 9.15.1 or later.
- The selected region and zone support Premium SSD v2 Managed Disks. For information about the supported regions, refer to [Microsoft Azure website: Products available by region](#).
- The subscription is registered for the Microsoft [Microsoft.Compute/VMOrchestratorZonalMultiFD feature](#).

## Disks

Customer data resides on locally-redundant storage (LRS) managed disks. Each node has access to the other node's storage. Additional storage is also required for [boot, root, partner root, core, and NVRAM data](#).

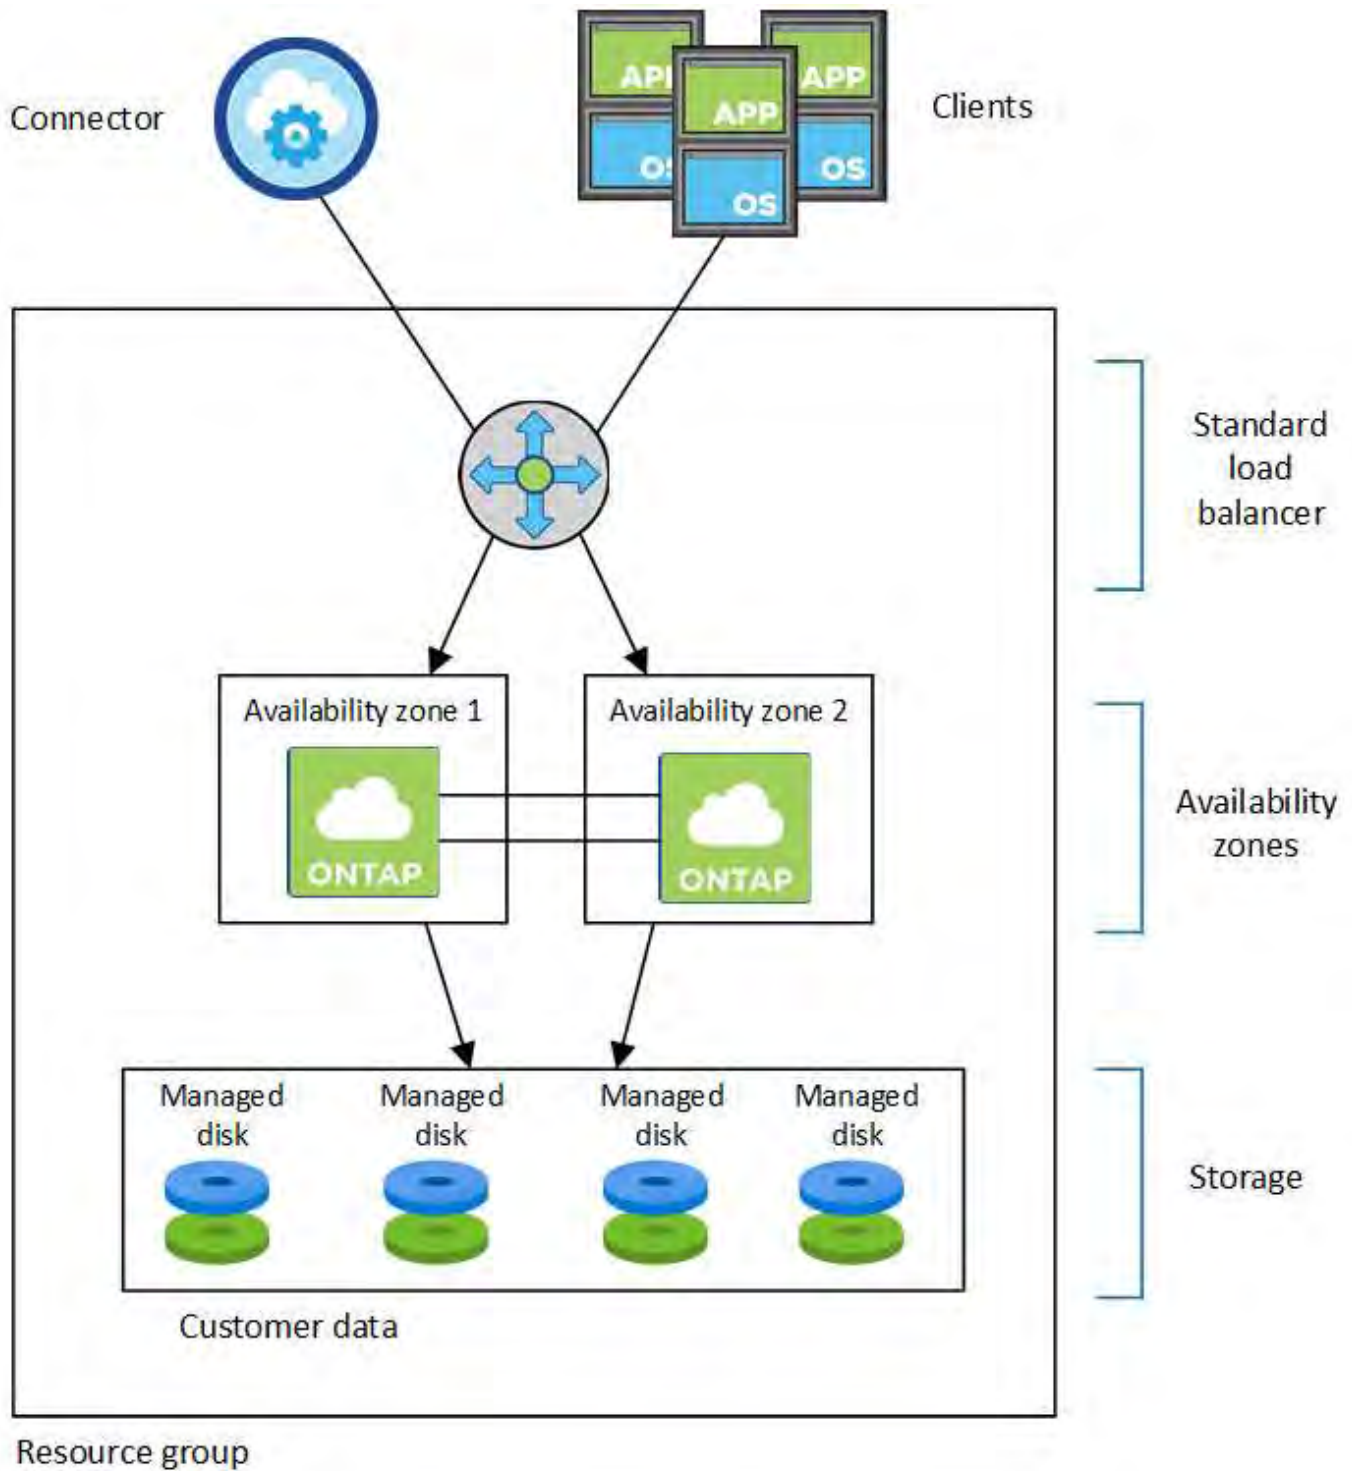
## Storage accounts

Storage accounts are used for managed disk based deployments to handle diagnostic logs and tiering to blob storage.

## HA multiple availability zone configuration

A Cloud Volumes ONTAP HA multiple availability zone configuration in Azure includes the following components:





Note the following about the Azure components that BlueXP deploys for you:

### Azure Standard Load Balancer

The load balancer manages incoming traffic to the Cloud Volumes ONTAP HA pair.

### Availability Zones

HA multiple availability zone configuration utilizes a deployment model where two Cloud Volumes ONTAP nodes are deployed into different availability zones, ensuring that the nodes are in different fault domains to provide redundancy and availability. To learn how Virtual Machine Scale Sets in Flexible orchestration mode can use availability zones in Azure, refer to [Microsoft Azure documentation: Create a Virtual Machine Scale](#)

[Set that uses Availability Zones.](#)

## Disks

Customer data resides on zone-redundant storage (ZRS) managed disks. Each node has access to the other node's storage. Additional storage is also required for [boot, root, partner root, and core data](#).

## Storage accounts

Storage accounts are used for managed disk based deployments to handle diagnostic logs and tiering to blob storage.

## RPO and RTO

An HA configuration maintains high availability of your data as follows:

- The recovery point objective (RPO) is 0 seconds.  
Your data is transactionally consistent with no data loss.
- The recovery time objective (RTO) is 120 seconds.  
In the event of an outage, data should be available in 120 seconds or less.

## Storage takeover and giveback

Similar to a physical ONTAP cluster, storage in an Azure HA pair is shared between nodes. Connections to the partner's storage allows each node to access the other's storage in the event of a *takeover*. Network path failover mechanisms ensure that clients and hosts continue to communicate with the surviving node. The partner *gives back* storage when the node is brought back on line.

For NAS configurations, data IP addresses automatically migrate between HA nodes if failures occur.

For iSCSI, Cloud Volumes ONTAP uses multipath I/O (MPIO) and Asymmetric Logical Unit Access (ALUA) to manage path failover between the active-optimized and non-optimized paths.



For information about which specific host configurations support ALUA, refer to the [NetApp Interoperability Matrix Tool](#) and the [SAN hosts and cloud clients guide](#) for your host operating system.

Storage takeover, resync, and giveback are all automatic by default. No user action is required.

## Storage configurations

You can use an HA pair as an active-active configuration, in which both nodes serve data to clients, or as an active-passive configuration, in which the passive node responds to data requests only if it has taken over the storage for the active node.

## Learn about Cloud Volumes ONTAP HA pairs in Google Cloud

A Cloud Volumes ONTAP high-availability (HA) configuration provides nondisruptive operations and fault tolerance. In Google Cloud, data is synchronously mirrored between the two nodes.

## HA components

Cloud Volumes ONTAP HA configurations in Google Cloud include the following components:

- Two Cloud Volumes ONTAP nodes whose data is synchronously mirrored between each other.
- A mediator instance that provides a communication channel between the nodes to assist in storage takeover and giveback processes.
- One zone or three zones (recommended).

If you choose three zones, the two nodes and mediator are in separate Google Cloud zones.

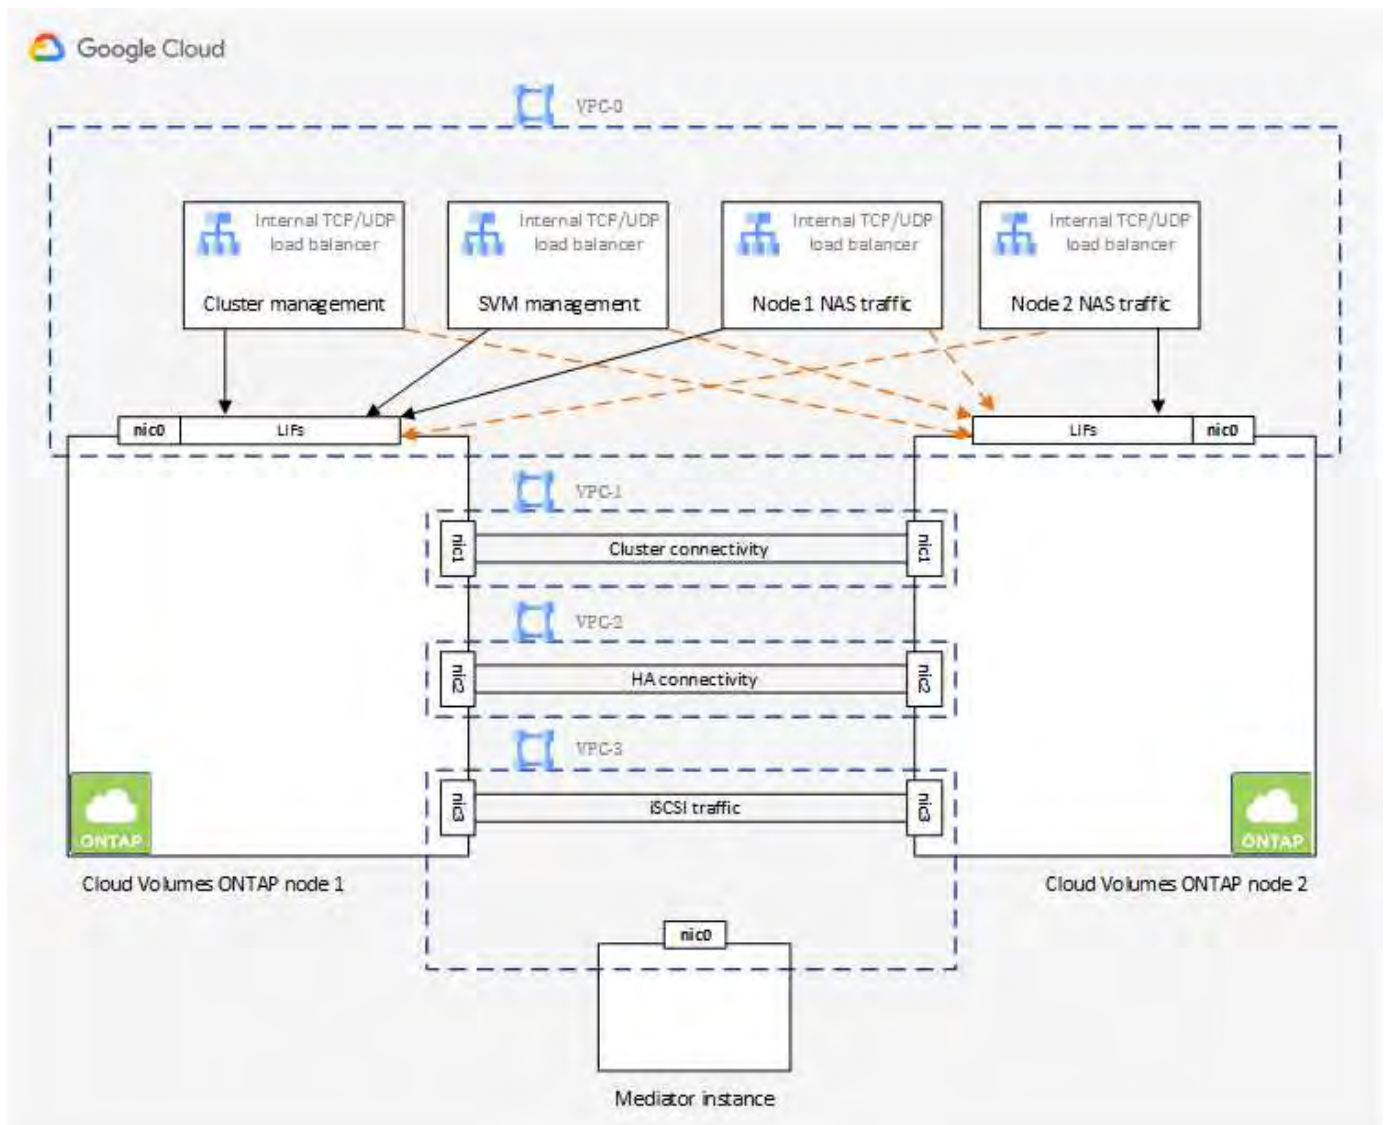
- Four Virtual Private Clouds (VPCs).

The configuration uses four VPCs because GCP requires that each network interface resides in a separate VPC network.

- Four Google Cloud internal load balancers (TCP/UDP) that manage incoming traffic to the Cloud Volumes ONTAP HA pair.

[Learn about networking requirements](#), including more details about load balancers, VPCs, internal IP addresses, subnets, and more.

The following conceptual image shows a Cloud Volumes ONTAP HA pair and its components:



## Mediator

Here are some key details about the mediator instance in Google Cloud:

### Instance type

e2-micro (an f1-micro instance was previously used)

### Disks

Two standard persistent disks that are 10 GiB each

### Operating system

Debian 11



For Cloud Volumes ONTAP 9.10.0 and earlier, Debian 10 was installed on the mediator.

### Upgrades

When you upgrade Cloud Volumes ONTAP, BlueXP also updates the mediator instance as needed.

### Access to the instance

For Debian, the default cloud user is `admin`. Google Cloud creates and adds a certificate for the `admin` user when SSH access is requested through the Google Cloud console or `gcloud` command line. You can specify `sudo` to gain root privileges.

### Third-party agents

Third-party agents or VM extensions are not supported on the mediator instance.

### Storage takeover and giveback

If a node goes down, the other node can serve data for its partner to provide continued data service. Clients can access the same data from the partner node because the data was synchronously mirrored to the partner.

After the node reboots, the partner must resync data before it can return the storage. The time that it takes to resync data depends on how much data was changed while the node was down.

Storage takeover, resync, and giveback are all automatic by default. No user action is required.

### RPO and RTO

An HA configuration maintains high availability of your data as follows:

- The recovery point objective (RPO) is 0 seconds.

Your data is transactionally consistent with no data loss.

- The recovery time objective (RTO) is 120 seconds.

In the event of an outage, data should be available in 120 seconds or less.

### HA deployment models

You can ensure the high availability of your data by deploying an HA configuration in multiple zones or in a single zone.

## Multiple zones (recommended)

Deploying an HA configuration across three zones ensures continuous data availability if a failure occurs within a zone. Note that write performance is slightly lower compared to using a single zone, but it's minimal.

## Single zone

When deployed in a single zone, a Cloud Volumes ONTAP HA configuration uses a spread placement policy. This policy ensures that an HA configuration is protected from a single point of failure within the zone, without having to use separate zones to achieve fault isolation.

This deployment model does lower your costs because there are no data egress charges between zones.

## How storage works in an HA pair

Unlike an ONTAP cluster, storage in a Cloud Volumes ONTAP HA pair in GCP is not shared between nodes. Instead, data is synchronously mirrored between the nodes so that the data is available in the event of failure.

### Storage allocation

When you create a new volume and additional disks are required, BlueXP allocates the same number of disks to both nodes, creates a mirrored aggregate, and then creates the new volume. For example, if two disks are required for the volume, BlueXP allocates two disks per node for a total of four disks.

### Storage configurations

You can use an HA pair as an active-active configuration, in which both nodes serve data to clients, or as an active-passive configuration, in which the passive node responds to data requests only if it has taken over storage for the active node.

### Performance expectations for an HA configuration

A Cloud Volumes ONTAP HA configuration synchronously replicates data between nodes, which consumes network bandwidth. As a result, you can expect the following performance in comparison to a single-node Cloud Volumes ONTAP configuration:

- For HA configurations that serve data from only one node, read performance is comparable to the read performance of a single-node configuration, whereas write performance is lower.
- For HA configurations that serve data from both nodes, read performance is higher than the read performance of a single-node configuration, and write performance is the same or higher.

For more details about Cloud Volumes ONTAP performance, refer to [Performance](#).

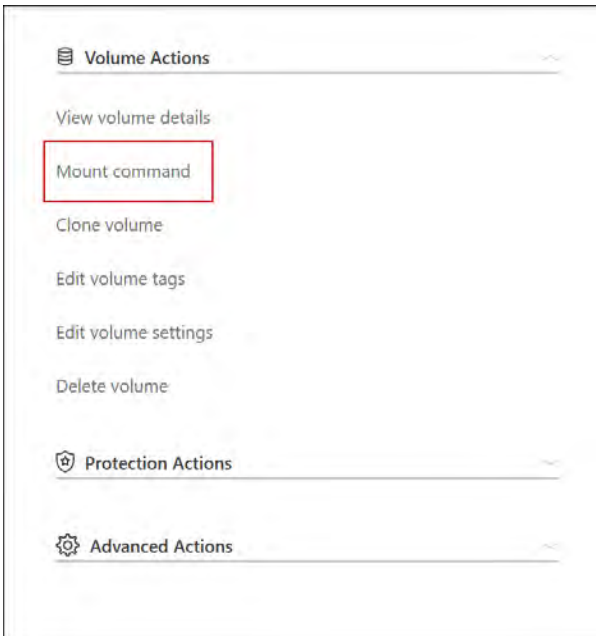
### Client access to storage

Clients should access NFS and CIFS volumes by using the data IP address of the node on which the volume resides. If NAS clients access a volume by using the IP address of the partner node, traffic goes between both nodes, which reduces performance.



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

You can easily identify the correct IP address through the *Mount Command* option under the manage volumes panel in BlueXP.



#### Related links

- [Learn about networking requirements](#)
- [Learn how to get started in GCP](#)

## Operations unavailable when a node in Cloud Volumes ONTAP HA pair is offline

When a node in an HA pair isn't available, the other node serves data for its partner to provide continued data service. This is called *storage takeover*. Several actions are unavailable until in storage giveback is complete.



When a node in an HA pair is unavailable, the state of the working environment in BlueXP is *Degraded*.

The following actions are unavailable from BlueXP storage takeover:

- Support registration
- License changes
- Instance or VM type changes
- Write speed changes
- CIFS setup
- Changing the location of configuration backups
- Setting the cluster password
- Managing disks and aggregates (advanced allocation)

These actions are available again after storage giveback completes and the state of the working environment changes back to normal.



# Learn about Cloud Volumes ONTAP data encryption and ransomware protection

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

## Encryption of data at rest

Cloud Volumes ONTAP supports the following encryption technologies:

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

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

## NetApp encryption solutions (NVE and NAE)

Cloud Volumes ONTAP supports [NetApp Volume Encryption \(NVE\)](#) and [NetApp Aggregate Encryption \(NAE\)](#). NVE and NAE are software-based solutions that enable (FIPS) 140-2-compliant data-at-rest encryption of volumes. Both NVE and NAE use AES 256-bit encryption.

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

Cloud Volumes ONTAP supports both NVE and NAE with external key management services (EKMs) provided by AWS, Azure, and Google Cloud, including third-party solutions, such as Fortanix. Unlike ONTAP, for Cloud Volumes ONTAP, encryption keys are generated at the cloud provider's side, not in ONTAP.

Cloud Volumes ONTAP uses the standard Key Management Interoperability Protocol (KMIP) services that ONTAP uses. For more information about the supported services, refer to the [Interoperability Matrix Tool](#).

If you use NVE, you have the option to use your cloud provider's key vault to protect ONTAP encryption keys:

- AWS Key Management Service (KMS)
- Azure Key Vault (AKV)
- Google Cloud Key Management Service

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

Setting up a supported key manager is the only required step. For set up instructions, refer to [Encrypt volumes with NetApp encryption solutions](#).

## AWS Key Management Service

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



You can't change the AWS data encryption method after you create a Cloud Volumes ONTAP system.

If you want to use this encryption option, then you must ensure that the AWS KMS is set up appropriately. For information, refer to [Setting up the AWS KMS](#).

## Azure Storage Service Encryption

Data is automatically encrypted on Cloud Volumes ONTAP in Azure using [Azure Storage Service Encryption](#) with a Microsoft-managed key.

You can use your own encryption keys if you prefer. [Learn how to set up Cloud Volumes ONTAP to use a customer-managed key in Azure](#).

## Google Cloud Platform default encryption

[Google Cloud Platform data-at-rest encryption](#) is enabled by default for Cloud Volumes ONTAP. No setup is required.

While Google Cloud Storage always encrypts your data before it's written to disk, you can use BlueXP APIs to create a Cloud Volumes ONTAP system that uses *customer-managed encryption keys*. These are keys that you generate and manage in GCP using the Cloud Key Management Service. [Learn more](#).

## ONTAP virus scanning

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

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

For information about the vendors, software, and versions supported by Vscan, refer to the [NetApp Interoperability Matrix](#).

For information about how to configure and manage the antivirus functionality on ONTAP systems, refer to the [ONTAP 9 Antivirus Configuration Guide](#).

## Ransomware protection

Ransomware attacks can cost a business time, resources, and reputation. BlueXP enables you to implement the NetApp solution for ransomware, which provides effective tools for visibility, detection, and remediation.

- BlueXP identifies volumes that are not protected by a Snapshot policy and enables you to activate the default Snapshot policy on those volumes.

Snapshot copies are read-only, which prevents ransomware corruption. They can also provide the granularity to create images of a single file copy or a complete disaster recovery solution.

- BlueXP also enables you to block common ransomware file extensions by enabling ONTAP's FPolicy



solution.

**Ransomware Protection**

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

**1 Enable Snapshot Copy Protection**

50 % Protection

1 Volumes without a Snapshot Policy

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

[Activate Snapshot Policy](#)

**2 Block Ransomware File Extensions**

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

[View Denied File Names](#)

[Activate FPolicy](#)

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

## Learn about performance monitoring for Cloud Volumes ONTAP workloads

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

### Performance technical reports

- Cloud Volumes ONTAP for AWS

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

- Cloud Volumes ONTAP for Microsoft Azure

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

- Cloud Volumes ONTAP for Google Cloud

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

### CPU performance

Cloud Volumes ONTAP nodes show as highly utilized (over 90%) from your cloud provider's monitoring tools. This is because ONTAP reserves all vCPUs presented to the virtual machine so that they are available when needed.

For information, refer to the [NetApp knowledgebase article about how to monitor ONTAP CPU utilization using the CLI](#)

# License management for node-based BYOL

Each Cloud Volumes ONTAP system that has a node-based bring your own license (BYOL) must have a system license installed with an active subscription. BlueXP simplifies the process by managing licenses for you and by displaying a warning before they expire.



A node-based license is the previous generation license for Cloud Volumes ONTAP. A node-based license could be procured from NetApp (BYOL) and is available for license renewals, only in specific cases.

[Learn more about Cloud Volumes ONTAP licensing options.](#)

[Learn more about how to manage node-based licenses.](#)

## BYOL system licenses

Node-based licenses could be procured from NetApp. The number of licenses that you can purchase for a single node system or HA pair is unlimited.



NetApp has restricted the purchase, extension, and renewal of BYOL licensing. For more information, refer to [Restricted availability of BYOL licensing for Cloud Volumes ONTAPP](#).

A node-based license provides up to 368 TiB of capacity for a single node or HA pair. You might have purchased multiple licenses for a Cloud Volumes ONTAP BYOL system to allocate more than 368 TiB of capacity. For example, you might have two licenses to allocate up to 736 TiB of capacity to Cloud Volumes ONTAP. Or you could have four licenses to get up to 1.4 PiB.

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

## License management for a new system

When you create a node-based BYOL system, BlueXP prompts you for the serial number of your license and your NetApp Support Site account. BlueXP uses the account to download the license file from NetApp and to install it on the Cloud Volumes ONTAP system.

[Learn how to add NetApp Support Site accounts to BlueXP.](#)

If BlueXP can't access the license file over the secure internet connection, you can [obtain the file yourself and then manually upload the file to BlueXP](#).

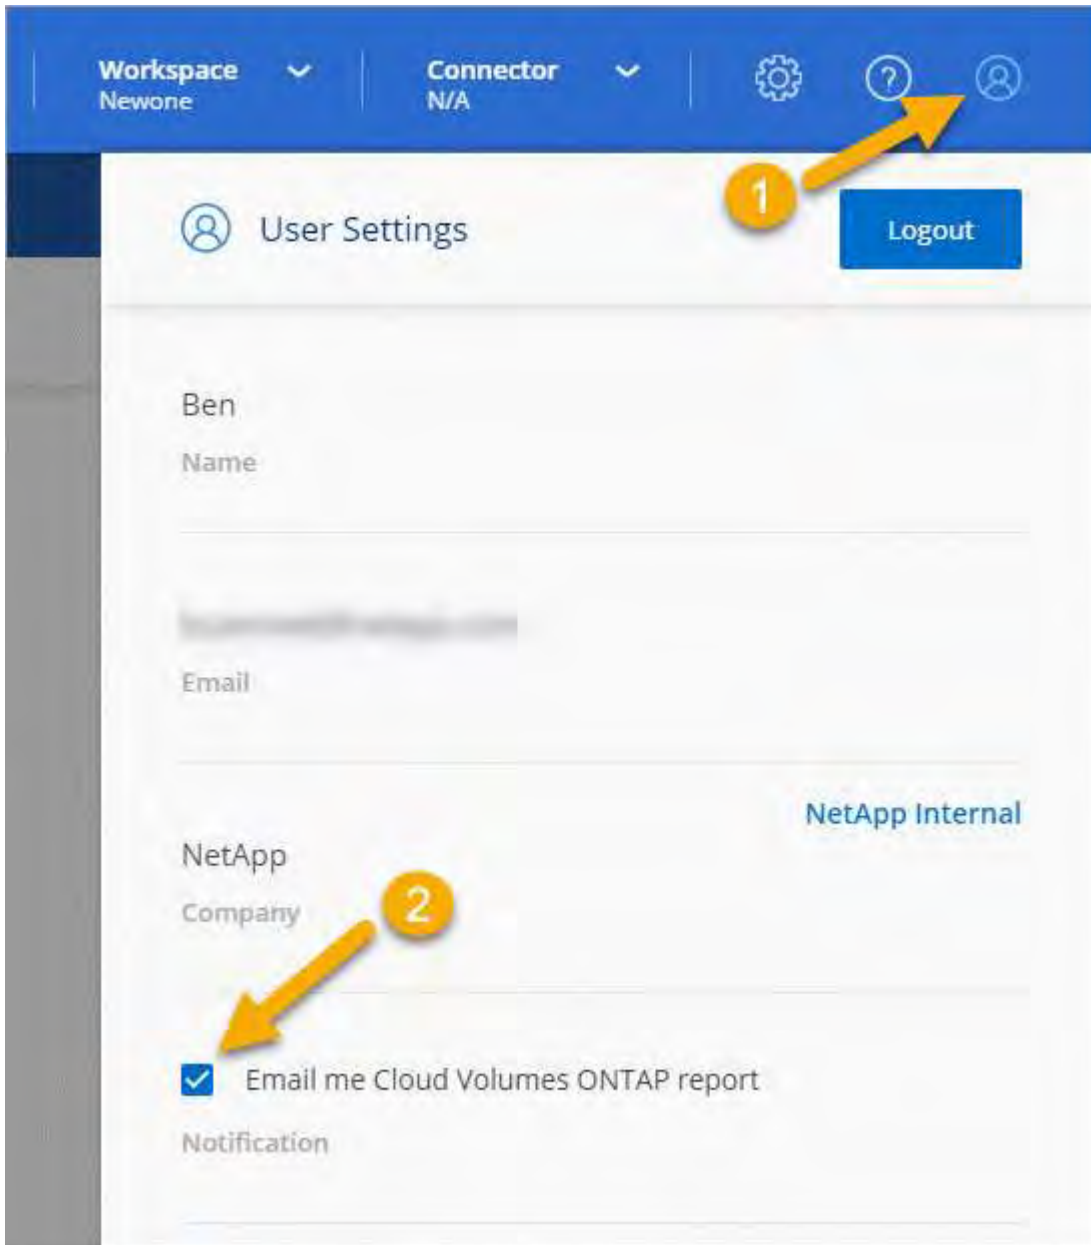
## License expiration

BlueXP displays a warning 30 days before a node-based license is due to expire and again when the license expires. The following image shows a 30-day expiration warning that appears in the user interface:



You can select the working environment to review the message.

BlueXP includes a license expiration warning in the Cloud Volumes ONTAP report that's emailed to you, if you are a BlueXP Organization or Account admin and you enabled the option:



The emailed report includes the license expiration warning every 2 weeks.

If you don't renew the license in time, the Cloud Volumes ONTAP system shuts itself down. If you restart it, it shuts itself down again.

## License renewal

If you renew a node-based BYOL subscription by contacting a NetApp representative, BlueXP automatically obtains the new license from NetApp and installs it on the Cloud Volumes ONTAP system.

If BlueXP can't access the license file over the secure internet connection, you can [obtain the file yourself and then manually upload the file to BlueXP](#).

## License transfer to a new system

A node-based BYOL license is transferable between Cloud Volumes ONTAP systems when you delete an existing system and then create a new one using the same license.

For example, you might want to delete an existing licensed system and then use the license with a new BYOL system in a different VPC/VNet or cloud provider. Note that only *cloud-agnostic* serial numbers work in any cloud provider. Cloud-agnostic serial numbers start with the *908xxxx* prefix.

It's important to note that your BYOL license is tied to your company and a specific set of NetApp Support Site credentials.

## Learn how AutoSupport and Digital Advisor are used for Cloud Volumes ONTAP

The AutoSupport component of ONTAP collects telemetry and sends it for analysis. Active IQ Digital Advisor (also known as Digital Advisor) analyzes the data from AutoSupport and provides proactive care and optimization. Using artificial intelligence, Digital Advisor can identify potential problems and help you resolve them before they impact your business.

Digital Advisor enables you to optimize your data infrastructure across your global hybrid cloud by delivering actionable predictive analytics and proactive support through a cloud-based portal and mobile app. Data-driven insights and recommendations from Digital Advisor are available to all NetApp customers with an active SupportEdge contract (features vary by product and support tier).

Here are some things you can do with Digital Advisor:

- Plan upgrades.

Digital Advisor identifies issues in your environment that can be resolved by upgrading to a newer version of ONTAP and the Upgrade Advisor component helps you plan for a successful upgrade.

- View system wellness.

Your Digital Advisor dashboard reports any issues with wellness and helps you correct those issues. Monitor system capacity to make sure you never run out of storage space. View support cases for your system.

- Manage performance.

Digital Advisor shows system performance over a longer period than you can see in ONTAP System Manager. Identify configuration and system issues that are impacting your performance. Maximize efficiency. View storage efficiency metrics and identify ways to store more data in less space.

- View inventory and configuration.

Digital Advisor displays complete inventory and software and hardware configuration information. See when service contracts are expiring and renew them to ensure you remain supported.

#### Related links

- [NetApp Documentation: Digital Advisor](#)
- [Launch Digital Advisor](#)
- [SupportEdge Services](#)

## Supported default configurations for Cloud Volumes ONTAP

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

### Default setup

- BlueXP creates one data-serving storage VM when it deploys Cloud Volumes ONTAP. Some configurations support additional storage VMs. [Learn more about managing storage VMs.](#)

Starting with the BlueXP 3.9.5 release, logical space reporting is enabled on the initial storage VM. When space is reported logically, ONTAP reports the volume space such that all the physical space saved by the storage efficiency features are also reported as used. For information about inline storage efficiency features, refer to the knowledge base article [KB: What Inline Storage Efficiency features are supported with CVO?](#)

- BlueXP automatically installs the following ONTAP feature licenses on Cloud Volumes ONTAP:
  - CIFS
  - FlexCache
  - FlexClone
  - iSCSI
  - Multi-tenant Encryption Key Management (MTEKM), starting with Cloud Volumes ONTAP 9.12.1 GA
  - NetApp Volume Encryption (only for bring your own license (BYOL) or registered pay-as-you-go (PAYGO) systems)
  - NFS
  - ONTAP S3

Starting with Cloud Volumes ONTAP 9.11.0 in AWS

Starting with Cloud Volumes ONTAP 9.9.1 in Azure

- SnapMirror
  - SnapRestore
  - SnapVault
- Several network interfaces are created by default:

- A cluster management LIF
- An intercluster LIF
- An SVM management LIF on HA systems in Azure
- An SVM management LIF on HA systems in Google Cloud
- An SVM management LIF on single node systems in AWS
- A node management LIF

In Google Cloud, this LIF is combined with the intercluster LIF.

- An iSCSI data LIF
- A CIFS and NFS data LIF



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

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

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

You can use the backups for reconfiguring your Cloud Volumes ONTAP system. For more information about configuration backups, refer to the [ONTAP documentation](#).

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

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

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

For information about these attributes, refer to [ONTAP volume create man page](#).

## Internal disks for system data

In addition to the storage for user data, BlueXP also purchases cloud storage for system data.

### AWS

- Three disks per node for boot, root, and core data:
  - 47 GiB io1 disk for boot data
  - 140 GiB gp3 disk for root data
  - 540 GiB gp2 disk for core data
- For HA pairs:
  - Two st1 EBS volumes for the mediator instance, one of approximately 8 GiB as root disk, and one of 4 GiB as data disk
  - One 140 GiB gp3 disk in each node to contain a copy of the root data of the other node



In some zones, the available EBS disk type can only be gp2.

- One EBS snapshot for each boot disk and root disk



Snapshots are created automatically upon reboot.

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



In AWS, NVRAM is on the boot disk.

### Azure (single node)

- Three Premium SSD disks:
  - One 10 GiB disk for boot data
  - One 140 GiB disk for root data
  - One 512 GiB disk for NVRAM

If the virtual machine that you chose for Cloud Volumes ONTAP supports Ultra SSDs, then the system uses a 32 GiB Ultra SSD for NVRAM, rather than a Premium SSD.

- One 1024 GiB Standard HDD disk for saving cores
- One Azure snapshot for each boot disk and root disk
- Every disk by default in Azure is encrypted at rest.

If the virtual machine that you chose for Cloud Volumes ONTAP supports Premium SSD v2 Managed Disk as data disks, the system uses a 32 GiB Premium SSD v2 Managed Disk for NVRAM, and another one as the root disk.

## Azure (HA pair)

### HA pairs with page blob

- Two 10 GiB Premium SSD disks for the boot volume (one per node)
- Two 140 GiB Premium Storage page blobs for the root volume (one per node)
- Two 1024 GiB Standard HDD disks for saving cores (one per node)
- Two 512 GiB Premium SSD disks for NVRAM (one per node)
- One Azure snapshot for each boot disk and root disk



Snapshots are created automatically upon reboot.

- Every disk by default in Azure is encrypted at rest.

### HA pairs with shared managed disks in multiple availability zones

- Two 10 GiB Premium SSD disks for the boot volume (one per node)
- Two 512 GiB Premium SSD disks for the root volume (one per node)
- Two 1024 GiB Standard HDD disks for saving cores (one per node)
- Two 512 GiB Premium SSD disks for NVRAM (one per node)
- One Azure snapshot for each boot disk and root disk



Snapshots are created automatically upon reboot.

- Every disk by default in Azure is encrypted at rest.

### HA pairs with shared managed disks in single availability zones

- Two 10 GiB Premium SSD disks for the boot volume (one per node)
- Two 512 GiB Premium SSD Shared Managed disks for the root volume (one per node)
- Two 1024 GiB Standard HDD disks for saving cores (one per node)
- Two 512 GiB Premium SSD Managed disks for NVRAM (one per node)

If your virtual machine supports Premium SSD v2 Managed Disks as data disks, it uses 32 GiB Premium SSD v2 Managed Disks for NVRAM and 512 GiB Premium SSD v2 Shared Managed disks for the root volume.

You can deploy HA pairs in a single single availability zone and use Premium SSD v2 Managed Disks when the following conditions are fulfilled:

- The version of Cloud Volumes ONTAP is 9.15.1 or later.
- The selected region and zone support Premium SSD v2 Managed Disks. For information about the supported regions, refer to [Microsoft Azure website: Products available by region](#).
- The subscription is registered for the Microsoft [Microsoft.Compute/VMOrchestratorZonalMultiFD](#) feature.

## Google Cloud (single node)

- One 10 GiB SSD persistent disk for boot data
- One 64 GiB SSD persistent disk for root data
- One 500 GiB SSD persistent disk for NVRAM



- One 315 GiB Standard persistent disk for saving cores
- Snapshots for boot and root data



Snapshots are created automatically upon reboot.

- Boot and root disks are encrypted by default.

### Google Cloud (HA pair)

- Two 10 GiB SSD persistent disks for boot data
- Four 64 GiB SSD persistent disk for root data
- Two 500 GiB SSD persistent disk for NVRAM
- Two 315 GiB Standard persistent disk for saving cores
- One 10 GiB Standard persistent disk for mediator data
- One 10 GiB Standard persistent disk for mediator boot data
- Snapshots for boot and root data



Snapshots are created automatically upon reboot.

- Boot and root disks are encrypted by default.

### Where the disks reside

BlueXP lays out the storage as follows:

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

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

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

# 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](#)
- [Google Cloud NetApp Volumes](#)

## Support registration overview

There are two forms of registration to activate support entitlement:

- Registering your BlueXP account serial number (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 BlueXP for NetApp support

To register for support and activate support entitlement, one user in your BlueXP organization (or 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 BlueXP organization 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 organization is not registered for support. As long as one user in the organization has followed these steps, then your organization 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.



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 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 organization 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 organization is different than the NSS account that is associated with a BlueXP user login.

These NSS credentials are associated with your specific BlueXP organization ID. Users who belong to the BlueXP organization 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 24/7, 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](#)
- [Google Cloud NetApp Volumes](#)

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 organization (or account), project (or 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 Working Environment


Select 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)  Upload 

No files selected  

### 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.

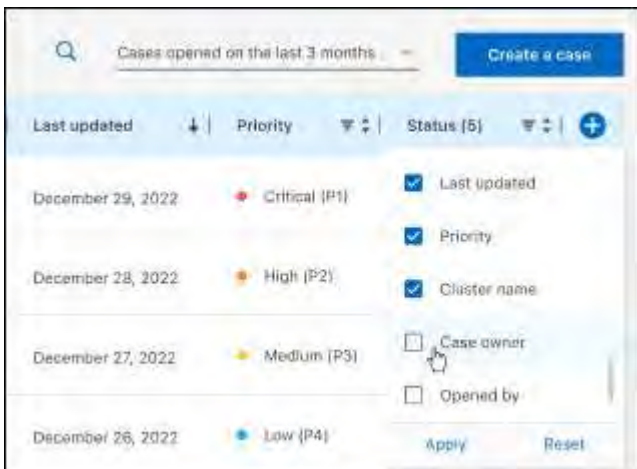
3. Optionally modify the information that displays in the table:
  - Under **Organization's cases**, select **View** to view all cases associated with your company.
  - Modify the date range by choosing an exact date range or by choosing a different time frame.



- Filter the contents of the columns.



- Change the columns that appear in the table by selecting  and then choosing the columns that you'd like to display.




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 Cloud Volumes ONTAP](#)
- [Notice for ONTAP](#)

## Copyright information

Copyright © 2025 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.



# **BlueXP classification documentation**

## **BlueXP classification**

NetApp  
August 13, 2025

# Table of Contents

- BlueXP classification documentation . . . . . 1
- Release notes . . . . . 2
  - What’s new in BlueXP classification . . . . . 2
    - 14 July 2025 . . . . . 2
    - 10 June 2025 . . . . . 2
    - 12 May 2025 . . . . . 3
    - 14 April 2025 . . . . . 4
    - 10 March 2025 . . . . . 4
    - 19 February 2025 . . . . . 4
    - 22 January 2025 . . . . . 5
    - 16 December 2024 . . . . . 6
    - 4 November 2024 . . . . . 6
    - 10 October 2024 . . . . . 6
    - 2 September 2024 . . . . . 7
    - 05 August 2024 . . . . . 7
    - 01 July 2024 . . . . . 7
    - 05 June 2024 . . . . . 8
    - 15 May 2024 . . . . . 8
    - 01 April 2024 . . . . . 8
    - 04 March 2024 . . . . . 9
    - 10 January 2024 . . . . . 9
    - 14 December 2023 . . . . . 10
    - 06 November 2023 . . . . . 10
    - 04 October 2023 . . . . . 10
    - 05 September 2023 . . . . . 10
    - 17 July 2023 . . . . . 11
    - 06 June 2023 . . . . . 11
    - 03 April 2023 . . . . . 12
    - 07 March 2023 . . . . . 12
    - 05 February 2023 . . . . . 13
    - 09 January 2023 . . . . . 14
  - Known limitations in BlueXP classification . . . . . 15
    - BlueXP classification disabled options . . . . . 15
    - BlueXP classification scanning . . . . . 15
- Get started . . . . . 17
  - Learn about BlueXP classification . . . . . 17
    - Features . . . . . 17
    - Supported working environments and data sources . . . . . 18
    - Cost . . . . . 18
    - The BlueXP classification instance . . . . . 19
    - How BlueXP classification scanning works . . . . . 20
    - What’s the difference between Mapping and Classification scans . . . . . 21
    - Information that BlueXP classification categorizes . . . . . 21

Networking overview	22
User roles in BlueXP classification	22
Access BlueXP classification	22
Deploy BlueXP classification	23
Which BlueXP classification deployment should you use?	23
Deploy BlueXP classification in the cloud using BlueXP	23
Install BlueXP classification on a host that has internet access	33
Install BlueXP classification on a Linux host with no internet access	43
Check that your Linux host is ready to install BlueXP classification	52
Activate scanning on your data sources	57
Scan data sources overview with BlueXP classification	57
Scan Azure NetApp Files volumes with BlueXP classification	61
Scan Amazon FSx for ONTAP volumes with BlueXP classification	64
Scan Cloud Volumes ONTAP and on-premises ONTAP volumes with BlueXP classification	69
Scan database schemas with BlueXP classification	73
Scan file shares with BlueXP classification	75
Scan StorageGRID data with BlueXP classification	79
Integrate your Active Directory with BlueXP classification	81
Supported data sources	82
Connect to your Active Directory server	82
Manage your Active Directory integration	84
Use BlueXP classification	85
View governance details about the data stored in your organization with BlueXP classification	85
Review the Governance dashboard	85
Create the Data Discovery Assessment Report	88
Create the Data Mapping Overview Report	88
View compliance details about the private data stored in your organization with BlueXP classification	91
View files that contain personal data	92
View files that contain sensitive personal data	94
View files by categories	96
View files by file types	96
Categories of private data in BlueXP classification	97
Types of personal data	97
Types of sensitive personal data	102
Types of categories	103
Types of files	104
Accuracy of information found	104
Create a custom classification in BlueXP classification	105
Create a custom classification	105
Investigate the data stored in your organization with BlueXP classification	107
Filter data in the Data Investigation page	107
View file metadata	110
View users' permissions for files and directories	111
Check for duplicate files in your storage systems	112
Create the Data Investigation Report	113



Create a saved search based on selected filters . . . . .	116
Manage saved searches with BlueXP classification . . . . .	117
View saved searches results in the Investigation page . . . . .	117
Create custom saved searches . . . . .	117
Edit saved searches . . . . .	119
Delete saved searches . . . . .	121
Default searches . . . . .	121
Change the BlueXP classification scan settings for your repositories . . . . .	121
View the scan status for your repositories . . . . .	122
Change the type of scanning for a repository . . . . .	123
Prioritize scans . . . . .	124
Stop scanning for a repository . . . . .	124
Pause and resume scanning for a repository . . . . .	125
View BlueXP classification compliance reports . . . . .	126
Select the working environments for reports . . . . .	127
Data Subject Access Request Report . . . . .	127
Health Insurance Portability and Accountability Act (HIPAA) Report . . . . .	129
Payment Card Industry Data Security Standard (PCI DSS) Report . . . . .	130
Privacy Risk Assessment Report . . . . .	131
Manage BlueXP classification . . . . .	134
Exclude specific directories from BlueXP classification scans . . . . .	134
Supported data sources . . . . .	134
Define the directories to exclude from scanning . . . . .	134
Examples . . . . .	135
Escaping special characters in folder names . . . . .	136
View the current exclusion list . . . . .	137
Define additional group IDs as open to organization in BlueXP classification . . . . .	137
Add the "open to organization" permission to group IDs . . . . .	137
View the current list of group IDs . . . . .	138
Remove data sources from BlueXP classification . . . . .	138
Deactivate compliance scans for a working environment . . . . .	138
Remove a database from BlueXP classification . . . . .	138
Remove a group of file shares from BlueXP classification . . . . .	139
Uninstall BlueXP classification . . . . .	139
Uninstall BlueXP classification from a cloud deployment . . . . .	139
Uninstall BlueXP classification from an on-premises deployment . . . . .	140
Deprecated features . . . . .	141
BlueXP classification deprecated features . . . . .	141
Supported data sources . . . . .	141
Compliance features . . . . .	141
Features to manage your data . . . . .	142
Deploy BlueXP classification deprecations . . . . .	142
Install BlueXP classification on multiple hosts for large configurations with no internet access . . . . .	143
Scan data deprecations . . . . .	144
Scan Amazon S3 buckets with BlueXP classification . . . . .	144

Scan OneDrive accounts with BlueXP classification . . . . .	151
Scan SharePoint accounts with BlueXP classification . . . . .	155
Scan Google Drive accounts with BlueXP classification . . . . .	159
Scan StorageGRID data with BlueXP classification . . . . .	161
Manage data deprecations . . . . .	164
View governance details about your data using the BlueXP classification Governance dashboard . . . . .	164
Organize your private data with BlueXP classification . . . . .	166
Manage your private data with BlueXP classification . . . . .	174
Add personal data identifiers to your BlueXP classification scans . . . . .	185
View the status of your compliance actions in BlueXP classification . . . . .	200
Audit the history of BlueXP classification actions . . . . .	201
Reducing the BlueXP classification scan speed . . . . .	202
Remove a OneDrive, SharePoint, or Google Drive account from BlueXP classification . . . . .	203
Reference . . . . .	204
Supported BlueXP classification instance types . . . . .	204
AWS instance types . . . . .	204
Azure instance types . . . . .	204
GCP instance types . . . . .	205
Metadata collected from data sources in BlueXP classification . . . . .	205
Last access time timestamp . . . . .	205
Log in to the BlueXP classification system . . . . .	206
BlueXP classification APIs . . . . .	207
Overview . . . . .	207
Accessing the Swagger API reference . . . . .	208
Example using the APIs . . . . .	208
Knowledge and support . . . . .	218
Register for BlueXP support . . . . .	218
Support registration overview . . . . .	218
Register BlueXP for NetApp support . . . . .	218
Associate NSS credentials for Cloud Volumes ONTAP support . . . . .	220
Get help for BlueXP classification . . . . .	222
Get support for a cloud provider file service . . . . .	222
Use self-support options . . . . .	222
Create a case with NetApp support . . . . .	222
Manage your support cases (Preview) . . . . .	225
Frequently asked questions about BlueXP classification . . . . .	228
BlueXP classification service . . . . .	228
How does BlueXP classification work? . . . . .	228
Does BlueXP classification have a REST API, and does it work with third-party tools? . . . . .	228
Is BlueXP classification available through the cloud marketplaces? . . . . .	228
BlueXP classification scanning and analytics . . . . .	228
How often does BlueXP classification scan my data? . . . . .	228
Does scan performance vary? . . . . .	229
Can I search my data using BlueXP classification? . . . . .	229
BlueXP classification management and privacy . . . . .	229

How do I enable or disable BlueXP classification? .....	229
Can the service exclude scanning data in certain directories? .....	230
Are snapshots that reside on ONTAP volumes scanned? .....	230
What happens if data tiering is enabled on your ONTAP volumes? .....	230
Types of source systems and data types .....	230
Are there any restrictions when deployed in a Government region? .....	230
What data sources can I scan if I install BlueXP classification in a site without internet access? .....	230
Which file types are supported? .....	230
What kinds of data and metadata does BlueXP classification capture? .....	231
Can I limit BlueXP classification information to specific users? .....	231
Can anyone access the private data sent between my browser and BlueXP classification? .....	231
How is sensitive data handled? .....	231
Where is the data stored? .....	232
How is the data accessed? .....	232
Licenses and costs .....	232
How much does BlueXP classification cost? .....	232
Connector deployment .....	232
What is the Connector? .....	232
Where does the Connector need to be installed? .....	232
Does BlueXP classification require access to credentials? .....	232
Does communication between the service and the Connector use HTTP? .....	232
BlueXP classification deployment .....	233
What deployment models does BlueXP classification support? .....	233
What type of instance or VM is required for BlueXP classification? .....	233
Can I deploy the BlueXP classification on my own host? .....	233
What about secure sites without internet access? .....	233
Legal notices .....	234
Copyright .....	234
Trademarks .....	234
Patents .....	234
Privacy policy .....	234
Open source .....	234

# BlueXP classification documentation

# Release notes

## What's new in BlueXP classification

Learn what's new in BlueXP classification.

### 14 July 2025

#### Version 1.45

This BlueXP classification release includes code changes that optimize resource utilization and:

#### Improved workflow to add file shares for scanning

The workflow to add files shares to a file share group has been simplified. The process also now differentiates CIFS protocol support based on authentication type (Kerberos or NTLM).

For more information, see [Scan file shares](#).

#### Enhanced file owner information

You can now view more information about file owners for files captured in the Investigation tab. When viewing metadata for a file in the Investigation tab, locate the file owner then select **View details** to see the username, email, and SAM account name. You can also view other items owned by this user. This feature is only available for working environments with Active Directory.

For more information, see [Investigate the data stored in your organization](#).

### 10 June 2025

#### Version 1.44

This BlueXP classification release includes:

#### Improved update times for the Governance dashboard

Update times for individual components of the Governance dashboard have been improved. The following table displays the frequency of updates for each component.

Component	Update times
Age of Data	24 hours
Categories	24 hours
Data Overview	5 minutes
Duplicate Files	2 hours
File Types	24 hours
Non-Business Data	2 hours
Open Permissions	24 hours
Saved Searches	2 hours
Sensitive Data and Wide Permissions	24 hours

Component	Update times
Size of Data	24 hours
Stale Data	2 hours
Top Data Repositories by Sensitivity Level	2 hours

You can view the time of the last update and manually update the Duplicate Files, Non-Business Data, Saved Searches, Stale Data, and Top Data Repositories by Sensitivity Level components. For more information about the Governance dashboard, see [View governance details about the data stored in your organization](#).

**Performance and security improvements**

Enhancements have been made to improve BlueXP classification’s performance, memory consumption, and security.

**Bug fixes**

Redis has been upgraded to improve the reliability of BlueXP classification. BlueXP classification now uses Elasticsearch to improve the accuracy of file count reporting during scans.

**12 May 2025**

**Version 1.43**

This BlueXP classification release includes:

**Prioritize classification scans**

BlueXP classification supports the ability to prioritize Map & Classify scans in addition to Mapping-only scans, enabling you to select which scans are completed first. Prioritization of Map & Classify scans is supported during and before the scans begin. If you choose to prioritize a scan while it’s in progress, both the mapping and classification scans are prioritized.

For more information, see [Prioritize scans](#).

**Support for Canadian personally identifiable information (PII) data categories**

BlueXP classification scans identify Canadian PII data categories. These categories include banking information, passport numbers, social insurance numbers, driver’s license numbers and health card numbers for all Canadian provinces and territories.

For more information, see [Personal data categories](#).

**Custom classification (preview)**

BlueXP classification supports custom classifications for Map & Classify scans. With custom classifications, you can tailor BlueXP scans to capture data specific to your organization using regular expressions. This feature is currently in preview.

For more information, see [Add custom classifications](#).

**Saved searches tab**

The **Policies** tab has been renamed **Saved searches**. The functionality is unchanged.

**Send scan events to BlueXP timeline**

BlueXP classification supports sending classification events (when a scan is initiated and when it ends) to the [BlueXP timeline](#).

## Security updates

- The Keras package has been updated, mitigating vulnerabilities (BDSA-2025-0107 and BDSA-2025-1984).
- The Docker containers configuration has been updated. The container no longer has access to the host's network interfaces for crafting raw network packets. By reducing unnecessary access, the update mitigates potential security risks.

## Performance enhancements

Code enhancements have been implemented to reduce RAM usage and improve the overall performance of BlueXP classification.

## Bug fixes

Bugs that caused StorageGRID scans to fail, the investigation page filter options to not load, and the Data Discovery Assessment to not download for high volume assessments have been fixed.

## 14 April 2025

### Version 1.42

This BlueXP classification release includes:

#### Bulk scanning for working environments

BlueXP classification supports bulk operations for working environments. You can choose to enable Mapping scans, enable Map & Classify scans, disable scans, or create a custom configuration across volumes in working environment. If you make a selection for an individual volume, it overrides the bulk selection. To perform a bulk operation, navigate to the **Configuration** page and make your selection.

#### Download investigation report locally

BlueXP classification supports the ability to download data investigation reports locally to view in the browser. If you choose the local option, the data investigation is only available in the CSV format and only displays the first 10,000 rows of data.

For more information, see [Investigate the data stored in your organization with BlueXP classification](#).

## 10 March 2025

### Version 1.41

This BlueXP classification release includes general improvements and bug fixes. It also includes:

#### Scan status

BlueXP classification tracks the real time progress of the *initial* mapping and classification scans on a volume. Separate progressive bars track the mapping and classification scans, presenting a percentage of total files scanned. You can also hover over a progress bar to view the number of files scanned and the total files. Tracking the status of your scans creates deeper insights into the scan progress, enabling you to better plan your scans and understand resource allocation.

To view the status of your scans, navigate to **Configuration** in BlueXP classification then select the **Working Environment configuration**. Progress is displayed in line for each volume.

## 19 February 2025

## Version 1.40

This BlueXP classification release includes the following updates.

### Support for RHEL 9.5

This release provides support for Red Hat Enterprise Linux v9.5 in addition to previously supported versions. This is applicable to any manual on-premises installation of BlueXP classification, including dark site deployments.

The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater: Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4, and 9.5.

### Prioritize mapping-only scans

When conducting Mapping-only scans, you can prioritize the most important scans. This feature helps when you have many working environments and want to ensure high priority scans are completed first.

By default, scans are queued based on the order in which they are initiated. With the ability to prioritize scans, you can move scans to the front of the queue. Multiple scans can be prioritized. Priority is designated in a first-in, first-out order, meaning the first scan you prioritize moves to the front of the queue; the second scan you prioritize becomes second in the queue, and so forth.

Priority is granted on a one-time basis. Automatic rescans of mapping data occur in the default order.

Prioritization is limited to [mapping-only scans](#); it's not available for map and classify scans.

For more information, see [Prioritize scans](#).

### Retry all scans

BlueXP classification supports the ability to batch retry all failed scans.

You can reattempt scans in a batch operation with the **Retry all** function. If classification scans are failing due to a temporary issue such as a network outage, you can retry all scans at the same time with one button instead of retrying them individually. Scans can be retried as many times as needed.

To retry all scans:

1. From the BlueXP classification menu, select **Configuration**.
2. To retry all failed scans, select **Retry all scans**.

### Improved categorization model accuracy

The accuracy of the machine learning model for [predefined categories](#) has improved by 11%.

## 22 January 2025

### Version 1.39

This BlueXP classification release updates the export process for the Data Investigation report. This export update is useful for performing additional analyses on your data, creating additional visualizations on the data, or sharing the results of your data investigation with others.

Previously, the Data Investigation report export was limited to 10,000 rows. With this release, the limit has been removed so that you can export all of your data. This change enables you to export more data from your Data Investigation reports, providing you with more flexibility in your data analysis.



You can choose the working environment, volumes, destination folder, and either JSON or CSV format. The exported filename includes a timestamp to help you identify when the data was exported.

The supported working environments include:

- Cloud Volumes ONTAP
- FSx for ONTAP
- ONTAP
- Share group

Exporting data from the Data Investigation report has the following limitations:

- The maximum number of records to download is 500 million. per type (files, directories, and tables)
- One million records are expected to take about 35 minutes to export.

For details about data investigation and the report, see [Investigate data stored in your organization](#).

## 16 December 2024

### Version 1.38

This BlueXP classification release includes general improvements and bug fixes.

## 4 November 2024

### Version 1.37

This BlueXP classification release includes the following updates.

#### Support for RHEL 8.10

This release provides support for Red Hat Enterprise Linux v8.10 in addition to previously supported versions. This is applicable to any manual on-premises installation of BlueXP classification, including dark site deployments.

The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater: Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, and 9.4.

Learn more about [BlueXP classification](#).

#### Support for NFS v4.1

This release provides support for NFS v4.1 in addition to previously supported versions.

Learn more about [BlueXP classification](#).

## 10 October 2024

### Version 1.36

#### Support for RHEL 9.4

This release provides support for Red Hat Enterprise Linux v9.4 in addition to previously supported versions. This is applicable to any manual on-premises installation of BlueXP classification, including dark site

deployments.

The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater: Red Hat Enterprise Linux version 8.8, 9.0, 9.1, 9.2, 9.3, and 9.4.

Learn more about [BlueXP classification deployments overview](#).

### **Improved scan performance**

This release provides improved scan performance.

## **2 September 2024**

### **Version 1.35**

#### **Scan StorageGRID data**

BlueXP classification supports scanning data in StorageGRID.

For details, refer to [Scan StorageGRID data](#).

## **05 August 2024**

### **Version 1.34**

This BlueXP classification release includes the following update.

#### **Change from CentOS to Ubuntu**

BlueXP classification has updated its Linux operating system for Microsoft Azure and Google Cloud Platform (GCP) from CentOS 7.9 to Ubuntu 22.04.

For deployment details, refer to [Install on a Linux host with internet access and prepare the Linux host system](#).

## **01 July 2024**

### **Version 1.33**

#### **Ubuntu supported**

This release supports the Ubuntu 24.04 Linux platform.

#### **Mapping scans gather metadata**

The following metadata is extracted from files during mapping scans and is displayed on the Governance, Compliance, and Investigation dashboards:

- Working environment
- Working environment type
- Storage repository
- File type
- Used capacity
- Number of files
- File size

- File creation
- File last access
- File last modified
- File discovered time
- Permissions extraction

### **Additional data in dashboards**

This release updates which data appears in the Governance, Compliance, and Investigation dashboards during mapping scans.

For details, see [What's the difference between mapping and classification scans](#).

## **05 June 2024**

### **Version 1.32**

#### **New Mapping status column in the Configuration page**

This release now shows a new Mapping status column in the Configuration page. The new column helps you identify if the mapping is running, queued, paused or more.

For explanations of the statuses, see [Change scan settings](#).

## **15 May 2024**

### **Version 1.31**

#### **Classification is available as a core service within BlueXP**

BlueXP classification is now available as a core capability within BlueXP at no additional charge for up to 500 TiB of scanned data per connector. No Classification license or paid subscription is required. As we focus BlueXP classification functionality on scanning NetApp storage systems with this new version, some legacy functionality will only be available to customers who had previously paid for a license. The use of those legacy features will expire when the paid contract reaches its end date.



BlueXP classification does not impose a limit on the amount of data it can scan. Each Connector supports scanning and displaying 500 TiB of data. To scan more than 500 TiB of data, [install another Connector](#) then [deploy another classification instance](#).

The BlueXP UI displays data from a single connector. For tips on viewing data from multiple Connectors, see [Work with multiple Connectors](#).

[Learn more about the deprecated features](#).

## **01 April 2024**

### **Version 1.30**

#### **Support added for RHEL v8.8 and v9.3 BlueXP classification**

This release provides support for Red Hat Enterprise Linux v8.8 and v9.3 in addition to previously supported 9.x, which requires Podman, rather than the Docker engine. This is applicable to any manual on-premises installation of BlueXP classification.

The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater: Red Hat Enterprise Linux version 8.8, 9.0, 9.1, 9.2, and 9.3.

Learn more about [BlueXP classification deployments overview](#).

BlueXP classification is supported if you install the Connector on a RHEL 8 or 9 host that resides on-premises. It's not supported if the RHEL 8 or 9 host resides in AWS, Azure, or Google Cloud.

#### **Option to activate audit log collection removed**

The option to activate audit log collection has been disabled.

#### **Scan speed improved**

Scan performance on secondary scanner nodes has been improved. You can add more scanner nodes if you need additional processing power for your scans. For details, refer to [Install BlueXP classification on a host that has internet access](#).

#### **Automatic upgrades**

If you deployed BlueXP classification on a system with internet access, the system upgrades automatically. Previously, the upgrade occurred after a specific time elapsed since the last user activity. With this release, BlueXP classification upgrades automatically if the local time is between 1:00 AM and 5:00 AM. If the local time is outside of these hours, the upgrade occurs after a specific time elapses since the last user activity. For details, refer to [Install on a Linux host with internet access](#).

If you deployed BlueXP classification without internet access, you'll need to upgrade manually. For details, refer to [Install BlueXP classification on a Linux host with no internet access](#).

## **04 March 2024**

### **Version 1.29**

#### **Now you can exclude scanning data that resides in certain data source directories**

If you want BlueXP classification to exclude scanning data that resides in certain data source directories, you can add these directory names to a configuration file that BlueXP classification processes. This feature enables you to avoid scanning directories that are unnecessary, or that would result in returning false positive personal data results.

[Learn more](#).

#### **Extra Large instance support is now qualified**

If you need BlueXP classification to scan more than 250 million files, you can use an Extra Large instance in your cloud deployment or on-premises installation. This type of system can scan up to 500 million files.

[Learn more](#).

## **10 January 2024**

### **Version 1.27**

#### **Investigation page results display the total size in addition to total number of items**

The filtered results in the Investigation page display the total size of the items in addition to the total number of files. This can help when moving files, deleting files, and more.

#### **Configure additional Group IDs as "Open to Organization"**

Now you can configure Group IDs in NFS to be considered as "Open to Organization" directly from BlueXP classification if the group had not initially been set with that permission. Any files and folders that have these group IDs attached will show as "Open to Organization" in the Investigation Details page. See how to [add additional Group IDs as "open to organization"](#).

## 14 December 2023

### Version 1.26.6

This release included some minor enhancements.

The release also removed the following options:

- The option to activate audit log collection has been disabled.
- During the Directories investigation, the option to calculate the number of personal identifiable information (PII) data by Directories is not available. Refer to [Investigate the data stored in your organization](#).
- The option to integrate data using Azure Information Protection (AIP) labels has been disabled. Refer to [Organize your private data](#).

## 06 November 2023

### Version 1.26.3

The following issues have been fixed in this release

- Fixed an inconsistency when presenting the number of files scanned by the system in dashboards.
- Improved the scanning behavior by handling and reporting on files and directories with special characters in the name and metadata.

## 04 October 2023

### Version 1.26

#### Support for on-premises installations of BlueXP classification on RHEL version 9

Red Hat Enterprise Linux versions 8 and 9 do not support the Docker engine; which was required for the BlueXP classification installation. We now support BlueXP classification installation on RHEL 9.0, 9.1, and 9.2 using Podman version 4 or greater as the container infrastructure. If your environment requires using the newest versions of RHEL, now you can install BlueXP classification (version 1.26 or greater) when using Podman.

At this time we don't supported dark site installations or distributed scanning environments (using a master and remote scanner nodes) when using RHEL 9.x.

## 05 September 2023

### Version 1.25

#### Small and medium deployments temporarily unavailable

When you deploy an instance of BlueXP classification in AWS, the option to select **Deploy > Configuration** and choose a small or medium-sized instance is unavailable at this time. You can still deploy the instance using the large instance size by selecting **Deploy > Deploy**.

## **Apply tags on up to 100,000 items from the Investigation Results page**

In the past you could only apply tags to a single page at a time in the Investigation Results page (20 items). Now you can select **all** items in the Investigation Results pages and apply tags to all the items - up to 100,000 items at a time. [See how](#).

## **Identify duplicated files with a minimum file size of 1 MB**

BlueXP classification used to identify duplicated files only when files were 50 MB or larger. Now duplicated files starting with 1 MB can be identified. You can use the Investigation page filters "File Size" along with "Duplicates" to see which files of a certain size are duplicated in your environment.

## **17 July 2023**

### **Version 1.24**

#### **Two new types of German personal data are identified by BlueXP classification**

BlueXP classification can identify and categorize files that contain the following types of data:

- German ID (Personalausweisnummer)
- German Social Security Number (Sozialversicherungsnummer)

[See all the types of personal data that BlueXP classification can identify in your data.](#)

#### **BlueXP classification is fully supported in Restricted mode and Private mode**

BlueXP classification is now fully supported in sites with no internet access (Private mode) and with limited outbound internet access (Restricted mode). [Learn more about BlueXP deployment modes for the Connector.](#)

#### **Ability to skip versions when upgrading a Private mode installation of BlueXP classification**

Now you can upgrade to a newer version of BlueXP classification even if it is not sequential. This means that the current limitation of upgrading BlueXP classification by one version at a time is no longer required. This feature is relevant starting from version 1.24 onwards.

#### **The BlueXP classification API is now available**

The BlueXP classification API enables you to perform actions, create queries, and export information about the data you are scanning. The interactive documentation is available using Swagger. The documentation is separated into multiple categories, including Investigation, Compliance, Governance, and Configuration. Each category is a reference to the tabs in the BlueXP classification UI.

[Learn more about the BlueXP classification APIs.](#)

## **06 June 2023**

### **Version 1.23**

#### **Japanese is now supported when searching for data subject names**

Japanese names can now be entered when searching for a subject's name in response to a Data Subject Access Request (DSAR). You can generate a [Data Subject Access Request report](#) with the resulting information. You can also enter Japanese names in the "[Data Subject](#)" filter in the [Data Investigation page](#) to identify files that contain the subject's name.

#### **Ubuntu is now a supported Linux distribution on which you can install BlueXP classification**

Ubuntu 22.04 has been qualified as a supported operating system for BlueXP classification. You can install BlueXP classification on a Ubuntu Linux host in your network, or on a Linux host in the cloud when using

version 1.23 of the installer. [See how to install BlueXP classification on a host with Ubuntu installed.](#)

### **Red Hat Enterprise Linux 8.6 and 8.7 are no longer supported with new BlueXP classification installations**

These versions are not supported with new deployments because Red Hat no longer supports Docker, which is a prerequisite. If you have an existing BlueXP classification machine running on RHEL 8.6 or 8.7, NetApp will continue to support your configuration.

### **BlueXP classification can be configured as an FPolicy Collector to receive FPolicy events from ONTAP systems**

You can enable file access audit logs to be collected on your BlueXP classification system for file access events detected on volumes in your working environments. BlueXP classification can capture the following types of FPolicy events and the users who performed the actions on your files: Create, Read, Write, Delete, Rename, Change owner/permissions, and Change SACL/DAACL.

### **Data Sense BYOL licenses are now supported in dark sites**

Now you can upload your Data Sense BYOL license into the BlueXP digital wallet in a dark site so that you are notified when your license is getting low.

## **03 April 2023**

### **Version 1.22**

#### **New Data Discovery Assessment Report**

The Data Discovery Assessment Report provides a high-level analysis of your scanned environment to highlight the system's findings and to show areas of concern and potential remediation steps. The goal of this report is to raise awareness of data governance concerns, data security exposures, and data compliance gaps of your data set. [See how to generate and use the Data Discovery Assessment Report.](#)

#### **Ability to deploy BlueXP classification on smaller instances in the cloud**

When deploying BlueXP classification from a BlueXP Connector in an AWS environment, now you can select from two smaller instance types than what is available with the default instance. If you are scanning a small environment this can help you save on cloud costs. However, there are some restrictions when using the smaller instance. [See the available instance types and limitations.](#)

#### **Standalone script is now available to qualify your Linux system prior to BlueXP classification installation**

If you would like to verify that your Linux system meets all prerequisites independently of running the BlueXP classification installation, there is a separate script you can download that only tests for the prerequisites. [See how to check if your Linux host is ready to install BlueXP classification.](#)

## **07 March 2023**

### **Version 1.21**

#### **New functionality to add your own custom categories from the BlueXP classification UI**

BlueXP classification now enables you to add your own custom categories so that BlueXP classification will identify the files that fit into those categories. BlueXP classification has many [predefined categories](#), so this feature enables you to add custom categories to identify where information that is unique to your organization are found in your data.

[Learn more.](#)

#### **Now you can add custom keywords from the BlueXP classification UI**

BlueXP classification has had the ability to add custom keywords that BlueXP classification will identify in future scans for a while. However, you needed to log into the BlueXP classification Linux host and use a command line interface to add the keywords. In this release, the ability to add custom keywords is in the BlueXP classification UI, making it very easy to add and edit these keywords.

[Learn more about adding custom keywords from the BlueXP classification UI.](#)

### **Ability to have BlueXP classification not scan files when the "last access time" will be changed**

By default, if BlueXP classification doesn't have adequate "write" permissions, the system won't scan files in your volumes because BlueXP classification can't revert the "last access time" to the original timestamp. However, if you don't care if the last access time is reset to the original time in your files, you can override this behavior in the Configuration page so that BlueXP classification will scan the volumes regardless of permissions.

In conjunction with this capability, a new filter named "Scan Analysis Event" has been added so you can view the files that were not classified because BlueXP classification couldn't revert last accessed time, or the files that were classified even though BlueXP classification couldn't revert last accessed time.

[Learn more about the "Last access time timestamp" and the permissions BlueXP classification requires.](#)

### **Three new types of personal data are identified by BlueXP classification**

BlueXP classification can identify and categorize files that contain the following types of data:

- Botswana Identity Card (Omang) Number
- Botswana Passport Number
- Singapore National Registration Identity Card (NRIC)

[See all the types of personal data that BlueXP classification can identify in your data.](#)

### **Updated functionality for directories**

- The "Light CSV Report" option for Data Investigation Reports now includes information from directories.
- The "Last Accessed" time filter now shows the last accessed time for both files and directories.

### **Installation enhancements**

- The BlueXP classification installer for sites without internet access (dark sites) now performs a pre-check to make sure your system and networking requirements are in place for a successful installation.
- Installation audit log files are saved now; they are written to `/ops/netapp/install_logs`.

## **05 February 2023**

### **Version 1.20**

#### **Ability to send Policy-based notification emails to any email address**

In earlier versions of BlueXP classification you could send email alerts to the BlueXP users in your account when certain critical Policies return results. This feature enables you to get notifications to protect your data when you're not online. Now you can also send email alerts from Policies to any other users - up to 20 email addresses - who are not in your BlueXP account.

[Learn more about sending email alerts based on Policy results.](#)

#### **Now you can add personal patterns from the BlueXP classification UI**



BlueXP classification has had the ability to add custom "personal data" that BlueXP classification will identify in future scans for a while. However, you needed to log into the BlueXP classification Linux host and use a command line to add the custom patterns. In this release, the ability to add personal patterns using a regex is in the BlueXP classification UI, making it very easy to add and edit these custom patterns.

[Learn more about adding custom patterns from the BlueXP classification UI.](#)

### **Ability to move 15 million files using BlueXP classification**

In the past you could have BlueXP classification move a maximum of 100,000 source files to any NFS share. Now you can move up to 15 million files at a time. [Learn more about moving source files using BlueXP classification.](#)

### **Ability to see the number of users who have access to SharePoint Online files**

The filter "Number of users with access" now supports files stored in SharePoint Online repositories. In the past only files on CIFS shares were supported. Note that SharePoint groups that are not active directory based will not be counted in this filter at this time.

### **New "Partial Success" status has been added to the Action Status panel**

The new "Partial Success" status indicates that a BlueXP classification action is finished and some items failed and some items succeeded, for example, when you are moving or deleting 100 files. Additionally, the "Finished" status has been renamed to "Success". In the past, the "Finished" status might list actions that succeeded and that failed. Now the "Success" status means that all actions succeeded on all items. [See how to view the Actions Status panel.](#)

## **09 January 2023**

### **Version 1.19**

### **Ability to view a chart of files that contain sensitive data and that are overly permissive**

The Governance dashboard has added a new *Sensitive Data and Wide Permissions* area that provides a heatmap of files that contain sensitive data (including both sensitive and sensitive personal data) and that are overly permissive. This can help you to see where you may have some risks with sensitive data. [Learn more.](#)

### **Three new filters are available in the Data Investigation page**

New filters are available to refine the results that display in the Data Investigation page:

- The "Number of users with access" filter shows which files and folders are open to a certain number of users. You can choose a number range to refine the results - for example, to see which files are accessible by 51-100 users.
- The "Created Time", "Discovered Time", "Last Modified", and "Last Accessed" filters now allow you to create a custom date range instead of just selecting a pre-defined range of days. For example, you can look for files with a "Created Time" "older than 6 months", or with a "Last Modified" date within the "last 10 days".
- The "File Path" filter now enables you to specify paths that you want to exclude from the filtered query results. If you enter paths to both include and exclude certain data, BlueXP classification finds all files in the included paths first, then it removes files from excluded paths, and then it displays the results.

[See the list of all the filters you can use to investigate your data.](#)

### **BlueXP classification can identify the Japanese Individual Number**

BlueXP classification can identify and categorize files that contain the Japanese Individual Number (also known as My Number). This includes both the Personal and Corporate My Number. [See all the types of](#)

personal data that BlueXP classification can identify in your data.

## Known limitations in BlueXP classification

Known limitations identify functions that are not supported or do not interoperate correctly in this release. Review these limitations carefully.

### BlueXP classification disabled options

The December 2023 (version 1.26.6) release removed the following options:

- The option to activate audit log collection has been disabled.
- During the Directories investigation, the option to calculate the number of personally identifiable information (PII) data by Directories is not available.
- The option to integrate data using Azure Information Protection (AIP) labels has been disabled.

### BlueXP classification scanning

The following limitations occur with BlueXP classifications scans.

#### BlueXP classification scans only one share under a volume

If you have multiple file shares under a single volume, BlueXP classification scans the share with the highest hierarchy. For example, if you have shares like the following:

- /A
- /A/B
- /C
- /D/E

In this configuration, only the data in /A is scanned. The data in /C and /D is not scanned.

#### Workaround

There is a workaround to make sure you are scanning data from all the shares in your volume. Follow these steps:

1. In the working environment, add the volume to be scanned.
2. After BlueXP classification has completed scanning the volume, go to the *Data Investigation* page and create a filter to see which share is being scanned:

Filter the data by "Working Environment Name" and "Directory Type = Share" to see which share is being scanned.

3. Get the complete list of shares that exist in the volume so you can see which shares are not being scanned.
4. [Add the remaining shares to a share group.](#)

Add all the shares individually, for example:

/C

/D

5. Perform these steps for each volume in the working environment that has multiple shares.

### **Last accessed timestamp**

When BlueXP classification conducts a scan of a directory, the scan impacts the directory's **Last accessed** field. When you view the **Last accessed** field, that metadata reflects either the date and time of the scan or the last time a user accessed the directory.

# Get started

## Learn about BlueXP classification

BlueXP classification (Cloud Data Sense) is a data governance service for BlueXP that scans your corporate on-premises and cloud data sources to map and classify data, and to identify private information. This can help reduce your security and compliance risk, decrease storage costs, and assist with your data migration projects.



Beginning with version 1.31, BlueXP classification is available as a core capability with BlueXP. There's no additional charge. No Classification license or subscription is required. If you've been using legacy version 1.30 or earlier, that version is available until your subscription expires. [See a list of deprecated features.](#)

### Features

BlueXP classification uses artificial intelligence (AI), natural language processing (NLP), and machine learning (ML) to understand the content that it scans in order to extract entities and categorize the content accordingly. This allows BlueXP classification to provide the following areas of functionality.

[Learn more about the use cases for BlueXP classification.](#)

#### Maintain compliance

BlueXP classification provides several tools that can help with your compliance efforts. You can use BlueXP classification to:

- Identify Personal Identifiable Information (PII).
- Identify a wide scope of sensitive personal information as required by GDPR, CCPA, PCI, and HIPAA privacy regulations.
- Respond to Data Subject Access Requests (DSAR) based on name or email address.

#### Strengthen security

BlueXP classification can identify data that is potentially at risk for being accessed for criminal purposes. You can use BlueXP classification to:

- Identify all the files and directories (shares and folders) with open permissions that are exposed to your entire organization or to the public.
- Identify sensitive data that resides outside of the initial, dedicated location.
- Comply with data retention policies.
- Use *Policies* to automatically detect new security issues so security staff can take action immediately.

#### Optimize storage usage

BlueXP classification provides tools that can help with your storage total cost of ownership (TCO). You can use BlueXP classification to:

- Increase storage efficiency by identifying duplicate or non-business-related data.
- Save storage costs by identifying inactive data that you can tier to less expensive object storage. [Learn more about tiering from Cloud Volumes ONTAP systems.](#) [Learn more about tiering from on-premises](#)

ONTAP systems.

## Supported working environments and data sources

BlueXP classification can scan and analyze structured and unstructured data from the following types of working environments and data sources:

### Working environments

- Amazon FSx for ONTAP
- Azure NetApp Files
- Cloud Volumes ONTAP (deployed in AWS, Azure, or GCP)
- On-premises ONTAP clusters
- StorageGRID

### Data sources

- NetApp file shares
- Databases:
  - Amazon Relational Database Service (Amazon RDS)
  - MongoDB
  - MySQL
  - Oracle
  - PostgreSQL
  - SAP HANA
  - SQL Server (MSSQL)

BlueXP classification supports NFS versions 3.x, 4.0, and 4.1, and CIFS versions 1.x, 2.0, 2.1, and 3.0.

## Cost

BlueXP classification is free to use. No Classification license or paid subscription is required.

### Infrastructure costs

- Installing BlueXP classification in the cloud requires deploying a cloud instance, which results in charges from the cloud provider where it is deployed. See [the type of instance that is deployed for each cloud provider](#). There is no cost if you install BlueXP classification on an on-premises system.
- BlueXP classification requires that you have deployed a BlueXP Connector. In many cases you already have a Connector because of other storage and services you are using in BlueXP. The Connector instance results in charges from the cloud provider where it is deployed. See the [type of instance that is deployed for each cloud provider](#). There is no cost if you install the Connector on an on-premises system.

### Data transfer costs


Data transfer costs depend on your setup. If the BlueXP classification instance and data source are in the same Availability Zone and region, then there are no data transfer costs. But if the data source, such as a Cloud Volumes ONTAP system, is in a *different* Availability Zone or region, then you'll be charged by your

cloud provider for data transfer costs. See these links for more details:

- [AWS: Amazon Elastic Compute Cloud \(Amazon EC2\) Pricing](#)
- [Microsoft Azure: Bandwidth Pricing Details](#)
- [Google Cloud: Storage Transfer Service pricing](#)

## The BlueXP classification instance

When you deploy BlueXP classification in the cloud, BlueXP deploys the instance in the same subnet as the Connector. [Learn more about Connectors.](#)



Note the following about the default instance:

- In AWS, BlueXP classification runs on an [m6i.4xlarge instance](#) with a 500 GiB GP2 disk. The operating system image is Amazon Linux 2. When deployed in AWS, you can choose a smaller instance size if you are scanning a small amount of data.
- In Azure, BlueXP classification runs on a [Standard\\_D16s\\_v3 VM](#) with a 500 GiB disk. The operating system image is Ubuntu 22.04.
- In GCP, BlueXP classification runs on an [n2-standard-16 VM](#) with a 500 GiB Standard persistent disk. The operating system image is Ubuntu 22.04.
- In regions where the default instance isn't available, BlueXP classification runs on an alternate instance. [See the alternate instance types.](#)
- The instance is named *CloudCompliance* with a generated hash (UUID) concatenated to it. For example: *CloudCompliance-16bb6564-38ad-4080-9a92-36f5fd2f71c7*
- Only one BlueXP classification instance is deployed per Connector.

You can also deploy BlueXP classification on a Linux host on your premises or on a host in your preferred cloud provider. The software functions exactly the same way regardless of which installation method you choose. Upgrades of BlueXP classification software are automated as long as the instance has internet access.



The instance should remain running at all times because BlueXP classification continuously scans the data.

## Deploy on different instance types

Review the following specifications for instance types:

System size	Specs	Limitations
Extra Large	32 CPUs, 128 GB RAM, 1 TiB SSD	Can scan up to 500 million files.
Large (default)	16 CPUs, 64 GB RAM, 500 GiB SSD	Can scan up to 250 million files.

When deploying BlueXP classification in Azure or GCP, email [ng-contact-data-sense@netapp.com](mailto:ng-contact-data-sense@netapp.com) for assistance if you want to use a smaller instance type.

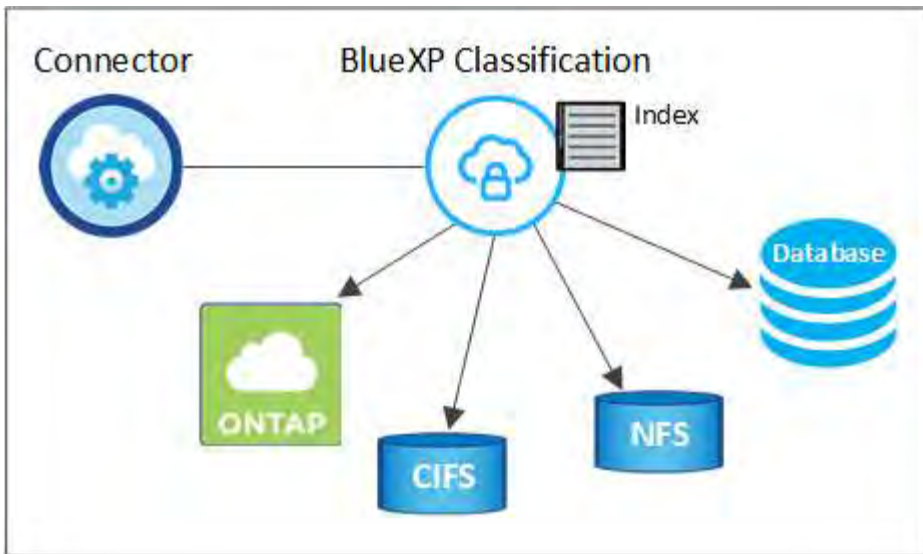
## How BlueXP classification scanning works

At a high-level, BlueXP classification scanning works like this:

1. You deploy an instance of BlueXP classification in BlueXP.
2. You enable high-level mapping (called *Mapping only* scans) or deep-level scanning (called *Map & Classify* scans) on one or more data sources.
3. BlueXP classification scans the data using an AI learning process.
4. You use the provided dashboards and reporting tools to help in your compliance and governance efforts.

After you enable BlueXP classification and select the repositories that you want to scan (these are the volumes, database schemas, or other user data), it immediately starts scanning the data to identify personal and sensitive data. You should focus on scanning live production data in most cases instead of backups, mirrors, or DR sites. Then BlueXP classification maps your organizational data, categorizes each file, and identifies and extracts entities and predefined patterns in the data. The result of the scan is an index of personal information, sensitive personal information, data categories, and file types.

BlueXP classification connects to the data like any other client by mounting NFS and CIFS volumes. NFS volumes are automatically accessed as read-only, while you need to provide Active Directory credentials to scan CIFS volumes.



After the initial scan, BlueXP classification continuously scans your data in a round-robin fashion to detect incremental changes. This is why it's important to keep the instance running.

You can enable and disable scans at the volume level or the database schema level.



BlueXP classification does not impose a limit on the amount of data it can scan. Each Connector supports scanning and displaying 500 TiB of data. To scan more than 500 TiB of data, [install another Connector](#) then [deploy another classification instance](#).

The BlueXP UI displays data from a single connector. For tips on viewing data from multiple Connectors, see [Work with multiple Connectors](#).

## What's the difference between Mapping and Classification scans

You can conduct two types of scans in BlueXP classification:

- **Mapping-only scans** provide only a high-level overview of your data and are performed on selected data sources. Mapping-only scans take less time than map and classify scans because they do not access files to see the data inside. You might want to do this initially to identify areas of research and then perform a Map & Classify scan on those areas.
- **Map & Classify scans** provide deep-level scanning of your data.

For details about the differences between Mapping and Classification scans, see [What's the difference between Mapping and Classification scans?](#)

## Information that BlueXP classification categorizes

BlueXP classification collects, indexes, and assigns categories to the following data:

- **Standard metadata** about files: the file type, its size, creation and modification dates, and so on.
- **Personal data:** Personally identifiable information (PII) such as email addresses, identification numbers, or credit card numbers, which BlueXP classification identifies using specific words, strings, and patterns in the files. [Learn more about personal data](#).
- **Sensitive personal data:** Special types of sensitive personal information (SPII), such as health data, ethnic origin, or political opinions, as defined by General Data Protection Regulation (GDPR) and other privacy regulations. [Learn more about sensitive personal data](#).



- **Categories:** BlueXP classification takes the data that it scanned and divides it into different types of categories. Categories are topics based on AI analysis of the content and metadata of each file. [Learn more about categories.](#)
- **Types:** BlueXP classification takes the data that it scanned and breaks it down by file type. [Learn more about types.](#)
- **Name entity recognition:** BlueXP classification uses AI to extract people's natural names from documents. [Learn about responding to Data Subject Access Requests.](#)

## Networking overview

BlueXP classification deploys a single server, or cluster, wherever you choose — in the cloud or on premises. The servers connect via standard protocols to the data sources and index the findings in an Elasticsearch cluster, which is also deployed on the same servers. This enables support for multi-cloud, cross-cloud, private cloud, and on-premises environments.

BlueXP deploys the BlueXP classification instance with a security group that enables inbound HTTP connections from the Connector instance.

When you use BlueXP in SaaS mode, the connection to BlueXP is served over HTTPS, and the private data sent between your browser and the BlueXP classification instance are secured with end-to-end encryption using TLS 1.2, which means NetApp and third parties can't read it.

Outbound rules are completely open. Internet access is needed to install and upgrade the BlueXP classification software and to send usage metrics.

If you have strict networking requirements, [learn about the endpoints that BlueXP classification contacts.](#)

## User roles in BlueXP classification

The role each user has been assigned provides different capabilities within BlueXP and within BlueXP classification. For details, refer to [BlueXP IAM roles](#) (when using BlueXP in standard mode).

## Access BlueXP classification

You can access the BlueXP classification service through NetApp BlueXP.

To sign in to BlueXP, you can use your NetApp Support Site credentials or you can sign up for a NetApp cloud login using your email and a password. [Learn more about logging in to BlueXP.](#)

Specific tasks require specific BlueXP user roles. [Learn about BlueXP access roles for all services.](#)

### Before you begin

- [You should add a BlueXP Connector.](#)
- [Understand which BlueXP classification deployment style suits your workload.](#)

### Steps

1. In a web browser, navigate to the [BlueXP console](#).

The NetApp BlueXP login page appears.

2. Sign in to BlueXP.

3. From the BlueXP left navigation menu, select **Governance > Classification**.
4. If this is your first time accessing BlueXP classification, the landing page appears.

Select **Deploy Classification On-Premises or Cloud** to begin deploying your classification instance. For more information, see [Which BlueXP classification deployment should you use?](#)

[A screenshot of selecting the button to activate BlueXP classification.]

Otherwise, the BlueXP classification Dashboard appears.

## Deploy BlueXP classification

### Which BlueXP classification deployment should you use?

You can deploy BlueXP classification in different ways. Learn which method meets your needs.

BlueXP classification can be deployed in the following ways:

- [Deploy in the cloud using BlueXP](#). BlueXP will deploy the BlueXP classification instance in the same cloud provider network as the BlueXP Connector.
- [Install on a Linux host with internet access](#). Install BlueXP classification on a Linux host in your network, or on a Linux host in the cloud, that has internet access. This type of installation may be a good option if you prefer to scan on-premises ONTAP systems using a BlueXP classification instance that's also located on premises—but this is not a requirement.
- [Install on a Linux host in an on-premises site without internet access](#), also known as *private mode*. This type of installation, which uses an installation script, has no connectivity to the BlueXP SaaS layer.

Both the installation on a Linux host with internet access and the on-premises installation on a Linux host without internet access use an installation script. The script starts by checking if the system and environment meet the prerequisites. If the prerequisites are met, the installation starts. If you would like to verify the prerequisites independently of running the BlueXP classification installation, there is a separate software package you can download that only tests for the prerequisites.

Refer to [Check that your Linux host is ready to install BlueXP classification](#).

### Deploy BlueXP classification in the cloud using BlueXP

Complete a few steps to deploy BlueXP classification in the cloud. BlueXP will deploy the BlueXP classification instance in the same cloud provider network as the BlueXP Connector.

Note that you can also [install BlueXP classification on a Linux host that has internet access](#). This type of installation may be a good option if you prefer to scan on-premises ONTAP systems using a BlueXP classification instance that's also located on premises — but this is not a requirement. The software functions exactly the same way regardless of which installation method you choose.

#### Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

**1**

### Create a Connector

If you don't already have a Connector, create a Connector now. See [creating a Connector in AWS](#), [creating a Connector in Azure](#), or [creating a Connector in GCP](#).

You can also [install the Connector on-premises](#) on a Linux host in your network or on a Linux host in the cloud.

**2**

### Review prerequisites

Ensure that your environment can meet the prerequisites. This includes outbound internet access for the instance, connectivity between the Connector and BlueXP classification over port 443, and more. [See the complete list](#).

**3**

### Deploy BlueXP classification

Launch the installation wizard to deploy the BlueXP classification instance in the cloud.

## Create a Connector

If you don't already have a Connector, create a Connector in your cloud provider. See [creating a Connector in AWS](#) or [creating a Connector in Azure](#), or [creating a Connector in GCP](#). In most cases you will probably have a Connector set up before you attempt to activate BlueXP classification because most [BlueXP features require a Connector](#), but there are cases where you'll need to set one up now.

There are some scenarios where you have to use a Connector that's deployed in a specific cloud provider:

- When scanning data in Cloud Volumes ONTAP in AWS or Amazon FSx for ONTAP buckets, you use a Connector in AWS.
- When scanning data in Cloud Volumes ONTAP in Azure or in Azure NetApp Files, you use a Connector in Azure.
  - For Azure NetApp Files, it must be deployed in the same region as the volumes you wish to scan.
- When scanning data in Cloud Volumes ONTAP in GCP, you use a Connector in GCP.

On-prem ONTAP systems, NetApp file shares, and databases can be scanned when using any of these cloud Connectors.

Note that you can also [install the Connector on-premises](#) on a Linux host in your network or in the cloud. Some users planning to install BlueXP classification on-prem may also choose to install the Connector on-prem.

As you can see, there may be some situations where you need to use [multiple Connectors](#).



BlueXP classification does not impose a limit on the amount of data it can scan. Each Connector supports scanning and displaying 500 TiB of data. To scan more than 500 TiB of data, [install another Connector](#) then [deploy another classification instance](#).

The BlueXP UI displays data from a single connector. For tips on viewing data from multiple Connectors, see [Work with multiple Connectors](#).

## Government region support

BlueXP classification is supported when the Connector is deployed in a Government region (AWS GovCloud,

Azure Gov, or Azure DoD). When deployed in this manner, BlueXP classification has the following restrictions:

[See more information about deploying the Connector in a Government region.](#)

### **Review prerequisites**

Review the following prerequisites to make sure that you have a supported configuration before you deploy BlueXP classification in the cloud. When you deploy BlueXP classification in the cloud, it's located in the same subnet as the Connector.

### **Enable outbound internet access from BlueXP classification**

BlueXP classification requires outbound internet access. If your virtual or physical network uses a proxy server for internet access, ensure that the BlueXP classification instance has outbound internet access to contact the following endpoints. The proxy must be non-transparent. Transparent proxies are not currently supported.

Review the appropriate table below depending on whether you are deploying BlueXP classification in AWS, Azure, or GCP.

### Required endpoints for AWS

Endpoints	Purpose
<a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a>	Communication with the BlueXP service, which includes NetApp accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://auth0.com">https://auth0.com</a>	Communication with the BlueXP website for centralized user authentication.
<a href="https://cloud-compliance-support-netapp.s3.us-west-2.amazonaws.com">https://cloud-compliance-support-netapp.s3.us-west-2.amazonaws.com</a> <a href="https://hub.docker.com">https://hub.docker.com</a> <a href="https://auth.docker.io">https://auth.docker.io</a> <a href="https://registry-1.docker.io">https://registry-1.docker.io</a> <a href="https://index.docker.io/">https://index.docker.io/</a> <a href="https://dseasb33srrn.cloudfront.net/">https://dseasb33srrn.cloudfront.net/</a> <a href="https://production.cloudflare.docker.com/">https://production.cloudflare.docker.com/</a>	Provides access to software images, manifests, and templates.
<a href="https://kinesis.us-east-1.amazonaws.com">https://kinesis.us-east-1.amazonaws.com</a>	Enables NetApp to stream data from audit records.
<a href="https://cognito-idp.us-east-1.amazonaws.com">https://cognito-idp.us-east-1.amazonaws.com</a> <a href="https://cognito-identity.us-east-1.amazonaws.com">https://cognito-identity.us-east-1.amazonaws.com</a> <a href="https://user-feedback-store-prod.s3.us-west-2.amazonaws.com">https://user-feedback-store-prod.s3.us-west-2.amazonaws.com</a> <a href="https://customer-data-production.s3.us-west-2.amazonaws.com">https://customer-data-production.s3.us-west-2.amazonaws.com</a>	Enables BlueXP classification to access and download manifests and templates, and to send logs and metrics.

### Required endpoints for Azure

Endpoints	Purpose
<a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a>	Communication with the BlueXP service, which includes NetApp accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://auth0.com">https://auth0.com</a>	Communication with the BlueXP website for centralized user authentication.
<a href="https://support.compliance.api.bluexp.netapp.com/">https://support.compliance.api.bluexp.netapp.com/</a> <a href="https://hub.docker.com">https://hub.docker.com</a> <a href="https://auth.docker.io">https://auth.docker.io</a> <a href="https://registry-1.docker.io">https://registry-1.docker.io</a> <a href="https://index.docker.io/">https://index.docker.io/</a> <a href="https://dseasb33srrn.cloudfront.net/">https://dseasb33srrn.cloudfront.net/</a> <a href="https://production.cloudflare.docker.com/">https://production.cloudflare.docker.com/</a>	Provides access to software images, manifests, templates, and to send logs and metrics.
<a href="https://support.compliance.api.bluexp.netapp.com/">https://support.compliance.api.bluexp.netapp.com/</a>	Enables NetApp to stream data from audit records.

### Required endpoints for GCP

Endpoints	Purpose
<a href="https://api.bluexp.netapp.com">https://api.bluexp.netapp.com</a>	Communication with the BlueXP service, which includes NetApp accounts.

Endpoints	Purpose
https://netapp-cloud-account.auth0.com https://auth0.com	Communication with the BlueXP website for centralized user authentication.
https://support.compliance.api.bluexp.netapp.com/ https://hub.docker.com https://auth.docker.io https://registry-1.docker.io https://index.docker.io/ https://dseasb33srrn.cloudfront.net/ https://production.cloudflare.docker.com/	Provides access to software images, manifests, templates, and to send logs and metrics.
https://support.compliance.api.bluexp.netapp.com/	Enables NetApp to stream data from audit records.

### Ensure that BlueXP has the required permissions

Ensure that BlueXP has permissions to deploy resources and create security groups for the BlueXP classification instance. You can find the latest BlueXP permissions in [the policies provided by NetApp](#).

### Ensure that the BlueXP Connector can access BlueXP classification

Ensure connectivity between the Connector and the BlueXP classification instance. The security group for the Connector must allow inbound and outbound traffic over port 443 to and from the BlueXP classification instance. This connection enables deployment of the BlueXP classification instance and enables you to view information in the Compliance and Governance tabs. BlueXP classification is supported in Government regions in AWS and Azure.

Additional inbound and outbound security group rules are required for AWS and AWS GovCloud deployments. See [Rules for the Connector in AWS](#) for details.

Additional inbound and outbound security group rules are required for Azure and Azure Government deployments. See [Rules for the Connector in Azure](#) for details.

### Ensure that you can keep BlueXP classification running

The BlueXP classification instance needs to stay on to continuously scan your data.

### Ensure web browser connectivity to BlueXP classification

After BlueXP classification is enabled, ensure that users access the BlueXP interface from a host that has a connection to the BlueXP classification instance.

The BlueXP classification instance uses a private IP address to ensure that the indexed data isn't accessible to the internet. As a result, the web browser that you use to access BlueXP must have a connection to that private IP address. That connection can come from a direct connection to your cloud provider (for example, a VPN), or from a host that's inside the same network as the BlueXP classification instance.

### Check your vCPU limits

Ensure that your cloud provider's vCPU limit allows for the deployment of an instance with the necessary number of cores. You'll need to verify the vCPU limit for the relevant instance family in the region where BlueXP is running. See [the required instance types](#).

See the following links for more details on vCPU limits:

- [AWS documentation: Amazon EC2 service quotas](#)
- [Azure documentation: Virtual machine vCPU quotas](#)
- [Google Cloud documentation: Resource quotas](#)

### **Deploy BlueXP classification in the cloud**


Follow these steps to deploy an instance of BlueXP classification in the cloud. The Connector will deploy the instance in the cloud, and then install BlueXP classification software on that instance.

In regions where the default instance type isn't available, BlueXP classification runs on an [alternate instance type](#).

## Deploy in AWS

### Steps

1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. Select **Deploy Classification On-Premises or Cloud**.



3. From the *Installation* page, select **Deploy > Deploy** to use the "Large" instance size and start the cloud deployment wizard.
4. The wizard displays progress as it goes through the deployment steps. It will stop and prompt for input if it runs into any issues.



5. When the instance is deployed and BlueXP classification is installed, select **Continue to configuration** to go to the *Configuration* page.

## Deploy in Azure

### Steps

1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. Select **Deploy Classification On-Premises or Cloud**.



**Classification**

## Classify and take control of your data with BlueXP Classification

Driven by powerful artificial intelligence, NetApp's BlueXP Classification gives you control of your data. Map, classify and understand all your cloud and on-premises data to stay secure and compliant, reduce storage costs, and get assistance with data migration projects.

How does it work? [🔗](#)

[Deploy Classification On-Premises or Cloud](#)

**!** You will be prompted to first deploy the BlueXP Connector, which is required for using Classification.

- Multiple Data Sources**  
Cloud and on-premises NetApp storage, databases and more.
- Take Control**  
Map and classify data, take action, set alerts and gain control.
- Safe**  
Data never leaves your network. Agentless solution.
- Now Available at No Cost**  
As part of BlueXP core capability. [Learn more](#)

3. Select **Deploy** to start the cloud deployment wizard.

## Install your Data Sense instance

Select your preferred deployment location:

[Learn more about deploying Data Sense](#)

### Cloud Environment

- I want BlueXP to deploy the instance and install Data Sense**
[Deploy](#)
  - BlueXP will deploy a new machine automatically in the chosen cloud environment.
  - You will be taken to an installation wizard where you can configure your Data Sense installation.
- I deployed an instance and I'm ready to install Data Sense**
[Deploy](#)

### On Premise

- I prepared a local machine and I'm ready to install Data Sense**
[Deploy](#)

4. The wizard displays progress as it goes through the deployment steps. It will stop and prompt for input if it runs into any issues.




5. When the instance is deployed and BlueXP classification is installed, select **Continue to configuration** to go to the *Configuration* page.

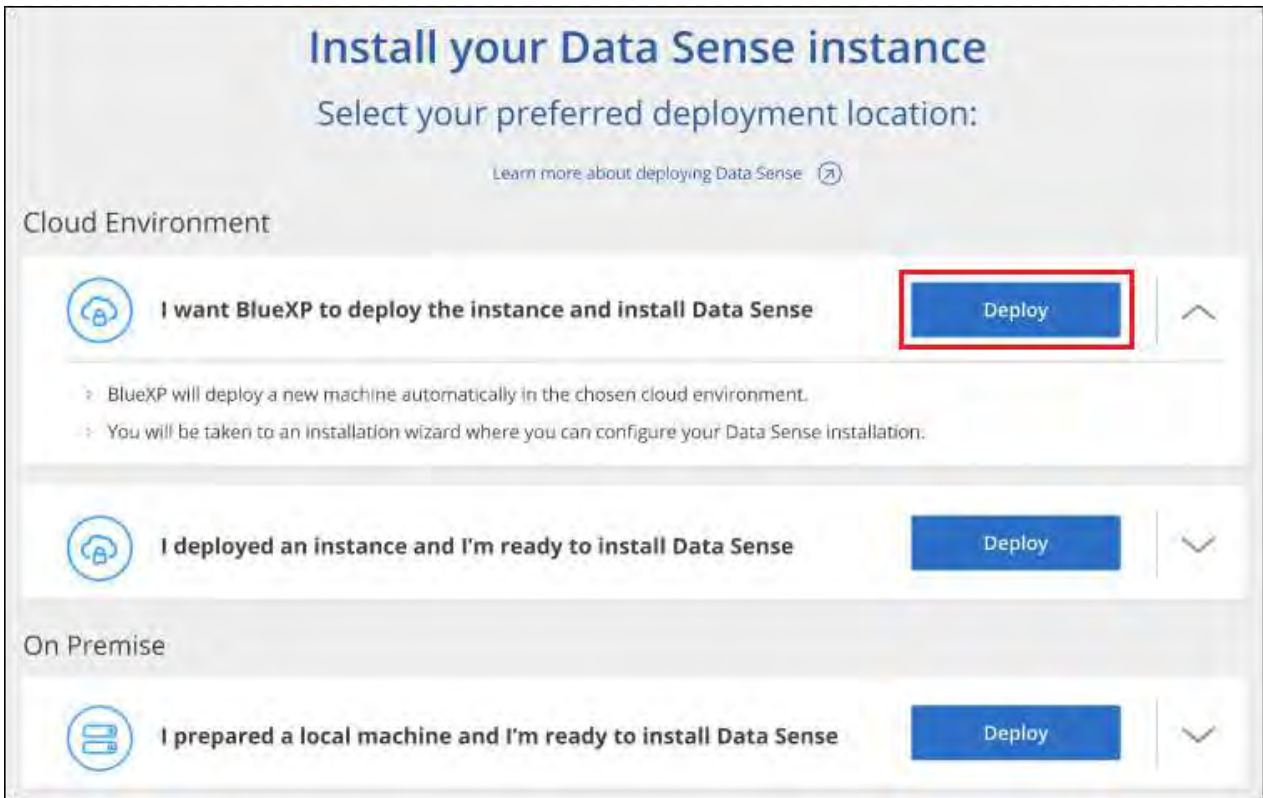
## Deploy in Google Cloud

### Steps

1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. Select **Deploy Classification On-Premises or Cloud**.



3. Select **Deploy** to start the cloud deployment wizard.



4. The wizard displays progress as it goes through the deployment steps. It will stop and prompt for input if it runs into any issues.



5. When the instance is deployed and BlueXP classification is installed, select **Continue to configuration** to go to the *Configuration* page.

## Result

BlueXP deploys the BlueXP classification instance in your cloud provider.

Upgrades to the BlueXP Connector and BlueXP classification software is automated as long as the instances have internet connectivity.

## What's Next

From the Configuration page you can select the data sources that you want to scan.

## Install BlueXP classification on a host that has internet access

Complete a few steps to install BlueXP classification on a Linux host in your network, or on a Linux host in the cloud, that has internet access. You'll need to deploy the Linux host manually in your network or in the cloud as part of this installation.

The on-premises installation may be a good option if you prefer to scan on-premises ONTAP systems using a BlueXP classification instance that's also located on premises — but this is not a requirement. The software functions exactly the same way regardless of which installation method you choose.

The BlueXP classification installation script starts by checking if the system and environment meet the required prerequisites. If the prerequisites are all met, then the installation starts. If you would like to verify the prerequisites independently of running the BlueXP classification installation, there is a separate software package you can download that only tests for the prerequisites. [See how to check if your Linux host is ready to install BlueXP classification.](#)

The typical installation on a Linux host *in your premises* has the following components and connections.

The typical installation on a Linux host *in the cloud* has the following components and connections.



For legacy versions 1.30 and earlier, if you need to install BlueXP classification on multiple hosts, refer to [Install BlueXP classification on multiple hosts with no internet access.](#)

You can also [install BlueXP classification in an on-premises site that doesn't have internet access.](#)

### Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

#### Create a Connector

If you don't already have a Connector, [deploy the Connector on-premises](#) on a Linux host in your network, or on a Linux host in the cloud.

You can also create a Connector with your cloud provider. See [creating a Connector in AWS](#), [creating a Connector in Azure](#), or [creating a Connector in GCP](#).

2

#### Review prerequisites

Ensure that your environment can meet the prerequisites. This includes outbound internet access for the instance, connectivity between the Connector and BlueXP classification over port 443, and more. [See the complete list.](#)

You also need a Linux system that meets the [following requirements](#).

3

#### Download and deploy BlueXP classification

Download the Cloud BlueXP classification software from the NetApp Support Site and copy the installer file to the Linux host you plan to use. Then launch the installation wizard and follow the prompts to deploy the BlueXP classification instance.

## Create a Connector

A BlueXP Connector is required before you can install and use BlueXP classification. In most cases you'll probably have a Connector set up before you attempt to activate BlueXP classification because most [BlueXP features require a Connector](#), but there are cases where you'll need to set one up now.

To create one in your cloud provider environment, see [creating a Connector in AWS](#), [creating a Connector in Azure](#), or [creating a Connector in GCP](#).

There are some scenarios where you have to use a Connector that's deployed in a specific cloud provider:

- When scanning data in Cloud Volumes ONTAP in AWS or Amazon FSx for ONTAP, you use a Connector in AWS.
- When scanning data in Cloud Volumes ONTAP in Azure or in Azure NetApp Files, you use a Connector in Azure.

For Azure NetApp Files, it must be deployed in the same region as the volumes you wish to scan.

- When scanning data in Cloud Volumes ONTAP in GCP, you use a Connector in GCP.

On-prem ONTAP systems, NetApp file shares and database accounts can be scanned using any of these cloud Connectors.

Note that you can also [deploy the Connector on-premises](#) on a Linux host in your network or on a Linux host in the cloud. Some users planning to install BlueXP classification on-prem may also choose to install the Connector on-prem.

You'll need the IP address or host name of the Connector system when installing BlueXP classification. You'll have this information if you installed the Connector in your premises. If the Connector is deployed in the cloud, you can find this information from the BlueXP console: click the Help icon, select **Support**, and click **BlueXP Connector**.

## Prepare the Linux host system

BlueXP classification software must run on a host that meets specific operating system requirements, RAM requirements, software requirements, and so on. The Linux host can be in your network, or in the cloud.

Ensure that you can keep BlueXP classification running. The BlueXP classification machine needs to stay on to continuously scan your data.

- BlueXP classification is not supported on a host that is shared with other applications—the host must be a dedicated host.
- When building the host system in your premises, you can choose among these system sizes depending on the size of the dataset that you plan to have BlueXP classification scan.

System size	CPU	RAM (swap memory must be disabled)	Disk
<b>Extra Large</b>	32 CPUs	128 GB RAM	<ul style="list-style-type: none"> <li>• 1 TiB SSD on /, or 100 GiB available on /opt</li> <li>• 895 GiB available on /var/lib/docker</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>
<b>Large</b>	16 CPUs	64 GB RAM	<ul style="list-style-type: none"> <li>• 500 GiB SSD on /, or 100 GiB available on /opt</li> <li>• 395 GiB available on /var/lib/docker or for Podman /var/lib/containers or for Podman /var/lib/containers</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>

- When deploying a compute instance in the cloud for your BlueXP classification installation, we recommend a system that meets the "Large" system requirements above:
  - **Amazon Elastic Compute Cloud (Amazon EC2) instance type:** We recommend "m6i.4xlarge". [See additional AWS instance types.](#)
  - **Azure VM size:** We recommend "Standard\_D16s\_v3". [See additional Azure instance types.](#)
  - **GCP machine type:** We recommend "n2-standard-16". [See additional GCP instance types.](#)
- **UNIX folder permissions:** The following minimum UNIX permissions are required:

Folder	Minimum Permissions
/tmp	rwXrwxrwt
/opt	rwXr-xr-x
/var/lib/docker	rwX-----
/usr/lib/systemd/system	rwXr-xr-x

- **Operating system:**
  - The following operating systems require using the Docker container engine:
    - Red Hat Enterprise Linux version 7.8 and 7.9
    - Ubuntu 22.04 (requires BlueXP classification version 1.23 or greater)
    - Ubuntu 24.04 (requires BlueXP classification version 1.23 or greater)



- The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater:
  - Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4, and 9.5
- Advanced Vector Extensions (AVX2) must be enabled on the host system.
- **Red Hat Subscription Management:** The host must be registered with Red Hat Subscription Management. If it's not registered, the system can't access repositories to update required 3rd-party software during installation.
- **Additional software:** You must install the following software on the host before you install BlueXP classification:
  - Depending on the OS you are using, you'll need to install one of the container engines:
    - Docker Engine version 19.3.1 or greater. [View installation instructions.](#)
    - Podman version 4 or greater. To install Podman, enter `(sudo yum install podman netavark -y)`.
- Python version 3.6 or greater. [View installation instructions.](#)
  - **NTP considerations:** NetApp recommends configuring the BlueXP classification system to use a Network Time Protocol (NTP) service. The time must be synchronized between the BlueXP classification system and the BlueXP Connector system.
- **Firewalld considerations:** If you are planning to use `firewalld`, we recommend that you enable it before installing BlueXP classification. Run the following commands to configure `firewalld` so that it is compatible with BlueXP classification:

```
firewall-cmd --permanent --add-service=http
firewall-cmd --permanent --add-service=https
firewall-cmd --permanent --add-port=80/tcp
firewall-cmd --permanent --add-port=8080/tcp
firewall-cmd --permanent --add-port=443/tcp
firewall-cmd --reload
```

If you're planning to use additional BlueXP classification hosts as scanner nodes, add these rules to your primary system at this time:

```
firewall-cmd --permanent --add-port=2377/tcp
firewall-cmd --permanent --add-port=7946/udp
firewall-cmd --permanent --add-port=7946/tcp
firewall-cmd --permanent --add-port=4789/udp
```

Note that you must restart Docker or Podman whenever you enable or update `firewalld` settings.



The IP address of the BlueXP classification host system can't be changed after installation.

## Enable outbound internet access from BlueXP classification

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

Endpoints	Purpose
https://api.bluexp.netapp.com	Communication with the BlueXP service, which includes NetApp accounts.
https://netapp-cloud-account.auth0.com https://auth0.com	Communication with the BlueXP website for centralized user authentication.
https://support.compliance.api.bluexp.netapp.com/ https://hub.docker.com https://auth.docker.io https://registry-1.docker.io https://index.docker.io/ https://dseasb33srnrn.cloudfront.net/ https://production.cloudflare.docker.com/	Provides access to software images, manifests, templates, and to send logs and metrics.
https://support.compliance.api.bluexp.netapp.com/	Enables NetApp to stream data from audit records.
https://github.com/docker https://download.docker.com	Provides prerequisite packages for docker installation.
http://packages.ubuntu.com/ http://archive.ubuntu.com	Provides prerequisite packages for Ubuntu installation.

## Verify that all required ports are enabled

You must ensure that all required ports are open for communication between the Connector, BlueXP classification, Active Directory, and your data sources.

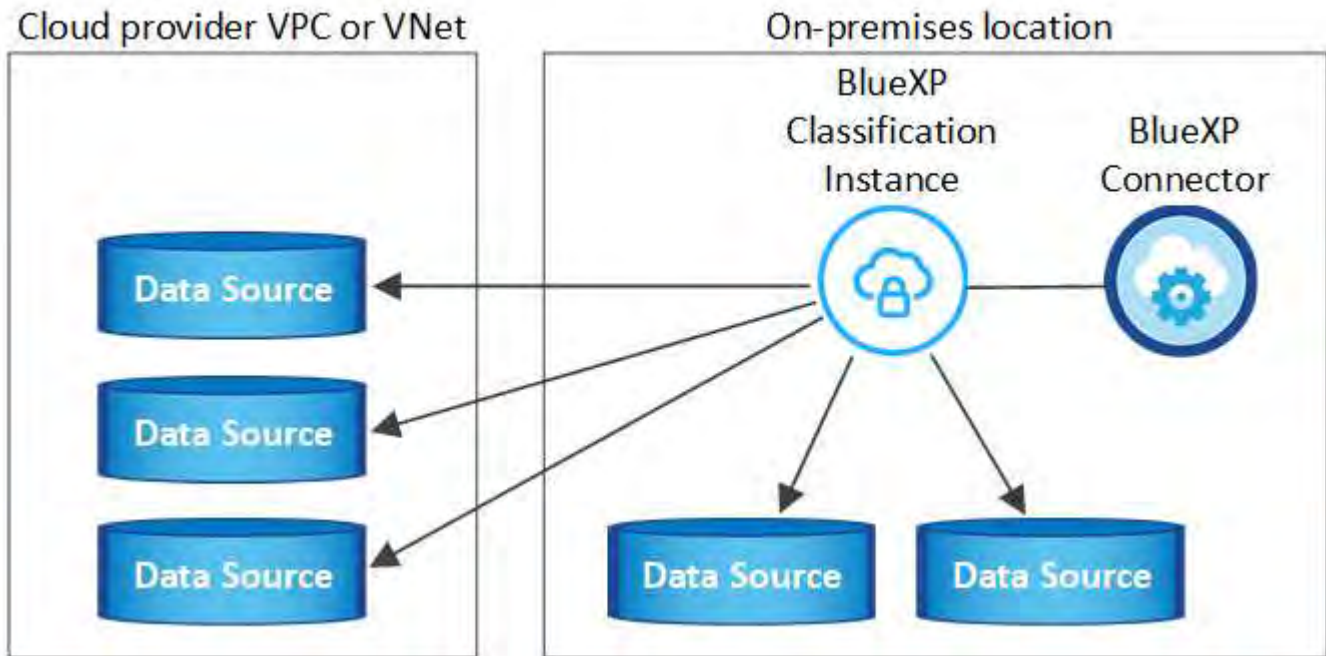
Connection Type	Ports	Description
Connector <> BlueXP classification	8080 (TCP), 443 (TCP), and 80. 9000	<p>The firewall or routing rules for the Connector must allow inbound and outbound traffic over port 443 to and from the BlueXP classification instance.</p> <p>Make sure port 8080 is open so you can see the installation progress in BlueXP.</p> <p>If a firewall is used on the Linux host, port 9000 is required for internal processes within an Ubuntu server.</p>



Connection Type	Ports	Description
Connector <> ONTAP cluster (NAS)	443 (TCP)	<p>BlueXP discovers ONTAP clusters using HTTPS. If you use custom firewall policies, they must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• The Connector host must allow outbound HTTPS access through port 443. If the Connector is in the cloud, all outbound communication is allowed by the predefined firewall or routing rules.</li> <li>• The ONTAP cluster must allow inbound HTTPS access through port 443. The default "mgmt" firewall policy allows inbound HTTPS access from all IP addresses. If you modified this default policy, or if you created your own firewall policy, you must associate the HTTPS protocol with that policy and enable access from the Connector host.</li> </ul>
BlueXP classification <> ONTAP cluster	<ul style="list-style-type: none"> <li>• For NFS - 111 (TCP\UDP) and 2049 (TCP\UDP)</li> <li>• For CIFS - 139 (TCP\UDP) and 445 (TCP\UDP)</li> </ul>	<p>BlueXP classification needs a network connection to each Cloud Volumes ONTAP subnet or on-prem ONTAP system. Firewalls or routing rules for Cloud Volumes ONTAP must allow inbound connections from the BlueXP classification instance.</p> <p>Make sure these ports are open to the BlueXP classification instance:</p> <ul style="list-style-type: none"> <li>• For NFS - 111 and 2049</li> <li>• For CIFS - 139 and 445</li> </ul> <p>NFS volume export policies must allow access from the BlueXP classification instance.</p>
BlueXP classification <> Active Directory	389 (TCP & UDP), 636 (TCP), 3268 (TCP), and 3269 (TCP)	<p>You must have an Active Directory already set up for the users in your company. Additionally, BlueXP classification needs Active Directory credentials to scan CIFS volumes.</p> <p>You must have the information for the Active Directory:</p> <ul style="list-style-type: none"> <li>• DNS Server IP Address, or multiple IP Addresses</li> <li>• User Name and Password for the server</li> <li>• Domain Name (Active Directory Name)</li> <li>• Whether you are using secure LDAP (LDAPS) or not</li> <li>• LDAP Server Port (typically 389 for LDAP, and 636 for secure LDAP)</li> </ul>

## Install BlueXP classification on the Linux host

For typical configurations you'll install the software on a single host system. [See those steps here.](#)



See [Preparing the Linux host system](#) and [Reviewing prerequisites](#) for the full list of requirements before you deploy BlueXP classification.

Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.



BlueXP classification is currently unable to scan S3 buckets, Azure NetApp Files, or FSx for ONTAP when the software is installed on premises. In these cases you'll need to deploy a separate Connector and instance of BlueXP classification in the cloud and [switch between Connectors](#) for your different data sources.

### Single-host installation for typical configurations

Review the requirements and follow these steps when installing BlueXP classification software on a single on-premises host.

[Watch this video](#) to see how to install BlueXP classification.

Note that all installation activities are logged when installing BlueXP classification. If you run into any issues during installation, you can view the contents of the installation audit log. It is written to `/opt/netapp/install_logs/`. [See more details here.](#)

### Before you begin

- Verify that your Linux system meets the [host requirements](#).
- Verify that the system has the two prerequisite software packages installed (Docker Engine or Podman, and Python 3).
- Make sure you have root privileges on the Linux system.
- If you're using a proxy for access to the internet:

- You'll need the proxy server information (IP address or host name, connection port, connection scheme: https or http, user name and password).
  - If the proxy is performing TLS interception, you'll need to know the path on the BlueXP classification Linux system where the TLS CA certificates are stored.
  - The proxy must be non-transparent. BlueXP classification does not currently support transparent proxies.
  - The user must be a local user. Domain users are not supported.
- Verify that your offline environment meets the required [permissions and connectivity](#).

## Steps


1. Download the BlueXP classification software from the [NetApp Support Site](#). The file you should select is named **DATASENSE-INSTALLER-<version>.tar.gz**.
2. Copy the installer file to the Linux host you plan to use (using `scp` or some other method).
3. Unzip the installer file on the host machine, for example:

```
tar -xzf DATASENSE-INSTALLER-V1.25.0.tar.gz
```

4. In BlueXP, select **Governance > Classification**.
5. Select **Deploy Classification On-Premises or Cloud**.

The screenshot shows the 'Classification' page in the BlueXP interface. The main heading is 'Classify and take control of your data with BlueXP Classification'. Below this, there is a brief description: 'Driven by powerful artificial intelligence, NetApp's BlueXP Classification gives you control of your data. Map, classify and understand all your cloud and on-premises data to stay secure and compliant, reduce storage costs, and get assistance with data migration projects.' A 'How does it work?' link is present. A prominent blue button reads 'Deploy Classification On-Premises or Cloud'. Below the button, a note states: 'You will be prompted to first deploy the BlueXP Connector, which is required for using Classification.' The page also features four key benefits with icons: 'Multiple Data Sources' (Cloud and on-premises NetApp storage, databases and more.), 'Take Control' (Map and classify data, take action, set alerts and gain control.), 'Safe' (Data never leaves your network. Agentless solution.), and 'Now Available at No Cost' (As part of BlueXP core capability. Learn more).

6. Depending on whether you are installing BlueXP classification on an instance you prepared in the cloud or on an instance you prepared in your premises, click the appropriate **Deploy** button to start the BlueXP classification installation.



7. The *Deploy Data Sense On Premises* dialog is displayed. Copy the provided command (for example: `sudo ./install.sh -a 12345 -c 27AG75 -t 2198qq`) and paste it in a text file so you can use it later. Then click **Close** to dismiss the dialog.
8. On the host machine, enter the command you copied and then follow a series of prompts, or you can provide the full command including all required parameters as command line arguments.

Note that the installer performs a pre-check to make sure your system and networking requirements are in place for a successful installation. [Watch this video](#) to understand the pre-check messages and implications.

Enter parameters as prompted:	Enter the full command:
<p>1. Paste the command you copied from step 7:  <code>sudo ./install.sh -a &lt;account_id&gt;  -c &lt;client_id&gt; -t &lt;user_token&gt;</code></p> <p>If you are installing on a cloud instance (not on your premises), add <code>--manual-cloud -install &lt;cloud_provider&gt;</code>.</p> <p>2. Enter the IP address or host name of the BlueXP classification host machine so it can be accessed by the Connector system.</p> <p>3. Enter the IP address or host name of the BlueXP Connector host machine so it can be accessed by the BlueXP classification system.</p> <p>4. Enter proxy details as prompted. If your BlueXP Connector already uses a proxy, there is no need to enter this information again here since BlueXP classification will automatically use the proxy used by the Connector.</p>	<p>Alternatively, you can create the whole command in advance, providing the necessary host and proxy parameters:</p> <pre>sudo ./install.sh -a &lt;account_id&gt; -c &lt;client_id&gt; -t &lt;user_token&gt; --host &lt;ds_host&gt; --manager-host &lt;cm_host&gt; --manual-cloud-install &lt;cloud_provider&gt; --proxy-host &lt;proxy_host&gt; --proxy-port &lt;proxy_port&gt; --proxy-scheme &lt;proxy_scheme&gt; --proxy -user &lt;proxy_user&gt; --proxy-password &lt;proxy_password&gt; --cacert-folder-path &lt;ca_cert_dir&gt;</pre>

Variable values:

- *account\_id* = NetApp Account ID
- *client\_id* = Connector Client ID (add the suffix "clients" to the client ID if it not already there)
- *user\_token* = JWT user access token
- *ds\_host* = IP address or host name of the BlueXP classification Linux system.
- *cm\_host* = IP address or host name of the BlueXP Connector system.
- *cloud\_provider* = When installing on a cloud instance, enter "AWS", "Azure", or "Gcp" depending on cloud provider.
- *proxy\_host* = IP or host name of the proxy server if the host is behind a proxy server.
- *proxy\_port* = Port to connect to the proxy server (default 80).
- *proxy\_scheme* = Connection scheme: https or http (default http).
- *proxy\_user* = Authenticated user to connect to the proxy server, if basic authentication is required. The user must be a local user - domain users are not supported.
- *proxy\_password* = Password for the user name that you specified.
- *ca\_cert\_dir* = Path on the BlueXP classification Linux system containing additional TLS CA certificate bundles. Only required if the proxy is performing TLS interception.

## Result

The BlueXP classification installer installs packages, registers the installation, and installs BlueXP classification. Installation can take 10 to 20 minutes.

If there is connectivity over port 8080 between the host machine and the Connector instance, you'll see the installation progress in the BlueXP classification tab in BlueXP.

## What's Next

From the Configuration page you can select the data sources that you want to scan.

## Install BlueXP classification on a Linux host with no internet access

Complete a few steps to install BlueXP classification on a Linux host in an on-premises site that doesn't have internet access - also known as *private mode*. This type of installation, which uses an installation script, has no connectivity to the BlueXP SaaS layer.

[Learn about the different deployment modes for the BlueXP Connector and BlueXP classification.](#)

You can also [deploy BlueXP classification in an on-premises site that has internet access.](#)

The BlueXP classification installation script starts by checking if the system and environment meet the required prerequisites. If the prerequisites are all met, then the installation starts. If you would like to verify the prerequisites independently of running the BlueXP classification installation, there is a separate software package you can download that only tests for the prerequisites. [See how to check if your Linux host is ready to install BlueXP classification.](#)



For legacy versions 1.30 and earlier, if you need to install BlueXP classification on multiple hosts, refer to [Install BlueXP classification on multiple hosts with no internet access.](#)

## Supported data sources

When installed private mode (sometimes called an "offline" or "dark" site), BlueXP classification can only scan data from data sources that are also local to the on-premises site. At this time, BlueXP classification can scan the following **local** data sources:

- On-premises ONTAP systems
- Database schemas

There is no support currently for scanning Cloud Volumes ONTAP, Azure NetApp Files, or FSx for ONTAP accounts when BlueXP classification is deployed in private mode.

## Limitations

Most BlueXP classification features work when it is deployed in a site with no internet access. However, certain features that require internet access are not supported, for example:

- Setting BlueXP roles for different users (for example, Account Admin or Compliance Viewer)
- Copying and synchronizing source files using BlueXP copy and sync
- Automated software upgrades from BlueXP

Both the BlueXP Connector and BlueXP classification will require periodic manual upgrades to enable new features. You can see the BlueXP classification version at the bottom of the BlueXP classification UI pages. Check the [BlueXP classification Release Notes](#) to see the new features in each release and whether you want those features. Then you can follow the steps to [upgrade the BlueXP Connector](#) and [upgrade your BlueXP classification software.](#)

## Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

### Install the BlueXP Connector

If you don't already have a Connector installed in private mode, [deploy the Connector](#) on a Linux host now.

2

### Review BlueXP classification prerequisites

Ensure that your Linux system meets the [host requirements](#), that it has all required software installed, and that your offline environment meets the required [permissions and connectivity](#).

3

### Download and deploy BlueXP classification

Download the BlueXP classification software from the NetApp Support Site and copy the installer file to the Linux host you plan to use. Then launch the installation wizard and follow the prompts to deploy the BlueXP classification instance.

## Install the BlueXP Connector

If you don't already have a BlueXP Connector installed in private mode, [deploy the Connector](#) on a Linux host in your offline site.

## Prepare the Linux host system

BlueXP classification software must run on a host that meets specific operating system requirements, RAM requirements, software requirements, and so on.

- BlueXP classification is not supported on a host that is shared with other applications—the host must be a dedicated host.
- When building the host system in your premises, you can choose among these system sizes depending on the size of the dataset that you plan to have BlueXP classification scan.

System size	CPU	RAM (swap memory must be disabled)	Disk
Extra Large	32 CPUs	128 GB RAM	<ul style="list-style-type: none"><li>• 1 TiB SSD on /, or 100 GiB available on /opt</li><li>• 895 GiB available on /var/lib/docker</li><li>• 5 GiB on /tmp</li><li>• <b>For Podman, 5 GB on /tmp</b></li><li>• <b>For Podman, 30 GB on /var/tmp</b></li></ul>

System size	CPU	RAM (swap memory must be disabled)	Disk
Large	16 CPUs	64 GB RAM	<ul style="list-style-type: none"> <li>• 500 GiB SSD on /, or 100 GiB available on /opt</li> <li>• 395 GiB available on /var/lib/docker or for Podman /var/lib/containers or for Podman /var/lib/containers</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>

- When deploying a compute instance in the cloud for your BlueXP classification installation, we recommend a system that meets the "Large" system requirements above:
  - **Amazon Elastic Compute Cloud (Amazon EC2) instance type:** We recommend "m6i.4xlarge". [See additional AWS instance types.](#)
  - **Azure VM size:** We recommend "Standard\_D16s\_v3". [See additional Azure instance types.](#)
  - **GCP machine type:** We recommend "n2-standard-16". [See additional GCP instance types.](#)
- **UNIX folder permissions:** The following minimum UNIX permissions are required:

Folder	Minimum Permissions
/tmp	rwXrwxrwt
/opt	rwXr-Xr-X
/var/lib/docker	rwX-----
/usr/lib/systemd/system	rwXr-Xr-X

- **Operating system:**
  - The following operating systems require using the Docker container engine:
    - Red Hat Enterprise Linux version 7.8 and 7.9
    - Ubuntu 22.04 (requires BlueXP classification version 1.23 or greater)
    - Ubuntu 24.04 (requires BlueXP classification version 1.23 or greater)
  - The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater:
    - Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4, and 9.5
  - Advanced Vector Extensions (AVX2) must be enabled on the host system.
- **Red Hat Subscription Management:** The host must be registered with Red Hat Subscription Management. If it's not registered, the system can't access repositories to update required 3rd-party software during installation.
- **Additional software:** You must install the following software on the host before you install BlueXP classification:



- Depending on the OS you are using, you'll need to install one of the container engines:
  - Docker Engine version 19.3.1 or greater. [View installation instructions.](#)
  - Podman version 4 or greater. To install Podman, enter `(sudo yum install podman netavark -y)`.
- Python version 3.6 or greater. [View installation instructions.](#)
  - **NTP considerations:** NetApp recommends configuring the BlueXP classification system to use a Network Time Protocol (NTP) service. The time must be synchronized between the BlueXP classification system and the BlueXP Connector system.
- **Firewalld considerations:** If you are planning to use `firewalld`, we recommend that you enable it before installing BlueXP classification. Run the following commands to configure `firewalld` so that it is compatible with BlueXP classification:

```
firewall-cmd --permanent --add-service=http
firewall-cmd --permanent --add-service=https
firewall-cmd --permanent --add-port=80/tcp
firewall-cmd --permanent --add-port=8080/tcp
firewall-cmd --permanent --add-port=443/tcp
firewall-cmd --reload
```

Note that you must restart Docker or Podman whenever you enable or update `firewalld` settings.



The IP address of the BlueXP classification host system can't be changed after installation.

## Verify BlueXP and BlueXP classification prerequisites

Review the following prerequisites to make sure that you have a supported configuration before you deploy BlueXP classification.

- Ensure that the Connector has permissions to deploy resources and create security groups for the BlueXP classification instance. You can find the latest BlueXP permissions in [the policies provided by NetApp](#).
- Ensure that you can keep BlueXP classification running. The BlueXP classification instance needs to stay on to continuously scan your data.
- Ensure web browser connectivity to BlueXP classification. After BlueXP classification is enabled, ensure that users access the BlueXP interface from a host that has a connection to the BlueXP classification instance.

The BlueXP classification instance uses a private IP address to ensure that the indexed data isn't accessible to others. As a result, the web browser that you use to access BlueXP must have a connection to that private IP address. That connection can come from a host that's inside the same network as the BlueXP classification instance.

## Verify that all required ports are enabled

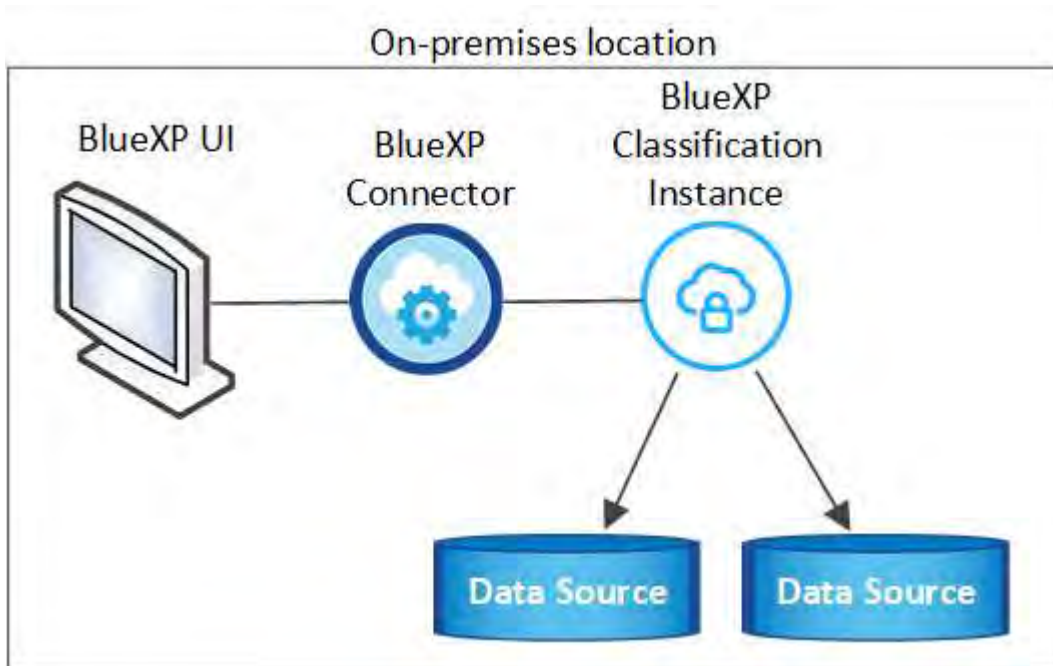
You must ensure that all required ports are open for communication between the Connector, BlueXP classification, Active Directory, and your data sources.

Connection Type	Ports	Description
Connector <> BlueXP classification	8080 (TCP), 6000 (TCP), 443 (TCP), and 80. 9000	<p>The security group for the Connector must allow inbound and outbound traffic over ports 6000 and 443 to and from the BlueXP classification instance.</p> <ul style="list-style-type: none"> <li>• Port 6000 is required so that the BlueXP classification BYOL license works in a dark site.</li> <li>• Port 8080 should be open so you can see the installation progress in BlueXP.</li> <li>• If a firewall is used on the Linux host, port 9000 is required for internal processes within an Ubuntu server.</li> </ul>
Connector <> ONTAP cluster (NAS)	443 (TCP)	<p>BlueXP discovers ONTAP clusters using HTTPS. If you use custom firewall policies, they must meet the following requirements:</p> <ul style="list-style-type: none"> <li>• The Connector host must allow outbound HTTPS access through port 443. If the Connector is in the cloud, all outbound communication is allowed by the predefined security group.</li> <li>• The ONTAP cluster must allow inbound HTTPS access through port 443. The default "mgmt" firewall policy allows inbound HTTPS access from all IP addresses. If you modified this default policy, or if you created your own firewall policy, you must associate the HTTPS protocol with that policy and enable access from the Connector host.</li> </ul>
BlueXP classification <> ONTAP cluster	<ul style="list-style-type: none"> <li>• For NFS - 111 (TCP\UDP) and 2049 (TCP\UDP)</li> <li>• For CIFS - 139 (TCP\UDP) and 445 (TCP\UDP)</li> </ul>	<p>BlueXP classification needs a network connection to each Cloud Volumes ONTAP subnet or on-prem ONTAP system. Security groups for Cloud Volumes ONTAP must allow inbound connections from the BlueXP classification instance.</p> <p>Make sure these ports are open to the BlueXP classification instance:</p> <ul style="list-style-type: none"> <li>• For NFS - 111 and 2049</li> <li>• For CIFS - 139 and 445</li> </ul> <p>NFS volume export policies must allow access from the BlueXP classification instance.</p>

Connection Type	Ports	Description
BlueXP classification <> Active Directory	389 (TCP & UDP), 636 (TCP), 3268 (TCP), and 3269 (TCP)	<p>You must have an Active Directory already set up for the users in your company. Additionally, BlueXP classification needs Active Directory credentials to scan CIFS volumes.</p> <p>You must have the information for the Active Directory:</p> <ul style="list-style-type: none"> <li>• DNS Server IP Address, or multiple IP Addresses</li> <li>• User Name and Password for the server</li> <li>• Domain Name (Active Directory Name)</li> <li>• Whether you are using secure LDAP (LDAPS) or not</li> <li>• LDAP Server Port (typically 389 for LDAP, and 636 for secure LDAP)</li> </ul>
If a firewall used on Linux host	9000	Needed for internal processes within an Ubuntu server.

### Install BlueXP classification on the on-premises Linux host

For typical configurations you'll install the software on a single host system.



### Single-host installation for typical configurations

Follow these steps when installing BlueXP classification software on a single on-premises host in an offline environment.

Note that all installation activities are logged when installing BlueXP classification. If you run into any issues during installation, you can view the contents of the installation audit log. It is written to

/opt/netapp/install\_logs/. [See more details here.](#)

## Before you begin

- Verify that your Linux system meets the [host requirements](#).
- Verify that you have installed the two prerequisite software packages (Docker Engine or Podman, and Python 3).
- Make sure you have root privileges on the Linux system.
- Verify that your offline environment meets the required [permissions and connectivity](#).

## Steps

1. On an internet-configured system, download the BlueXP classification software from the [NetApp Support Site](#). The file you should select is named **DataSense-offline-bundle-<version>.tar.gz**.
2. Copy the installer bundle to the Linux host you plan to use in private mode.
3. Unzip the installer bundle on the host machine, for example:

```
tar -xzf DataSense-offline-bundle-v1.25.0.tar.gz
```

This extracts required software and the actual installation file **cc\_onprem\_installer.tar.gz**.


4. Unzip the installation file on the host machine, for example:

```
tar -xzf cc_onprem_installer.tar.gz
```

5. Launch BlueXP and select **Governance > Classification**.
6. Select **Deploy Classification On-Premises or Cloud**.

The screenshot shows the 'Classification' page in the BlueXP interface. The main heading is 'Classify and take control of your data with BlueXP Classification'. Below this, there is a paragraph describing the capabilities of BlueXP Classification, followed by a 'How does it work?' link. A prominent blue button reads 'Deploy Classification On-Premises or Cloud'. Below the button, a note states: 'You will be prompted to first deploy the BlueXP Connector, which is required for using Classification.' The page also features a dashboard preview with various charts and a 'View All' link. At the bottom, there are four key benefits highlighted with icons: 'Multiple Data Sources' (Cloud and on-premises NetApp storage, databases and more), 'Take Control' (Map and classify data, take action, set alerts and gain control), 'Safe' (Data never leaves your network. Agentless solution), and 'Now Available at No Cost' (As part of BlueXP core capability. Learn more).

7. Click **Deploy** to start the on-prem installation.



8. The *Deploy Data Sense On Premises* dialog is displayed. Copy the provided command (for example: `sudo ./install.sh -a 12345 -c 27AG75 -t 2198qq --darksite`) and paste it in a text file so you can use it later. Then click **Close** to dismiss the dialog.
9. On the host machine, enter the command you copied and then follow a series of prompts, or you can provide the full command including all required parameters as command line arguments.

Note that the installer performs a pre-check to make sure your system and networking requirements are in place for a successful installation.

Enter parameters as prompted:	Enter the full command:
<ol style="list-style-type: none"> <li>1. Paste the information you copied from step 8:  <code>sudo ./install.sh -a &lt;account_id&gt;  -c &lt;client_id&gt; -t &lt;user_token&gt;  --darksite</code> </li> <li>2. Enter the IP address or host name of the BlueXP classification host machine so it can be accessed by the Connector system.</li> <li>3. Enter the IP address or host name of the BlueXP Connector host machine so it can be accessed by the BlueXP classification system.</li> </ol>	<p>Alternatively, you can create the whole command in advance, providing the necessary host parameters:</p> <pre>sudo ./install.sh -a &lt;account_id&gt; -c &lt;client_id&gt; -t &lt;user_token&gt; --host &lt;ds_host&gt; --manager-host &lt;cm_host&gt; --no-proxy --darksite</pre>

Variable values:

- *account\_id* = NetApp Account ID

- *client\_id* = Connector Client ID (add the suffix "clients" to the client ID if it not already there)
- *user\_token* = JWT user access token
- *ds\_host* = IP address or host name of the BlueXP classification system.
- *cm\_host* = IP address or host name of the BlueXP Connector system.

## Result

The BlueXP classification installer installs packages, registers the installation, and installs BlueXP classification. Installation can take 10 to 20 minutes.

If there is connectivity over port 8080 between the host machine and the Connector instance, you'll see the installation progress in the BlueXP classification tab in BlueXP.

## What's Next

From the Configuration page you can select the local [on-prem ONTAP clusters](#) and [databases](#) that you want to scan.

## Upgrade BlueXP classification software

Since BlueXP classification software is updated with new features on a regular basis, you should get into a routine to check for new versions periodically to make sure you're using the newest software and features. You'll need to upgrade BlueXP classification software manually because there's no internet connectivity to perform the upgrade automatically.

### Before you begin

- We recommend that your BlueXP Connector software is upgraded to the newest available version. [See the Connector upgrade steps](#).
- Starting with BlueXP classification version 1.24 you can perform upgrades to any future version of software.

If your BlueXP classification software is running a version prior to 1.24, you can upgrade only one major version at a time. For example, if you have version 1.21.x installed, you can upgrade only to 1.22.x. If you are a few major versions behind, you'll need to upgrade the software multiple times.

### Steps

1. On an internet-configured system, download the BlueXP classification software from the [NetApp Support Site](#). The file you should select is named **DataSense-offline-bundle-<version>.tar.gz**.
2. Copy the software bundle to the Linux host where BlueXP classification is installed in the dark site.
3. Unzip the software bundle on the host machine, for example:

```
tar -xvf DataSense-offline-bundle-v1.25.0.tar.gz
```

This extracts the installation file **cc\_onprem\_installer.tar.gz**.

4. Unzip the installation file on the host machine, for example:

```
tar -xzf cc_onprem_installer.tar.gz
```

This extracts the upgrade script **start\_darksite\_upgrade.sh** and any required third-party software.

5. Run the upgrade script on the host machine, for example:

```
start_darksite_upgrade.sh
```

## Result

The BlueXP classification software is upgraded on your host. The update can take 5 to 10 minutes.

You can verify that the software has been updated by checking the version at the bottom of the BlueXP classification UI pages.

## Check that your Linux host is ready to install BlueXP classification

Before installing BlueXP classification manually on a Linux host, optionally run a script on the host to verify that all the prerequisites are in place for installing BlueXP classification. You can run this script on a Linux host in your network, or on a Linux host in the cloud. The host can be connected to the internet, or the host can reside in a site that doesn't have internet access (a *dark site*).

There is also a prerequisite test script that is part of the BlueXP classification installation script. The script described here is specifically designed for users who want to verify the Linux host independently of running the BlueXP classification installation script.

## Getting Started

You'll perform the following tasks.

1. Optionally, install a BlueXP Connector if you don't already have one installed. You can run the test script without having a Connector installed, but the script checks for connectivity between the Connector and the BlueXP classification host machine - so it is recommended that you have a Connector.
2. Prepare the host machine and verify that it meets all the requirements.
3. Enable outbound internet access from the BlueXP classification host machine.
4. Verify that all required ports are enabled on all systems.
5. Download and run the Prerequisite test script.

## Create a Connector

A BlueXP Connector is required before you can install and use BlueXP classification. You can, however, run the Prerequisites script without a Connector.

You can [install the Connector on-premises](#) on a Linux host in your network or on a Linux host in the cloud. Some users planning to install BlueXP classification on-prem may also choose to install the Connector on-prem.

To create a Connector in your cloud provider environment, see [creating a Connector in AWS](#), [creating a Connector in Azure](#), or [creating a Connector in GCP](#).

You'll need the IP address or host name of the Connector system when running the Prerequisites script. You'll



have this information if you installed the Connector in your premises. If the Connector is deployed in the cloud, you can find this information from the BlueXP console: click the Help icon, select **Support**, and click **BlueXP Connector**.

### Verify host requirements

BlueXP classification software must run on a host that meets specific operating system requirements, RAM requirements, software requirements, and so on.

- BlueXP classification is not supported on a host that is shared with other applications—the host must be a dedicated host.
- When building the host system in your premises, you can choose among these system sizes depending on the size of the dataset that you plan to have BlueXP classification scan.

System size	CPU	RAM (swap memory must be disabled)	Disk
<b>Extra Large</b>	32 CPUs	128 GB RAM	<ul style="list-style-type: none"> <li>• 1 TiB SSD on /, or 100 GiB available on /opt</li> <li>• 895 GiB available on /var/lib/docker</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>
<b>Large</b>	16 CPUs	64 GB RAM	<ul style="list-style-type: none"> <li>• 500 GiB SSD on /, or 100 GiB available on /opt</li> <li>• 395 GiB available on /var/lib/docker or for Podman /var/lib/containers or for Podman /var/lib/containers</li> <li>• 5 GiB on /tmp</li> <li>• <b>For Podman, 5 GB on /tmp</b></li> <li>• <b>For Podman, 30 GB on /var/tmp</b></li> </ul>

- When deploying a compute instance in the cloud for your BlueXP classification installation, we recommend a system that meets the "Large" system requirements above:
  - **Amazon Elastic Compute Cloud (Amazon EC2) instance type:** We recommend "m6i.4xlarge". [See additional AWS instance types.](#)
  - **Azure VM size:** We recommend "Standard\_D16s\_v3". [See additional Azure instance types.](#)
  - **GCP machine type:** We recommend "n2-standard-16". [See additional GCP instance types.](#)
- **UNIX folder permissions:** The following minimum UNIX permissions are required:



Folder	Minimum Permissions
/tmp	rwXrwxrwt
/opt	rwXr-Xr-X
/var/lib/docker	rwX-----
/usr/lib/systemd/system	rwXr-Xr-X

- **Operating system:**

- The following operating systems require using the Docker container engine:
  - Red Hat Enterprise Linux version 7.8 and 7.9
  - Ubuntu 22.04 (requires BlueXP classification version 1.23 or greater)
  - Ubuntu 24.04 (requires BlueXP classification version 1.23 or greater)
- The following operating systems require using the Podman container engine, and they require BlueXP classification version 1.30 or greater:
  - Red Hat Enterprise Linux version 8.8, 8.10, 9.0, 9.1, 9.2, 9.3, 9.4, and 9.5
- Advanced Vector Extensions (AVX2) must be enabled on the host system.

- **Red Hat Subscription Management:** The host must be registered with Red Hat Subscription Management. If it's not registered, the system can't access repositories to update required 3rd-party software during installation.

- **Additional software:** You must install the following software on the host before you install BlueXP classification:

- Depending on the OS you are using, you'll need to install one of the container engines:
  - Docker Engine version 19.3.1 or greater. [View installation instructions.](#)
  - Podman version 4 or greater. To install Podman, enter `(sudo yum install podman netavark -y)`.

- Python version 3.6 or greater. [View installation instructions.](#)

- **NTP considerations:** NetApp recommends configuring the BlueXP classification system to use a Network Time Protocol (NTP) service. The time must be synchronized between the BlueXP classification system and the BlueXP Connector system.

- **Firewalld considerations:** If you are planning to use `firewalld`, we recommend that you enable it before installing BlueXP classification. Run the following commands to configure `firewalld` so that it is compatible with BlueXP classification:

```
firewall-cmd --permanent --add-service=http
firewall-cmd --permanent --add-service=https
firewall-cmd --permanent --add-port=80/tcp
firewall-cmd --permanent --add-port=8080/tcp
firewall-cmd --permanent --add-port=443/tcp
firewall-cmd --reload
```

If you're planning to use additional BlueXP classification hosts as scanner nodes (in a distributed model),

add these rules to your primary system at this time:

```
firewall-cmd --permanent --add-port=2377/tcp
firewall-cmd --permanent --add-port=7946/udp
firewall-cmd --permanent --add-port=7946/tcp
firewall-cmd --permanent --add-port=4789/udp
```

Note that you must restart Docker or Podman whenever you enable or update `firewalld` settings.

## Enable outbound internet access from BlueXP classification

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



This section is not required for host systems installed in sites without internet connectivity.

Endpoints	Purpose
<a href="https://api.bluelxp.netapp.com">https://api.bluelxp.netapp.com</a>	Communication with the BlueXP service, which includes NetApp accounts.
<a href="https://netapp-cloud-account.auth0.com">https://netapp-cloud-account.auth0.com</a> <a href="https://auth0.com">https://auth0.com</a>	Communication with the BlueXP website for centralized user authentication.
<a href="https://support.compliance.api.bluelxp.netapp.com/">https://support.compliance.api.bluelxp.netapp.com/</a> <a href="https://hub.docker.com">https://hub.docker.com</a> <a href="https://auth.docker.io">https://auth.docker.io</a> <a href="https://registry-1.docker.io">https://registry-1.docker.io</a> <a href="https://index.docker.io/">https://index.docker.io/</a> <a href="https://dseasb33srrn.cloudfront.net/">https://dseasb33srrn.cloudfront.net/</a> <a href="https://production.cloudflare.docker.com/">https://production.cloudflare.docker.com/</a>	Provides access to software images, manifests, templates, and to send logs and metrics.
<a href="https://support.compliance.api.bluelxp.netapp.com/">https://support.compliance.api.bluelxp.netapp.com/</a>	Enables NetApp to stream data from audit records.
<a href="https://github.com/docker">https://github.com/docker</a> <a href="https://download.docker.com">https://download.docker.com</a>	Provides prerequisite packages for docker installation.
<a href="http://packages.ubuntu.com/">http://packages.ubuntu.com/</a> <a href="http://archive.ubuntu.com">http://archive.ubuntu.com</a>	Provides prerequisite packages for Ubuntu installation.

## Verify that all required ports are enabled

You must ensure that all required ports are open for communication between the Connector, BlueXP classification, Active Directory, and your data sources.

Connection Type	Ports	Description
Connector <> BlueXP classification	8080 (TCP), 443 (TCP), and 80. 9000	<p>The firewall or routing rules for the Connector must allow inbound and outbound traffic over port 443 to and from the BlueXP classification instance.</p> <p>Make sure port 8080 is open so you can see the installation progress in BlueXP.</p> <p>If a firewall is used on the Linux host, port 9000 is required for internal processes within an Ubuntu server.</p>
Connector <> ONTAP cluster (NAS)	443 (TCP)	BlueXP discovers ONTAP clusters using HTTPS. If you use custom firewall policies, the Connector host must allow outbound HTTPS access through port 443. If the Connector is in the cloud, all outbound communication is allowed by the predefined firewall or routing rules.

### Run the BlueXP classification Prerequisites script

Follow these steps to run the BlueXP classification Prerequisites script.

[Watch this video](#) to see how to run the Prerequisites script and interpret the results.

#### Before you begin

- Verify that your Linux system meets the [host requirements](#).
- Verify that the system has the two prerequisite software packages installed (Docker Engine or Podman, and Python 3).
- Make sure you have root privileges on the Linux system.

#### Steps

1. Download the BlueXP classification Prerequisites script from the [NetApp Support Site](#). The file you should select is named **standalone-pre-requisite-tester-<version>**.
2. Copy the file to the Linux host you plan to use (using `scp` or some other method).
3. Assign permissions to run the script.

```
chmod +x standalone-pre-requisite-tester-v1.25.0
```

4. Run the script using the following command.

```
./standalone-pre-requisite-tester-v1.25.0 <--darksite>
```

Add the option "--darksite" only if you are running the script on a host that doesn't have internet access. Certain prerequisite tests are skipped when the host is not connected to the internet.

5. The script prompts you for the IP address of the BlueXP classification host machine.

- Enter the IP address or host name.
6. The script prompts whether you have an installed BlueXP Connector.
    - Enter **N** if you do not have an installed Connector.
    - Enter **Y** if you do have an installed Connector. And then enter the IP address or host name of the BlueXP Connector so the test script can test this connectivity.
  7. The script runs a variety of tests on the system and it displays results as it progresses. When it finishes it writes a log of the session to a file named `prerequisites-test-<timestamp>.log` in the directory `/opt/netapp/install_logs`.

## Result

If all the prerequisites tests ran successfully, you can install BlueXP classification on the host when you are ready.

If any issues were discovered, they are categorized as "Recommended" or "Required" to be fixed. Recommended issues are typically items that would make the BlueXP classification scanning and categorizing tasks run slower. These items do not need to be corrected - but you may want to address them.

If you have any "Required" issues, you should fix the issues and run the Prerequisites test script again.

# Activate scanning on your data sources

## Scan data sources overview with BlueXP classification

BlueXP classification scans the data in the repositories (the volumes, database schemas, or other user data) that you select to identify personal and sensitive data. BlueXP classification then maps your organizational data, categorizes each file, and identifies predefined patterns in the data. The result of the scan is an index of personal information, sensitive personal information, data categories, and file types.

After the initial scan, BlueXP classification continuously scans your data in a round-robin fashion to detect incremental changes. This is why it's important to keep the instance running.

You can enable and disable scans at the volume level or at the database schema level.

## What's the difference between Mapping and Classification scans

You can conduct two types of scans in BlueXP classification:

- **Mapping-only scans** provide only a high-level overview of your data and are performed on selected data sources. Mapping-only scans take less time than map and classify scans because they do not access files to see the data inside. You might want to do this initially to identify areas of research and then perform a Map & Classify scan on those areas.
- **Map & Classify scans** provide deep-level scanning of your data.

The table below shows some of the differences:

Feature	Map & classify scans	Mapping-only scans
Scan speed	Slow	Fast

Feature	Map & classify scans	Mapping-only scans
Pricing	Free	Free
Capacity	Limited to 500 TiB*	Limited to 500 TiB*
List of file types and used capacity	Yes	Yes
Number of files and used capacity	Yes	Yes
Age and size of files	Yes	Yes
Ability to run a <a href="#">Data Mapping Report</a>	Yes	Yes
Data Investigation page to view file details	Yes	No
Search for names within files	Yes	No
Create <a href="#">saved searches</a> that provide custom search results	Yes	No
Ability to run other reports	Yes	No
Ability to see metadata from files*	No	Yes

\* include::\_include/connector-limit.adoc[]

\*The following metadata is extracted from files during mapping scans:

- Working environment
- Working environment type
- Storage repository
- File type
- Used capacity
- Number of files
- File size
- File creation
- File last access
- File last modified
- File discovered time
- Permissions extraction

**Governance dashboard differences:**

<b>Feature</b>	<b>Map &amp; Classify</b>	<b>Map</b>
Stale data	Yes	Yes
Non-business data	Yes	Yes
Duplicated files	Yes	Yes
Predefined saved searches	Yes	No
Default saved searches	Yes	Yes
DDA report	Yes	Yes
Mapping report	Yes	Yes
Sensitivity level detection	Yes	No
Sensitive data with wide permissions	Yes	No
Open permissions	Yes	Yes
Age of data	Yes	Yes
Size of data	Yes	Yes
Categories	Yes	No
File types	Yes	Yes

**Compliance dashboard differences:**

<b>Feature</b>	<b>Map &amp; Classify</b>	<b>Map</b>
Personal information	Yes	No
Sensitive personal information	Yes	No
Privacy risk assessment report	Yes	No
HIPAA report	Yes	No
PCI DSS report	Yes	No

### Investigation filters differences:

Feature	Map & Classify	Map
Saved searches	Yes	Yes
Working environment type	Yes	Yes
Working environment	Yes	Yes
Storage repository	Yes	Yes
File type	Yes	Yes
File size	Yes	Yes
Created time	Yes	Yes
Discovered time	Yes	Yes
Last modified	Yes	Yes
Last access	Yes	Yes
Open permissions	Yes	Yes
File directory path	Yes	Yes
Category	Yes	No
Sensitivity level	Yes	No
Number of identifiers	Yes	No
Personal data	Yes	No
Sensitive personal data	Yes	No
Data subject	Yes	No
Duplicates	Yes	Yes
Classification status	Yes	Status is always "Limited insights"
Scan analysis event	Yes	Yes
File hash	Yes	Yes
Number of users with access	Yes	Yes
User/group permissions	Yes	Yes
File owner	Yes	Yes
Directory type	Yes	Yes

### How quickly does BlueXP classification scan data

The scan speed is affected by network latency, disk latency, network bandwidth, environment size, and file distribution sizes.

- When performing Mapping-only scans, BlueXP classification can scan between 100-150 TiBs of data per

day.

- When performing Map & classify scans, BlueXP classification can scan between 15-40 TiBs of data per day.

## Scan Azure NetApp Files volumes with BlueXP classification

Complete a few steps to get started with BlueXP classification for Azure NetApp Files.

### Discover the Azure NetApp Files system that you want to scan

If the Azure NetApp Files system you want to scan is not already in BlueXP as a working environment, you can add it to the canvas at this time.

[See how to discover the Azure NetApp Files system in BlueXP.](#)

### Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

BlueXP classification must be deployed in the cloud when scanning Azure NetApp Files volumes, and it must be deployed in the same region as the volumes you wish to scan.

**Note:** Deploying BlueXP classification in an on-premises location is not currently supported when scanning Azure NetApp Files volumes.

### Enable BlueXP classification in your working environments

You can enable BlueXP classification on your Azure NetApp Files volumes.

1. From the BlueXP left navigation menu, click **Governance > Classification**.
2. From the BlueXP classification menu, select **Configuration**.



3. Select how you want to scan the volumes in each working environment. [Learn about mapping and classification scans](#):
  - To map all volumes, select **Map all Volumes**.
  - To map and classify all volumes, select **Map & Classify all Volumes**.
  - To customize scanning for each volume, select **Or select scanning type for each volume**, and then choose the volumes you want to map and/or classify.

See [Enable and disable compliance scans on volumes](#) for details.



4. In the confirmation dialog box, select **Approve** to have BlueXP classification start scanning your volumes.

## Result

BlueXP classification starts scanning the volumes you selected in the working environment. Results will be available in the Compliance dashboard as soon as BlueXP classification finishes the initial scans. The time that it takes depends on the amount of data—it could be a few minutes or hours. You can track the progress of the initial scan by navigating to the **Configuration** menu then selecting the **Working Environment configuration**. The progress of each scan is shown as a progress bar. You can also hover over the progress bar to see the number of files scanned relative to the total files in the volume.



- By default, if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, the system won't scan the files in your volumes because BlueXP classification can't revert the "last access time" to the original timestamp. If you don't care if the last access time is reset, click **Or select scanning type for each volume**. The resulting page has a setting you can enable so that BlueXP classification will scan the volumes regardless of permissions.
- BlueXP classification scans only one file share under a volume. If you have multiple shares in your volumes, you'll need to scan those other shares separately as a shares group. [See more details about this BlueXP classification limitation.](#)

## Verify that BlueXP classification has access to volumes

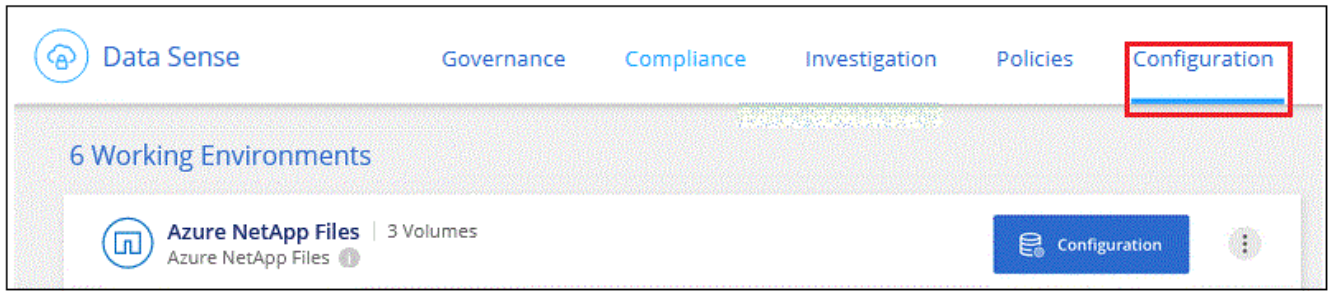
Make sure that BlueXP classification can access volumes by checking your networking, security groups, and export policies. You'll need to provide BlueXP classification with CIFS credentials so it can access CIFS volumes.



For Azure NetApp Files, BlueXP classification can only scan volumes that are in the same region as BlueXP.

## Steps

1. Make sure that there's a network connection between the BlueXP classification instance and each network that includes volumes for Azure NetApp Files.
2. Ensure the following ports are open to the BlueXP classification instance:
  - For NFS – ports 111 and 2049.
  - For CIFS – ports 139 and 445.
3. Ensure that NFS volume export policies include the IP address of the BlueXP classification instance so it can access the data on each volume.
4. If you use CIFS, provide BlueXP classification with Active Directory credentials so it can scan CIFS volumes.
  - a. From the BlueXP left navigation menu, select **Governance > Classification**.
5. From the BlueXP classification menu, select **Configuration**.




- a. For each working environment, select **Edit CIFS Credentials** and enter the user name and password that BlueXP classification needs to access CIFS volumes on the system.

The credentials can be read-only, but providing admin credentials ensures that BlueXP classification can read any data that requires elevated permissions. The credentials are stored on the BlueXP classification instance.


If you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, it's recommended the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, configure the Active Directory user as part of a parent group in the organization which has permissions to all files.

After you enter the credentials, you should see a message that all CIFS volumes were authenticated successfully.



6. On the Configuration page, select **View Details** to review the status for each CIFS and NFS volume and correct any errors.

For example, the following image shows four volumes; one of which BlueXP classification can't scan due to network connectivity issues between the BlueXP classification instance and the volume.



## Enable and disable compliance scans on volumes

You can start or stop mapping-only scans, or mapping and classification scans, in a working environment at any time from the Configuration page. You can also change from mapping-only scans to mapping and classification scans, and vice-versa. We recommend that you scan all volumes.



New volumes added to the working environment are automatically scanned only when you have set the **Map** or **Map & Classify** setting in the heading area. When set to **Custom** or **Off** in the heading area, you'll need to activate mapping and/or full scanning on each new volume you add in the working environment.

The switch at the top of the page for **Scan when missing "write attributes" permissions** is disabled by default. This means that if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, that the system won't scan the files because BlueXP classification can't revert the "last access time" to the original timestamp. If you don't care if the last access time is reset, turn the switch ON and all files are scanned regardless of the permissions. [Learn more](#).

Scan	Storage Repository (Volume)	Type	Status	Required Action
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	AdiNFSVoL_copy	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	AdiProtest2501	NFS	Continuously Scanning	
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	AlexTest	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	AlexTestSecond	NFS	Not Scanning	

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. Do one of the following:
  - To enable mapping-only scans on a volume, in the volume area, select **Map**. To enable on all volumes, in the heading area, select **Map**.
  - To enable full scanning on a volume, in the volume area, select **Map & Classify**. To enable on all volumes, in the heading area, select **Map & Classify**.
  - To disable scanning on a volume, in the volume area, select **Off**. To disable scanning on all volumes, in the heading area, select **Off**.

## Scan Amazon FSx for ONTAP volumes with BlueXP classification

Complete a few steps to get started scanning Amazon FSx for ONTAP volume with BlueXP classification.

### Before you begin

- You need an active Connector in AWS to deploy and manage BlueXP classification.
- The security group you selected when creating the working environment must allow traffic from the BlueXP

classification instance. You can find the associated security group using the ENI connected to the FSx for ONTAP file system and edit it using the AWS Management Console.

[AWS security groups for Linux instances](#)

[AWS security groups for Windows instances](#)

[AWS elastic network interfaces \(ENI\)](#)

- Ensure the following ports are open to the BlueXP classification instance:
  - For NFS – ports 111 and 2049.
  - For CIFS – ports 139 and 445.

## Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

You should deploy BlueXP classification in the same AWS network as the Connector for AWS and the FSx volumes you wish to scan.

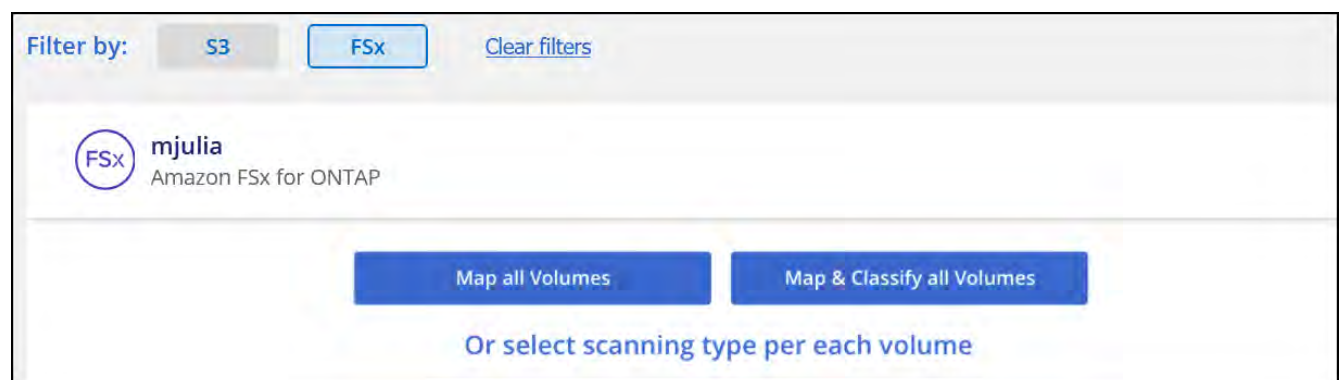
**Note:** Deploying BlueXP classification in an on-premises location is not currently supported when scanning FSx volumes.

Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

## Enable BlueXP classification in your working environments

You can enable BlueXP classification for FSx for ONTAP volumes.

1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. From the BlueXP classification menu, select **Configuration**.



3. Select how you want to scan the volumes in each working environment. [Learn about mapping and classification scans](#):
  - To map all volumes, click **Map all Volumes**.
  - To map and classify all volumes, click **Map & Classify all Volumes**.
  - To customize scanning for each volume, click **Or select scanning type for each volume**, and then choose the volumes you want to map and/or classify.
4. In the confirmation dialog box, click **Approve** to have BlueXP classification start scanning your volumes.

## Result

BlueXP classification starts scanning the volumes you selected in the working environment. Results will be available in the Compliance dashboard as soon as BlueXP classification finishes the initial scans. The time that it takes depends on the amount of data—it could be a few minutes or hours. You can track the progress of the initial scan by navigating to the **Configuration** menu then selecting the **Working Environment configuration**. The progress of each scan is shown as a progress bar. You can also hover over the progress bar to see the number of files scanned relative to the total files in the volume.



- By default, if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, the system won't scan the files in your volumes because BlueXP classification can't revert the "last access time" to the original timestamp. If you don't care if the last access time is reset, click **Or select scanning type for each volume**. The resulting page has a setting you can enable so that BlueXP classification will scan the volumes regardless of permissions.
- BlueXP classification scans only one file share under a volume. If you have multiple shares in your volumes, you'll need to scan those other shares separately as a shares group. [See more details about this BlueXP classification limitation](#).

## Verify that BlueXP classification has access to volumes

Make sure BlueXP classification can access volumes by checking your networking, security groups, and export policies.

You'll need to provide BlueXP classification with CIFS credentials so it can access CIFS volumes.

## Steps

1. From the BlueXP classification menu, select **Configuration**.
2. On the Configuration page, select **View Details** to review the status and correct any errors.

For example, the following image shows a volume BlueXP classification can't scan due to network connectivity issues between the BlueXP classification instance and the volume.

Scan	Storage Repository (Volume)	Type	Status	Required Action
Off	jrmclone	NFS	No Access	Check network connectivity between the Data Sense ...

3. Make sure there's a network connection between the BlueXP classification instance and each network that includes volumes for FSx for ONTAP.



For FSx for ONTAP, BlueXP classification can scan volumes only in the same region as BlueXP.

4. Ensure NFS volume export policies include the IP address of the BlueXP classification instance so it can access the data on each volume.
5. If you use CIFS, provide BlueXP classification with Active Directory credentials so it can scan CIFS volumes.
  - a. From the BlueXP classification menu, select **Configuration**.
  - b. For each working environment, select **Edit CIFS Credentials** and enter the user name and password that BlueXP classification needs to access CIFS volumes on the system.

The credentials can be read-only, but providing admin credentials ensures that BlueXP classification



can read any data that requires elevated permissions. The credentials are stored on the BlueXP classification instance.

If you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, it's recommended the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, configure the Active Directory user as part of a parent group in the organization which has permissions to all files.

After you enter the credentials, you should see a message that all CIFS volumes were authenticated successfully.

## Enable and disable compliance scans on volumes

You can start or stop mapping-only scans, or mapping and classification scans, in a working environment at any time from the Configuration page. You can also change from mapping-only scans to mapping and classification scans, and vice-versa. We recommend that you scan all volumes.

The switch at the top of the page for **Scan when missing "write attributes" permissions** is disabled by default. This means that if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, that the system won't scan the files because BlueXP classification can't revert the "last access time" to the original timestamp. If you don't care if the last access time is reset, turn the switch ON and all files are scanned regardless of the permissions. [Learn more.](#)

Scan	Storage Repository (Volume)	Type	Status	Required Action
Off   Map   <b>Map &amp; Classify</b>	AdiNFSVol_copy	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
Off   Map   <b>Map &amp; Classify</b>	AdiProtest2501	NFS	Continuously Scanning	
Off   <b>Map</b>   Map & Classify	AlexTest	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
<b>Off</b>   Map   Map & Classify	AlexTestSecond	NFS	Not Scanning	

1. From the BlueXP classification menu, select **Configuration**.
2. In the Configuration page, locate the working environment with the volumes you want to scan.
3. Do one of the following:
  - To enable mapping-only scans on a volume, in the volume area, select **Map**. Or, to enable on all volumes, in the heading area, select **Map**.  
To enable full scanning on a volume, in the volume area, select **Map & Classify**. Or, to enable on all volumes, in the heading area, select **Map & Classify**.
  - To disable scanning on a volume, in the volume area, select **Off**. To disable scanning on all volumes, in the heading area, select **Off**.




New volumes added to the working environment are automatically scanned only when you have set the **Map** or **Map & Classify** setting in the heading area. When set to **Custom** or **Off** in the heading area, you'll need to activate mapping and/or full scanning on each new volume you add in the working environment.

## Scan data protection volumes

By default, data protection (DP) volumes are not scanned because they are not exposed externally and BlueXP classification cannot access them. These are the destination volumes for SnapMirror operations from an FSx for ONTAP file system.

Initially, the volume list identifies these volumes as *Type DP* with the *Status Not Scanning* and the *Required Action Enable Access to DP volumes*.



## Steps

If you want to scan these data protection volumes:

1. From the BlueXP classification menu, select **Configuration**.
2. Select **Enable Access to DP volumes** at the top of the page.
3. Review the confirmation message and select **Enable Access to DP volumes** again.
  - Volumes that were initially created as NFS volumes in the source FSx for ONTAP file system are enabled.
  - Volumes that were initially created as CIFS volumes in the source FSx for ONTAP file system require that you enter CIFS credentials to scan those DP volumes. If you already entered Active Directory credentials so that BlueXP classification can scan CIFS volumes you can use those credentials, or you can specify a different set of Admin credentials.

The screenshot shows the 'Provide Active Directory Credentials' dialog box. The 'Use existing CIFS Scanning Credentials (user1@domain2)' radio button is selected and highlighted with a red box. Below it are fields for 'Active Directory Domain' and 'DNS IP Address'. At the bottom, there are 'Enable Access to DP Volumes' and 'Cancel' buttons.

The screenshot shows the 'Provide Active Directory Credentials' dialog box. The 'Use Custom Credentials' radio button is selected and highlighted with a red box. Below it are fields for 'Username', 'Password', 'Active Directory Domain', and 'DNS IP Address'. At the bottom, there are 'Enable Access to DP Volumes' and 'Cancel' buttons.

4. Activate each DP volume that you want to scan.

## Result

Once enabled, BlueXP classification creates an NFS share from each DP volume that was activated for

scanning. The share export policies only allow access from the BlueXP classification instance.

If you had no CIFS data protection volumes when you initially enabled access to DP volumes, and later add some, the button **Enable Access to CIFS DP** appears at the top of the Configuration page. Select this button and add CIFS credentials to enable access to these CIFS DP volumes.



Active Directory credentials are registered only in the storage VM of the first CIFS DP volume, so all DP volumes on that SVM will be scanned. Any volumes that reside on other SVMs will not have the Active Directory credentials registered, so those DP volumes won't be scanned.

## Scan Cloud Volumes ONTAP and on-premises ONTAP volumes with BlueXP classification

Complete a few steps to start scanning your Cloud Volumes ONTAP and on-premises ONTAP volumes using BlueXP classification.

### Prerequisites

Before you enable BlueXP classification, make sure you have a supported configuration.

- If you are scanning Cloud Volumes ONTAP and on-premises ONTAP systems that are accessible over the internet, you can [deploy BlueXP classification in the cloud](#) or [in an on-premises location that has internet access](#).
- If you are scanning on-premises ONTAP systems that have been installed in a dark site that has no internet access, you need to [deploy BlueXP classification in the same on-premises location that has no internet access](#). This also requires that the BlueXP Connector is deployed in that same on-premises location.


### Enable BlueXP classification scanning in your working environments

You can enable BlueXP classification scanning on Cloud Volumes ONTAP systems in any supported cloud provider, and on on-premises ONTAP clusters.

### Steps

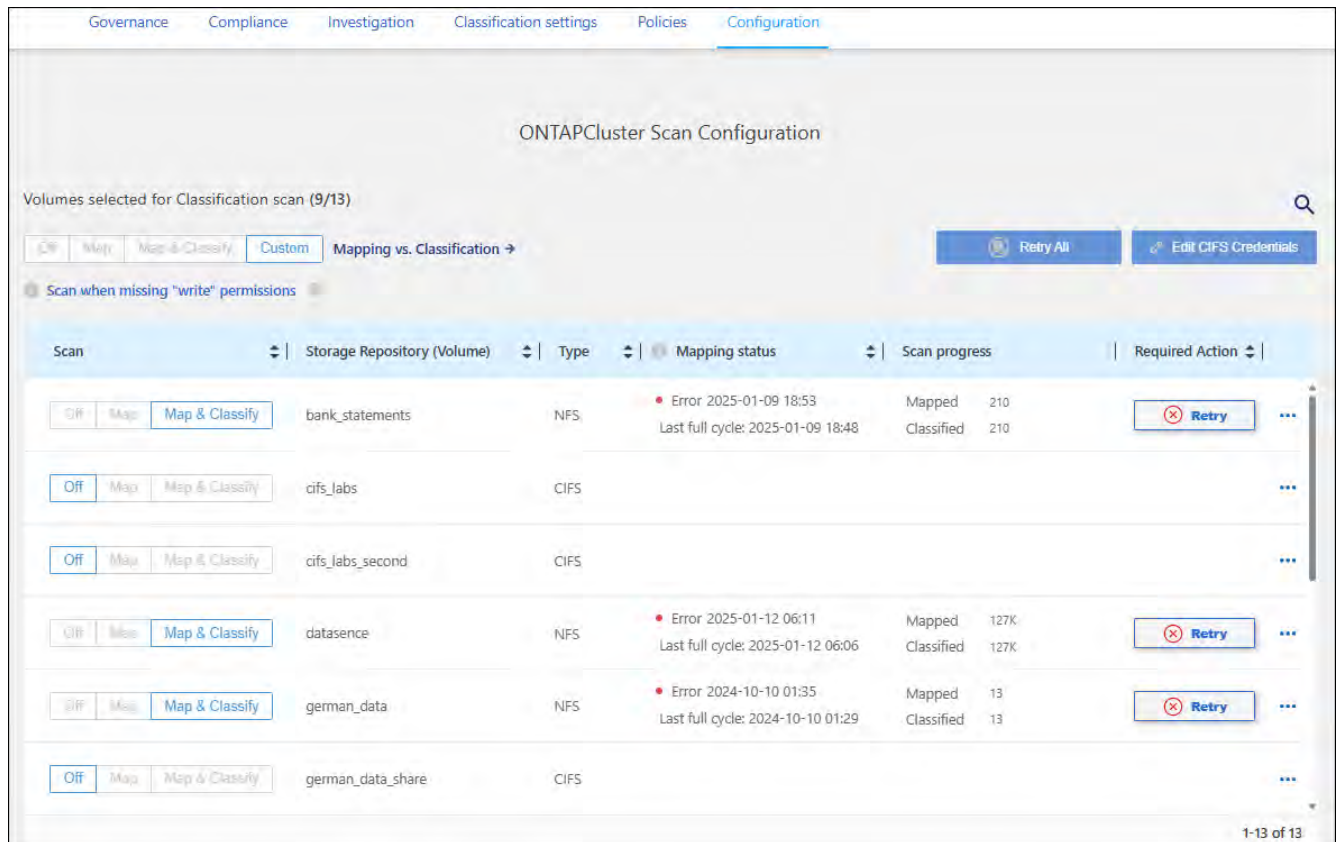
1. From the BlueXP left navigation menu, select **Governance > Classification**.
2. From the BlueXP classification menu, select **Configuration**.

The Configuration page shows multiple working environments.



3. Choose a working environment and select **Configuration**.





- If you don't care if the last access time is reset, turn the **Scan when missing "write attributes" permissions** switch ON and all files are scanned regardless of the permissions.

The switch at the top of the page for **Scan when missing "write attributes" permissions** is disabled by default. This means that if BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, that the system won't classify the files because BlueXP classification can't revert the "last access time" to the original timestamp. [Learn more.](#)

- Select how you want to scan the volumes in each working environment. [Learn about mapping and classification scans:](#)
  - To map all volumes, select **Map**.
  - To map and classify all volumes, select **Map & Classify**.
  - To customize scanning for each volume, select **Custom**, and then choose the volumes you want to map and/or classify.
- In the confirmation dialog box, select **Approve** to have BlueXP classification start scanning your volumes.

## Result

BlueXP classification starts scanning the volumes you selected in the working environment. Results start to appear in the Compliance dashboard as soon as BlueXP classification starts the scan. The time that it takes to complete depends on the amount of data—it could be a few minutes or hours.



BlueXP classification scans only one file share under a volume. If you have multiple shares in your volumes, you'll need to scan those other shares separately as a shares group. [See more details about this BlueXP classification limitation.](#)

## Verify that BlueXP classification has access to volumes

Make sure that BlueXP classification can access volumes by checking your networking, security groups, and export policies. You'll need to provide BlueXP classification with CIFS credentials so it can access CIFS volumes.

### Steps

1. Make sure that there's a network connection between the BlueXP classification instance and each network that includes volumes for Cloud Volumes ONTAP or on-prem ONTAP clusters.
2. Ensure that the security group for Cloud Volumes ONTAP allows inbound traffic from the BlueXP classification instance.

You can either open the security group for traffic from the IP address of the BlueXP classification instance, or you can open the security group for all traffic from inside the virtual network.

3. Ensure that NFS volume export policies include the IP address of the BlueXP classification instance so it can access the data on each volume.
4. If you use CIFS, provide BlueXP classification with Active Directory credentials so it can scan CIFS volumes.
  - a. From the BlueXP left navigation menu, select **Governance > Classification**.
  - b. From the BlueXP classification menu, select **Configuration**.

The screenshot displays the 'ONTAPCluster Scan Configuration' page in the BlueXP interface. The page shows a table of volumes selected for classification scan (9/13). The table has columns for 'Scan', 'Storage Repository (Volume)', 'Type', 'Mapping status', 'Scan progress', and 'Required Action'. The 'Scan' column contains buttons for 'Off', 'Map', and 'Map & Classify'. The 'Mapping status' column shows error messages for some volumes, such as 'Error 2025-01-09 18:53' for 'bank\_statements' and 'Error 2025-01-12 06:11' for 'datasence'. The 'Scan progress' column shows 'Mapped' and 'Classified' counts. The 'Required Action' column contains 'Retry' buttons. The page also includes a search bar, a 'Retry All' button, and an 'Edit CIFS Credentials' button.

Scan	Storage Repository (Volume)	Type	Mapping status	Scan progress	Required Action
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	bank_statements	NFS	<span style="color: red;">•</span> Error 2025-01-09 18:53 Last full cycle: 2025-01-09 18:48	Mapped 210 Classified 210	<input type="button" value="Retry"/> <input type="button" value="More"/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	cifs_labs	CIFS			<input type="button" value="More"/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	cifs_labs_second	CIFS			<input type="button" value="More"/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	datasence	NFS	<span style="color: red;">•</span> Error 2025-01-12 06:11 Last full cycle: 2025-01-12 06:06	Mapped 127K Classified 127K	<input type="button" value="Retry"/> <input type="button" value="More"/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	german_data	NFS	<span style="color: red;">•</span> Error 2024-10-10 01:35 Last full cycle: 2024-10-10 01:29	Mapped 13 Classified 13	<input type="button" value="Retry"/> <input type="button" value="More"/>
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map &amp; Classify"/>	german_data_share	CIFS			<input type="button" value="More"/>

- c. For each working environment, select **Edit CIFS Credentials** and enter the user name and password that BlueXP classification needs to access CIFS volumes on the system.

The credentials can be read-only, but providing admin credentials ensures that BlueXP classification can read any data that requires elevated permissions. The credentials are stored on the BlueXP classification instance.

If you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, it's recommended the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, configure the Active Directory user as part of a parent group in the organization which has permissions to all files.

After you enter the credentials, you should see a message that all CIFS volumes were authenticated successfully.

5. On the Configuration page, select **Configuration** to review the status for each CIFS and NFS volume and correct any errors.

### Disable compliance scans on volumes

You can start or stop mapping-only scans, or mapping and classification scans, in a working environment at any time from the Configuration page. You can also change from mapping-only scans to mapping and classification scans, and vice-versa. We recommend that you scan all volumes.



New volumes added to the working environment are automatically scanned only when you have set the **Map** or **Map & Classify** setting in the heading area. When the option is set to **Custom** or **Off** in the heading area, you'll need to activate mapping and/or full scanning on each new volume you add in the working environment.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. Select the **Configuration** button for the working environment that you want to change.

The screenshot shows the 'ONTAPCluster Scan Configuration' page. At the top, there are navigation tabs: Governance, Compliance, Investigation, Classification settings, Policies, and Configuration. Below the tabs, the page title is 'ONTAPCluster Scan Configuration'. Underneath, it says 'Volumes selected for Classification scan (9/13)'. There are several controls: a search icon, a 'Mapping vs. Classification' dropdown menu with options 'Off', 'Map', 'Map & Classify', and 'Custom', and buttons for 'Retry All' and 'Edit CIFS Credentials'. A checkbox labeled 'Scan when missing "write" permissions' is also visible. The main part of the page is a table with columns: Scan, Storage Repository (Volume), Type, Mapping status, Scan progress, and Required Action. The table lists several volumes, some with error messages in the 'Mapping status' column and 'Retry' buttons in the 'Required Action' column.

Scan	Storage Repository (Volume)	Type	Mapping status	Scan progress	Required Action
Off   Map   Map & Classify	bank_statements	NFS	Error 2025-01-09 18:53 Last full cycle: 2025-01-09 18:48	Mapped 210 Classified 210	Retry
Off   Map   Map & Classify	cifs_labs	CIFS			
Off   Map   Map & Classify	cifs_labs_second	CIFS			
Off   Map   Map & Classify	datasense	NFS	Error 2025-01-12 06:11 Last full cycle: 2025-01-12 06:06	Mapped 127K Classified 127K	Retry
Off   Map   Map & Classify	german_data	NFS	Error 2024-10-10 01:35 Last full cycle: 2024-10-10 01:29	Mapped 13 Classified 13	Retry
Off   Map   Map & Classify	german_data_share	CIFS			

3. Do one of the following:

- To disable scanning on a volume, in the volume area, select **Off**.
- To disable scanning on all volumes, in the heading area, select **Off**.

## Scan database schemas with BlueXP classification

Complete a few steps to start scanning your database schemas with BlueXP classification.

### Review prerequisites

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.

### Supported databases

BlueXP classification can scan schemas from the following databases:

- Amazon Relational Database Service (Amazon RDS)
- MongoDB
- MySQL
- Oracle
- PostgreSQL
- SAP HANA
- SQL Server (MSSQL)



The statistics gathering feature **must be enabled** in the database.

### Database requirements

Any database with connectivity to the BlueXP classification instance can be scanned, regardless of where it is hosted. You just need the following information to connect to the database:

- IP Address or host name
- Port
- Service name (only for accessing Oracle databases)
- Credentials that allow read access to the schemas

When choosing a user name and password, it's important to choose one that has full read permissions to all the schemas and tables you want to scan. We recommend that you create a dedicated user for the BlueXP classification system with all the required permissions.

**Note:** For MongoDB, a read-only Admin role is required.

### Deploy the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

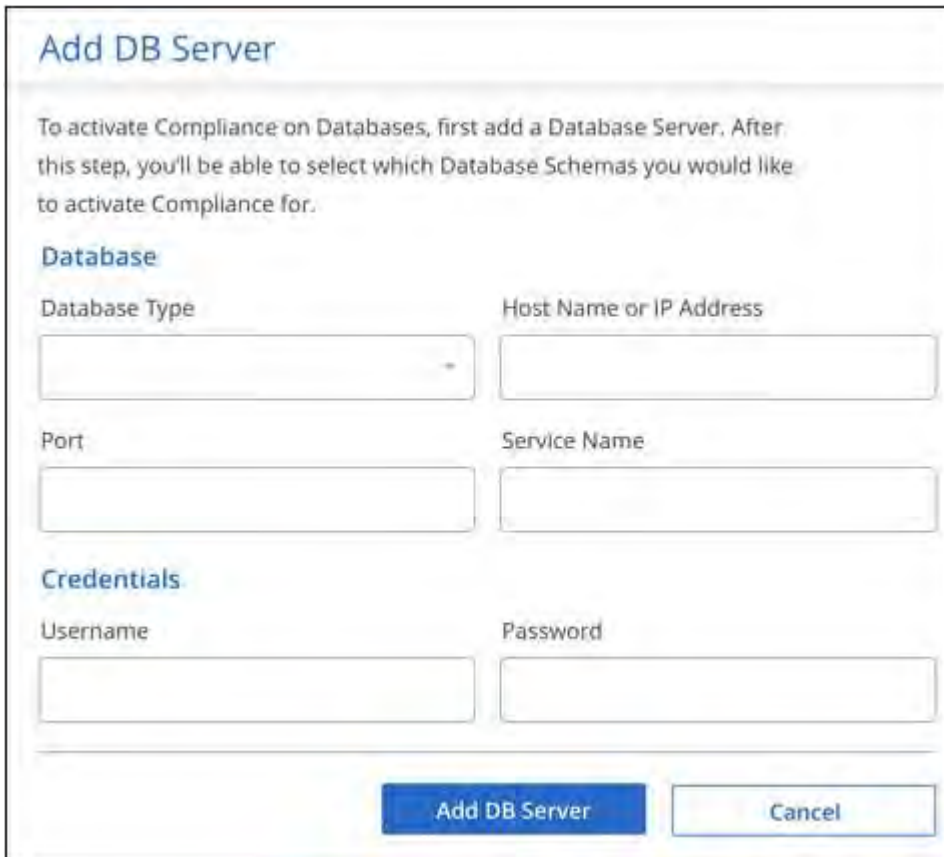
If you are scanning database schemas that are accessible over the internet, you can [deploy BlueXP classification in the cloud](#) or [deploy BlueXP classification in an on-premises location that has internet access](#).

If you are scanning database schemas that have been installed in a dark site that has no internet access, you need to [deploy BlueXP classification in the same on-premises location that has no internet access](#). This also requires that the BlueXP Connector is deployed in that same on-premises location.

### Add the database server

Add the database server where the schemas reside.

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration page, select **Add Working Environment > Add Database Server**.
3. Enter the required information to identify the database server.
  - a. Select the database type.
  - b. Enter the port and the host name or IP address to connect to the database.
  - c. For Oracle databases, enter the Service name.
  - d. Enter the credentials so that BlueXP classification can access the server.
  - e. Click **Add DB Server**.



The screenshot shows a web form titled "Add DB Server". At the top, there is a blue header with the title. Below the header is a paragraph of instructional text: "To activate Compliance on Databases, first add a Database Server. After this step, you'll be able to select which Database Schemas you would like to activate Compliance for." The form is divided into three sections: "Database", "Credentials", and "Add DB Server" buttons. The "Database" section contains four input fields: "Database Type" (a dropdown menu), "Host Name or IP Address", "Port", and "Service Name". The "Credentials" section contains two input fields: "Username" and "Password". At the bottom right, there are two buttons: "Add DB Server" (a blue button) and "Cancel" (a white button with a blue border).

The database is added to the list of working environments.


### Enable and disable compliance scans on database schemas

You can stop or start full scanning of your schemas at any time.

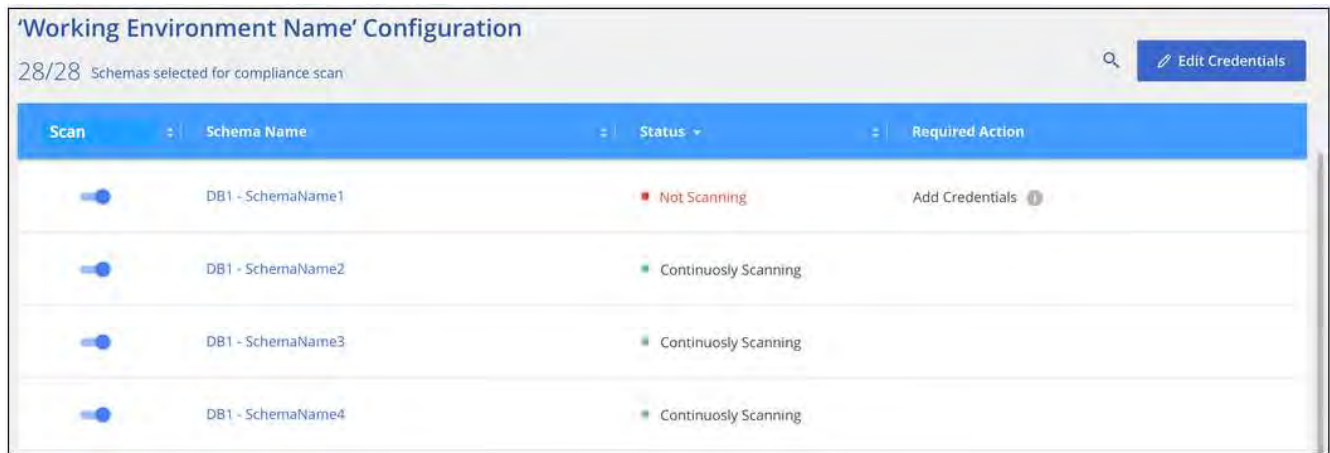


There is no option to select mapping-only scans for database schemas.

1. From the Configuration page, select the **Configuration** button for the database you want to configure.



2. Select the schemas that you want to scan by moving the slider to the right.



## Result

BlueXP classification starts scanning the database schemas that you enabled. You can track the progress of the initial scan by navigating to the **Configuration** menu then selecting the **Working Environment configuration**. The progress of each scan is shown as a progress bar. You can also hover over the progress bar to see the number of files scanned relative to the total files in the volume. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

BlueXP classification scans your databases once per day; databases are not continuously scanned like other data sources.

## Scan file shares with BlueXP classification

To scan file shares, you must first create a file shares group in BlueXP classification. File shares groups are for NFS or CIFS (SMB) shares hosted on-premises or in the cloud.



Scanning data from non-NetApp file shares is not supported in the BlueXP classification core version.

## Prerequisites

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.



- The shares can be hosted anywhere, including in the cloud or on-premises. CIFS shares from older NetApp 7-Mode storage systems can be scanned as file shares.
  - BlueXP classification can't extract permissions or the "last access time" from 7-Mode systems.
  - Because of a known issue between some Linux versions and CIFS shares on 7-Mode systems, you must configure the share to use only SMBv1 with NTLM authentication enabled.
- There needs to be network connectivity between the BlueXP classification instance and the shares.
- You can add a DFS (Distributed File System) share as a regular CIFS share. Because BlueXP classification is unaware that the share is built upon multiple servers/volumes combined as a single CIFS share, you might receive permission or connectivity errors about the share when the message really only applies to one of the folders/shares that is located on a different server/volume.
- For CIFS (SMB) shares, ensure that you have Active Directory credentials that provide read access to the shares. Admin credentials are preferred in case BlueXP classification needs to scan any data that requires elevated permissions.

If you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, it's recommended the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, configure the Active Directory user as part of a parent group in the organization which has permissions to all files.

- All CIFS file shares in a group must use the same Active Directory credentials.
- You can mix NFS and CIFS (using either Kerberos or NTLM) shares. You must add the shares to the group separately. That is, you must complete the process twice—once per protocol.
  - You cannot create a file shares group that mixes CIFS authentication types (Kerberos and NTLM).
- If you're using CIFS with Kerberos authentication, ensure the IP address provided is accessible to the BlueXP classification service. The file shares can't be added if the IP address is unreachable.

## Create a file shares group

When you add file shares to the group, you must use the format `<host_name>:/<share_path>`.

+

You can add file shares individually or you can enter a line-separated list of the file shares you want to scan. You can add up to 100 shares at a time.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration page, select **Add Working Environment > Add File Shares Group**.
3. In the Add File Shares Group dialog, enter the name for the group of shares then select **Continue**.
4. Select the protocol for the file shares you are adding.

**.If you're adding CIFS shares with NTLM authentication, enter the Active Directory credentials to access the CIFS volumes. Although read-only credentials are supported, it's recommended you provide full access with administrator credentials. Select Save.**

1. Add the file shares that you want to scan (one file share per line). Then select **Continue**.
2. A confirmation dialog displays the number of shares that were added.

If the dialog lists any shares that could not be added, capture this information so that you can resolve the issue. If the issue pertains to a naming convention, you can re-add the share with a corrected name.

### 3. Configure scanning on the volume:

- To enable mapping-only scans on file shares, select **Map**.
- To enable full scans on file shares, select **Map & Classify**.
- To disable scanning on file shares, select **Off**.



The switch at the top of the page for **Scan when missing "write attributes" permissions** is disabled by default. This means that if BlueXP classification doesn't have write attributes permissions in CIFS or write permissions in NFS, the system won't scan the files because BlueXP classification can't revert the "last access time" to the original timestamp.

If you switch **Scan when missing "write attributes" permissions** to **On**, the scan resets the last accessed time and scans all files regardless of permissions.

To learn more about the last accessed time stamp, see [xref:./Metadata collected from data sources in BlueXP classification](#).

## Result

BlueXP classification starts scanning the files in the file shares you added. You can [Track the scanning progress](#) and view the results of the scan in the **Dashboard**.



If the scan doesn't complete successfully for a CIFS configuration with Kerberos authentication, check the **Configuration** tab for errors.

## Edit a file shares group

After you create a file shares group, you can edit the CIFS protocol or add and remove file shares.

### Edit the CIFS protocol configuration

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration page, select the file shares group you want to modify.
3. Select **Edit CIFS Credentials**.



## Edit CIFS Authentication

Classification requires Active Directory credentials to access CIFS Volumes in Micky.

The credentials can be read-only, but providing admin credentials ensures that Classification can read any data that requires elevated permissions.

### Select Authentication Method

NTLM

Kerberos

Username

Password

domain\user or user@domain

Password

Save

Cancel

4. Choose the authentication method: **NTLM** or **Kerberos**.
5. Enter the Active Directory **Username** and **Password**.
6. Select **Save** to complete the process.

### Add file shares to compliance scans

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration page, select the file shares group you want to modify.
3. Select **+ Add shares**.
4. Select the protocol for the file shares you are adding.

If you're adding file shares to a protocol you've already configured, no changes are required.

If you're adding file shares with a second protocol, ensure you've properly configured the authentication properly as detailed in the [prerequisites](#).

5. Add the file shares you want to scan (one file share per line) using the format `<host_name>:/<share_path>`.
6. Select **Continue** to complete adding the file shares.

### Remove a file share from compliance scans

1. From the BlueXP classification menu, select **Configuration**.
2. Select the working environment from which you want to remove file shares.
3. Select **Configuration**.
4. From the Configuration page, select the Actions **...** for the file share you want to remove.
5. From the Actions menu, select **Remove Share**.

### Track the scanning progress

You can track the progress of the initial scan.

1. Select the **Configuration** menu.
2. Select the **Working Environment Configuration**.

The progress of each scan is shown as a progress bar.

3. Hover over the progress bar to see the number of files scanned relative to the total files in the volume.

### Scan StorageGRID data with BlueXP classification

Complete a few steps to start scanning data within StorageGRID directly with BlueXP classification.

#### Review StorageGRID requirements

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.

- You need to have the endpoint URL to connect with the object storage service.
- You need to have the Access Key and Secret Key from StorageGRID so that BlueXP classification can access the buckets.

#### Deploy the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

If you are scanning data from StorageGRID that is accessible over the internet, you can [deploy BlueXP classification in the cloud](#) or [deploy BlueXP classification in an on-premises location that has internet access](#).

If you are scanning data from StorageGRID that has been installed in a dark site that has no internet access, you need to [deploy BlueXP classification in the same on-premises location that has no internet access](#). This also requires that the BlueXP Connector is deployed in that same on-premises location.

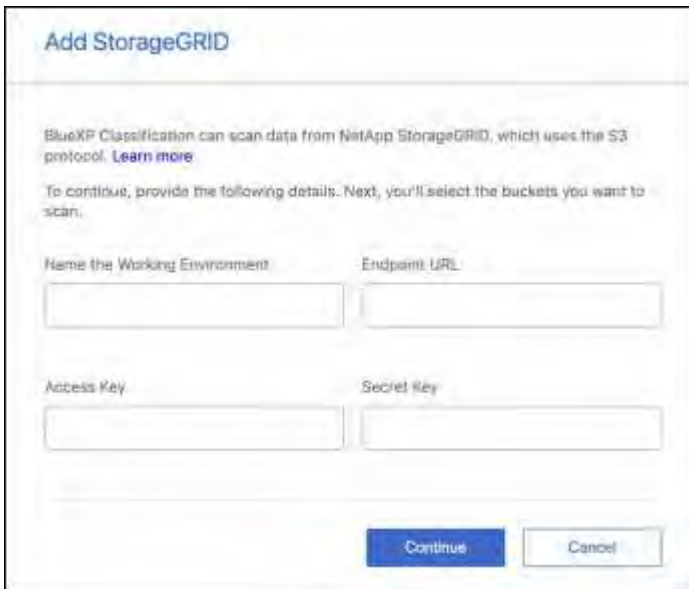
#### Add the StorageGRID service to BlueXP classification

Add the StorageGRID service.

##### Steps

1. From the BlueXP classification menu, select the **Configuration** option.
2. From the Configuration page, select **Add Working Environment > Add StorageGRID**.
3. In the Add StorageGRID Service dialog, enter the details for the StorageGRID service and click **Continue**.

- a. Enter the name you want to use for the Working Environment. This name should reflect the name of the StorageGRID service to which you are connecting.
- b. Enter the Endpoint URL to access the object storage service.
- c. Enter the Access Key and Secret Key so that BlueXP classification can access the buckets in StorageGRID.



The screenshot shows a web form titled "Add StorageGRID". Below the title, there is explanatory text: "BlueXP Classification can scan data from NetApp StorageGRID, which uses the S3 protocol. [Learn more](#)". This is followed by instructions: "To continue, provide the following details. Next, you'll select the buckets you want to scan." The form contains four input fields arranged in a 2x2 grid. The top row has "Name the Working Environment" and "Endpoint URL". The bottom row has "Access Key" and "Secret Key". At the bottom right of the form are two buttons: "Continue" (highlighted in blue) and "Cancel".

## Result

StorageGRID is added to the list of working environments.

## Enable and disable compliance scans on StorageGRID buckets

After you enable BlueXP classification on StorageGRID, the next step is to configure the buckets that you want to scan. BlueXP classification discovers those buckets and displays them in the working environment you created.

## Steps

1. In the Configuration page, locate the StorageGRID working environment.
2. On the StorageGRID working environment tile, select **Configuration**.

Buckets selected for Classification scan (5/8)

Scan	Storage Repository (Bucket)	Mapping status	Classification status	Required Action
Off   Map   <b>Map &amp; Classify</b>	bucketadipro	Finished 2024-09-05 10:33 Last full cycle: 2024-09-05 10:33	Mapped: 84 Classified: 5	...
Off   Map   <b>Map &amp; Classify</b>	datasense-0-files	Finished 2024-09-05 08:00 Last full cycle: 2024-09-05 08:00		...
Off   Map   <b>Map &amp; Classify</b>	datasense-10tb	Running 2024-09-04 07:25	Mapped: 3.7M Classified: 2.1M	...
Off   <b>Map</b>   Map & Classify	datasense-1tb	Running 2024-09-05 09:05 Last full cycle: 2024-09-05 03:04	Mapped: 1.3M	...
Off   <b>Map</b>   Map & Classify	datasense-1tb-2	Running 2024-09-05 09:06 Last full cycle: 2024-09-05 03:05	Mapped: 1.3M	...
Off   Map   Map & Classify	datasense-1tb-3	Not scanning		...

3. Complete one of the following steps to enable or disable scanning:

- To enable mapping-only scans on a bucket, select **Map**.
- To enable full scans on a bucket, select **Map & Classify**.
- To disable scanning on a bucket, select **Off**.

### Result

BlueXP classification starts scanning the buckets that you enabled. You can track the progress of the initial scan by navigating to the **Configuration** menu then selecting the **Working Environment configuration**. The progress of each scan is shown as a progress bar. You can also hover over the progress bar to see the number of files scanned relative to the total files in the volume. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

## Integrate your Active Directory with BlueXP classification

You can integrate a global Active Directory with BlueXP classification to enhance the results that BlueXP classification reports about file owners and which users and groups have access to your files.

When you set up certain data sources (listed below), you need to enter Active Directory credentials in order for BlueXP classification to scan CIFS volumes. This integration provides BlueXP classification with file owner and permissions details for the data that resides in those data sources. The Active Directory entered for those data sources might differ from the global Active Directory credentials you enter here. BlueXP classification will look in all integrated Active Directories for user and permission details.

This integration provides additional information in the following locations in BlueXP classification:

- You can use the "File Owner" [filter](#) and see results in the file's metadata in the Investigation pane. Instead of the file owner containing the SID (Security Identifier), it is populated with the actual user name.

You can also view more details about the file owner: account name, email address, and SAM account name, or view items owned by that user.

- You can see [full file permissions](#) for each file and directory when you click the "View all Permissions"

button.

- In the [Governance dashboard](#), the Open Permissions panel will show a greater level of detail about your data.



Local user SIDs, and SIDs from unknown domains, are not translated to the actual user name.

## Supported data sources

An Active Directory integration with BlueXP classification can identify data from within the following data sources:

- On-premises ONTAP systems
- Cloud Volumes ONTAP
- Azure NetApp Files
- FSx for ONTAP
- OneDrive accounts and SharePoint accounts (for legacy versions 1.30 and earlier)

There is no support for identifying user and permission information from Database schemas, Google Drive accounts, Amazon S3 accounts, or Object Storage that uses the Simple Storage Service (S3) protocol.

## Connect to your Active Directory server

After you've deployed BlueXP classification and have activated scanning on your data sources, you can integrate BlueXP classification with your Active Directory. Active Directory can be accessed using a DNS Server IP address or an LDAP Server IP address.

The Active Directory credentials can be read-only, but providing admin credentials ensures that BlueXP classification can read any data that requires elevated permissions. The credentials are stored on the BlueXP classification instance.

For CIFS volumes/file shares, if you want to make sure your files "last accessed times" are unchanged by BlueXP classification scans, we recommend that the user has Write Attributes permission. If possible, we recommend making the Active Directory configured user part of a parent group in the organization which has permissions to all files.

### Requirements

- You must have an Active Directory already set up for the users in your company.
- You must have the information for the Active Directory:
  - DNS Server IP address, or multiple IP addressesor
  - LDAP Server IP address, or multiple IP addresses
  - User Name and Password to access the server
  - Domain Name (Active Directory Name)
  - Whether you are using secure LDAP (LDAPS) or not
  - LDAP Server Port (typically 389 for LDAP, and 636 for secure LDAP)

- The following ports must be open for outbound communication by the BlueXP classification instance:

Protocol	Port	Destination	Purpose
TCP & UDP	389	Active Directory	LDAP
TCP	636	Active Directory	LDAP over SSL
TCP	3268	Active Directory	Global Catalog
TCP	3269	Active Directory	Global Catalog over SSL

## Steps

1. From the BlueXP classification Configuration page, click **Add Active Directory**.



2. In the Connect to Active Directory dialog, enter the Active Directory details and click **Connect**.

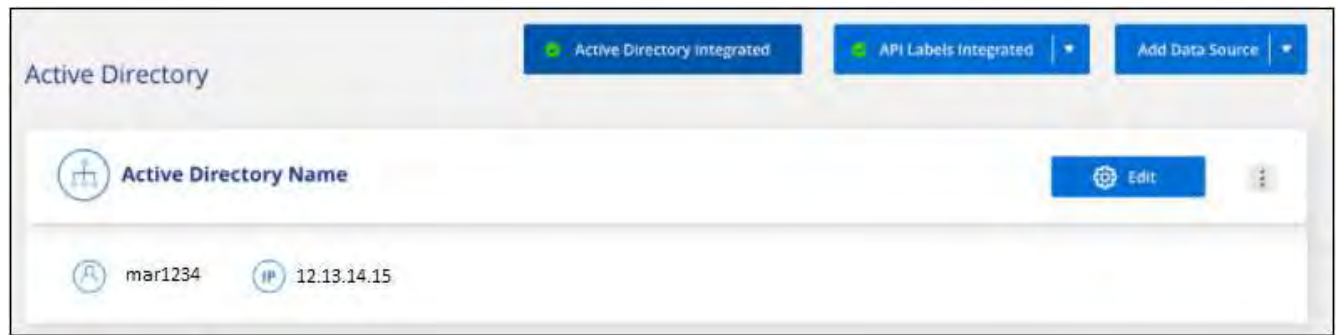
You can add multiple IP addresses, if required, by clicking **Add IP**.

 A screenshot of the "Connect to Active Directory" dialog box. The title is "Connect to Active Directory". It has several input fields:
 

- "Username" with the value "mar1234".
- "Password" with masked characters "\*\*\*\*\*".
- "DNS Server IP address:" with a radio button selected, containing the value "12.20.70.00" and a "+ Add IP" button.
- "Domain Name" with the value "mar@netapp.com".
- "LDAP Server IP Address:" with a radio button unselected and an "+ Add IP" button.
- "LDAP Server Port" with the value "389".
- "LDAP Secure Connection" checkbox, which is unchecked.


 At the bottom right, there are two buttons: "Connect" (highlighted with a red rectangular box) and "Cancel".

BlueXP classification integrates to the Active Directory, and a new section is added to the Configuration page.



## Manage your Active Directory integration

If you need to modify any values in your Active Directory integration, click the **Edit** button and make the changes.

You can also delete the integration selecting the  button then **Remove Active Directory**.

# Use BlueXP classification

## View governance details about the data stored in your organization with BlueXP classification

Gain control of the costs related to the data on your organization's storage resources. BlueXP classification identifies the amount of stale data, non-business data, duplicate files, and very large files in your systems so you can decide whether you want to remove or tier some files to less expensive object storage.

This is where you should begin your research. From the Governance dashboard, you can select an area for further investigation.

Additionally, if you're planning to migrate data from on-premises locations to the cloud, you can view the size of the data and whether any of the data contains sensitive information before moving it.

### Review the Governance dashboard

The Governance dashboard provides information so that you can increase the efficiency and control the costs related to the data stored on your storage resources.



**NetApp BlueXP** | Home | BlueXP Search | Organization: Demo SIM | Project: group1 | Connector: OCCMassDem... | Settings | Help | Logout

**Classification** | Governance | Compliance | Investigation | Classification settings | Policies | Configuration

---

### Savings Opportunities

Scale Data

216.2K items | 227.8 GB

Optimize Storage

Non-Business Data

18.6K items | 24.3 GB

Optimize Storage

Duplicate Files

159.8K items | 188.6 GB

Optimize Storage

### Policies

Policy	Items
bb1	132.3K
Data 3-5 years old	105.9K

[View All](#)

---

### Data Overview

[Data Discovery Assessment Report](#) | 
 [Full Data Mapping Overview Report](#) | 
 Scanned: 265.3 GB | 
 270.6K Files | 
 141 Tables

#### Top Data Repositories by Sensitivity Level

Repository	Non-Sensitive	Personal	Sensitive	Total Items
CVD	~50%	~30%	~20%	133.8K
Amazon S3	~50%	~30%	~20%	131.5K
File Shares	~50%	~30%	~20%	5.2K
Database	~50%	~30%	~20%	141

---

#### Sensitive Data and Wide Permissions

Y-axis: Detected sensitive classes (Sensitive to Not sensitive)  
X-axis: Access level (Restrictive to Permissive)  
Z-axis: Number of files (0 to 10K)

#### Open Permissions

■ 51% - No Open Permissions  
■ 48% - Open to Organization  
■ 1% - Open to Public

---

#### Age of Data

Modified | Created | Last Accessed

Y-axis: Number of files (0 to 120K)  
X-axis: Age ranges (30 days to Over 7 years)

#### Size of Data

Y-axis: Size (0 to 240K)  
X-axis: File types (PDF files, CSV files, etc.)

---

### Classification

#### 40 Categories

Miscellaneous	18.2K items
Code	18K items
Bank Statements	13K items
Miscellaneous Spreadsheets	5.2K items

[View All](#)

#### 127 File Types

PDF	99.9K items
TXT	88.9K items
DOCX	31.4K items
MAP	16K items

[View All](#)

Classification Documentation | v1.41.0

## Steps

1. From the BlueXP menu, select **Governance > Classification**.
2. Select **Governance**.

The Governance dashboard appears.

## Review savings opportunities

The *Saving Opportunities* component shows data that you can delete or tier to less expensive object storage. The data in *Saving Opportunities* update every 2 hours and can be manually updated.

## Steps

1. From the BlueXP classification menu, select **Governance**.
2. Within each Savings Opportunities tile of the Governance dashboard, select **Optimize Storage** to view the filtered results in the Investigation page. To discover any data you should delete or tier to less expensive storage, investigate the the *Saving Opportunities*.
  - **Stale Data** - Data that was last modified over 3 years ago.
  - **Non-Business Data** - Data considered not to be business related, based on their Category or File Type. This includes:
    - Application Data
    - Audio
    - Executables
    - Images
    - Logs
    - Videos
    - Miscellaneous (general "other" category)
  - **Duplicate Files** - Files that are duplicated in other locations in the data sources you are scanning. [See what types of duplicate files are displayed.](#)



If any of your data sources implement data tiering, old data that already resides in object storage can be identified in the *Stale Data* category.

## Review saved searches with the largest number of results

In the *Saved searches* tab, the searches with the greatest number of results appear at the top of the list. This data updates every two hours.

For details about saved searches, see [Create saved searches](#).

## Steps

1. From the BlueXP classification menu, select **Governance**.
2. In the Governance dashboard, locate the Saved Searches tile. Select the name of a saved search to display the results in the Investigation page.
3. Select **View All** to view the list of all available saved searches.  
In the *Saved searches* area, the searches with the greatest number of results appear at the top of the list.

## Create the Data Discovery Assessment Report

The Data Discovery Assessment Report provides a high-level analysis of the scanned environment to show areas of concern and potential remediation steps. The results are based on both mapping and classifying your data. The goal of this report is to raise awareness of three significant aspects of your dataset:


Feature	Description
Data governance concerns	A detailed picture of all the data you own and areas where you may be able to reduce the amount of data to save costs.
Data security exposures	Areas where your data is accessible to internal or external attacks because of broad access permissions.
Data compliance gaps	Where your personal or sensitive personal information is located for both security and for DSARs (data subject access requests).

Using this report, you might take the following actions:

- Reduce storage costs by changing your retention policy, or by moving or deleting certain data (stale, duplicate, or non-business data)
- Protect your data that has broad permissions by revising global group management policies
- Protect your data that has personal or sensitive personal information by moving PII to more secure data stores

### Steps

1. From the BlueXP menu, select **Governance > Classification**.
2. Select **Governance**.
3. Select **Data Discovery Assessment Report**.



### Result

BlueXP classification generates a PDF report that you can review and send to other groups as needed.

## Create the Data Mapping Overview Report


The Data Mapping Overview Report provides an overview of the data being stored in your corporate data sources to assist you with decisions of migration, back up, security, and compliance processes. The report summarizes all working environments and data sources. It also provides an analysis for each working environment.


The report includes the following information:

Category	Description
Usage Capacity	For all working environments: Lists the number of files and the used capacity for each working environment. For single working environments: Lists the files that are using the most capacity.
Age of Data	Provides three charts and graphs for when files were created, last modified, or last accessed. Lists the number of files, and their used capacity, based on certain date ranges.
Size of Data	Lists the number of files that exist within certain size ranges in your working environments.
File Types	Lists the total number of files and the used capacity for each type of file being stored in your working environments.

### Steps

1. From the BlueXP menu, select **Governance > Classification**.
2. Select **Governance**.
3. Select **Full Data Mapping Overview Report**.



4. To customize the company name that appears on the first page of the report, from the top right of the BlueXP classification page, select . Then select **Change company name**. The next time you generate the report, it will include the new name.

### Result

BlueXP classification generates a .pdf report that you can review and send to other groups as needed.

If the report is larger than 1 MB, the .pdf file is retained on the BlueXP classification instance and you'll see a pop-up message about the exact location. When BlueXP classification is installed on a Linux machine on your premises, or on a Linux machine you deployed in the cloud, you can navigate directly to the .pdf file. When BlueXP classification is deployed in the cloud, you'll need to SSH to the BlueXP classification instance to download the .pdf file.

### Review the top data repositories listed by data sensitivity

The *Top Data Repositories by Sensitivity Level* area of the Data Mapping Overview report lists the top four data repositories (working environments and data sources) that contain the most sensitive items. The bar chart for each working environment is divided into:

- Non-Sensitive data
- Personal data
- Sensitive Personal data

This data refreshes every two hours and can be manually refreshed.

### Steps

1. To see the total number of items in each category, position your cursor over each section of the bar.
2. To filter results that will appear in the Investigation page, select each area in the bar and investigate further.

### Review sensitive data and wide permissions

The *Sensitive Data and Wide Permissions* area of the Data Mapping Overview report shows the percentage of files that contain sensitive data and have wide permissions. The chart shows the following types of permissions:

- From the most restrictive permissions to the most permissive restrictions on the horizontal axis.
- From the least sensitive data to the most sensitive data on the vertical axis.

#### Steps

1. To see the total number of files in each category, position your cursor over each box.
2. To filter results that will appear in the Investigation page, select a box and investigate further.

### Review data listed by types of open permissions

The *Open Permissions* area of the Data Mapping Overview report shows the percentage for each type of permissions that exist for all files that are being scanned. The chart shows the following types of permissions:

- No Open Permissions
- Open to Organization
- Open to Public
- Unknown Access

#### Steps

1. To see the total number of files in each category, position your cursor over each box.
2. To filter results that will appear in the Investigation page, select a box and investigate further.

### Review the age and size of data

You might want to investigate the items in the *Age* and *Size* graphs of the Data Mapping Overview report to see if there is any data you should delete or tier to less expensive object storage.

#### Steps

1. In the Age of Data chart, to see details about the age of the data, position your cursor over a point in the chart.
2. To filter by an age or size range, select that age or size.
  - **Age of Data graph** - Categorizes data based on the time it was created, the last time it was accessed, or the last time it was modified.
  - **Size of Data graph** - Categorizes data based on size.



If any of your data sources implement data tiering, old data that already resides in object storage might be identified in the *Age of Data* graph.

## Review the most identified data classifications in your data

The *Classification* area of the Data Mapping Overview report provides a list of the most identified [Categories](#) and [File types](#) in your scanned data.

Categories can help you understand what's happening with your data by showing you the types of information that you have. For example, a category like "resumes" or "employee contracts" can include sensitive data. When you investigate the results, you might find that employee contracts are stored in a nonsecure location. You can then correct that issue.

See [Viewing files by categories](#) for more information.

### Steps

1. From the BlueXP menu, select **Governance > Classification**.
2. Select **Governance** then the **Data Discovery Assessment Report** button.

### Result

BlueXP classification generates a .pdf report that you can review and send to other groups as needed.

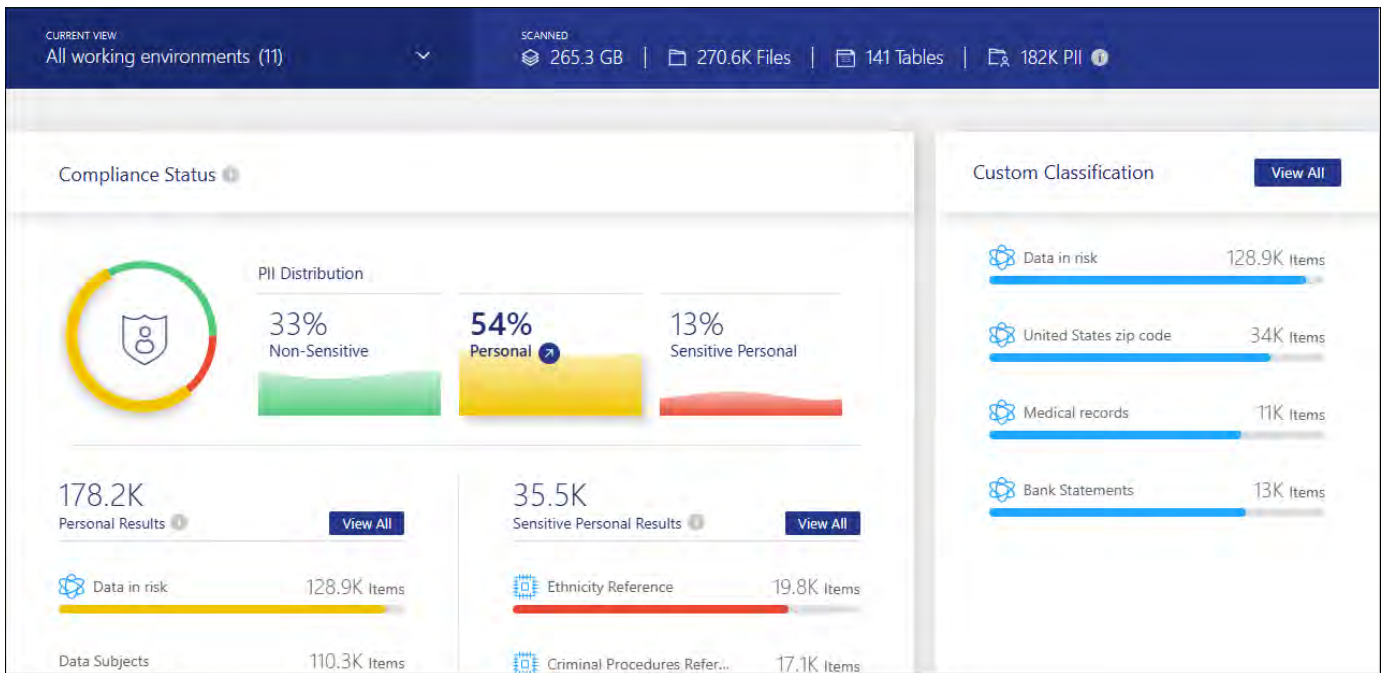
## View compliance details about the private data stored in your organization with BlueXP classification

Gain control of your private data by viewing details about the personal data (PII) and sensitive personal data (SPII) in your organization. You can also gain visibility by reviewing the categories and file types that BlueXP classification found in your data.



The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.

By default, the BlueXP classification dashboard displays compliance data for all working environments and databases. To see data for only some of the working environments, select them.



Filter the results from the Data Investigation page and download a report of the results as a CSV file. See [Filtering data in the Data Investigation page](#) for details.

## View files that contain personal data

BlueXP classification automatically identifies specific words, strings, and patterns (Regex) inside the data. For example, Personal Identification Information (PII), credit card numbers, social security numbers, bank account numbers, passwords, and more. [See the full list](#). BlueXP classification identifies this type of information in individual files, in files within directories (shares and folders), and in database tables.

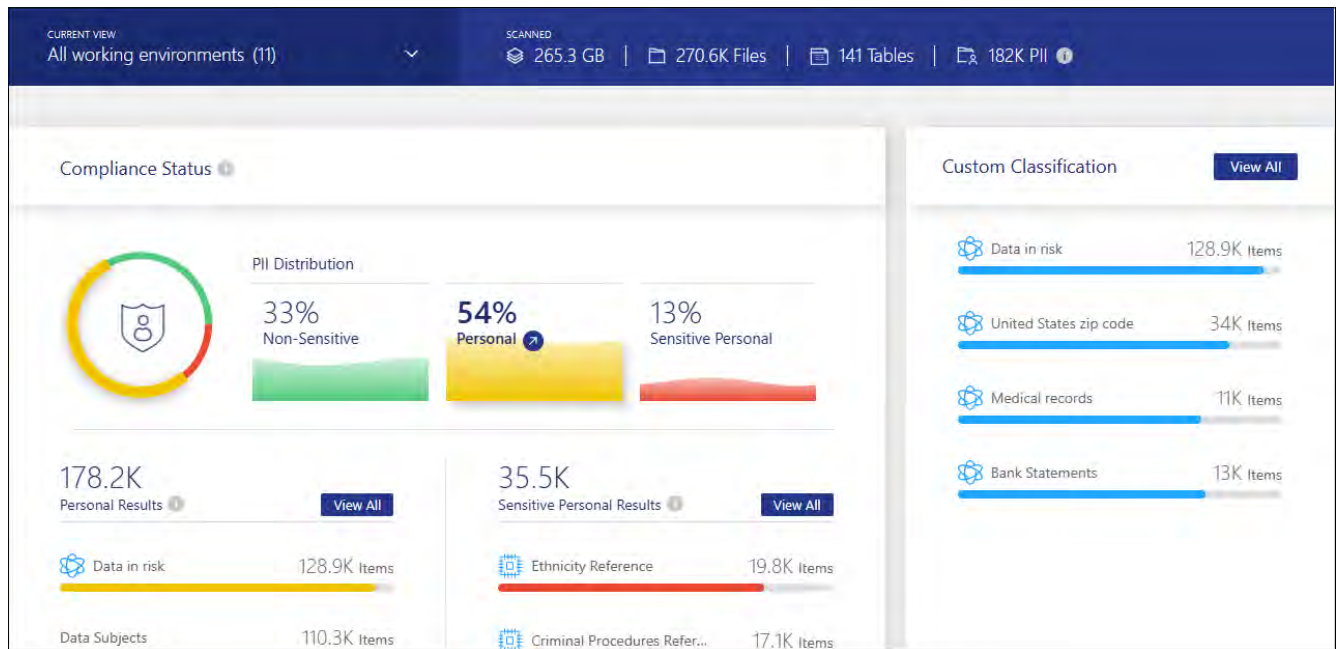
You can also create custom search terms to identify personal data specific to your organization. For more information, see [Create a custom classification](#).

For some types of personal data, BlueXP classification uses *proximity validation* to validate its findings. The validation occurs by looking for one or more predefined keywords in proximity to the personal data that was found. For example, BlueXP classification identifies a U.S. social security number (SSN) as an SSN if it sees a proximity word next to it—for example, *SSN* or *social security*. [The table of personal data](#) shows when BlueXP classification uses proximity validation.


### Steps

1. From the BlueXP classification menu, select the **Compliance** tab.
2. To investigate the details for all personal data, select the icon next to the personal data percentage.





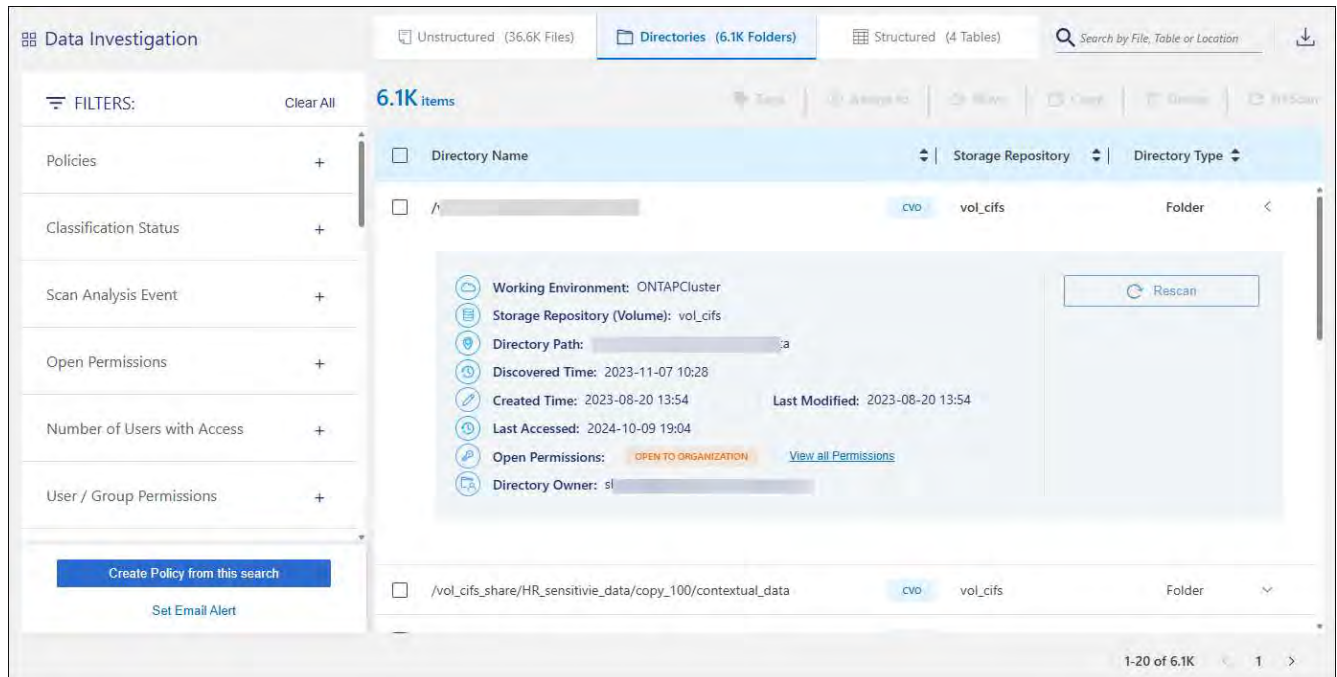
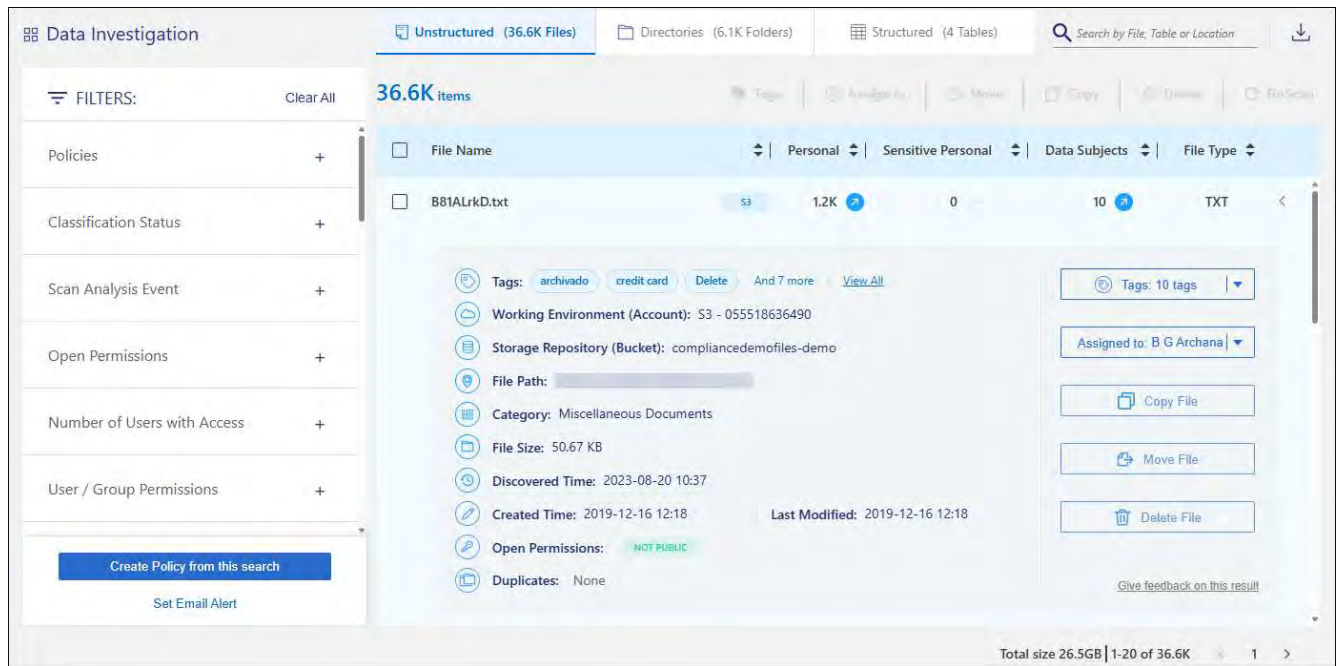
- To investigate the details for a specific type of personal data, select **View All** and then select the **Investigate Results** arrow icon for a specific type of personal data, for example, email addresses.



- Investigate the data by searching, sorting, expanding details for a specific file, selecting the **Investigate Results** arrow to see masked information, or by downloading the file list.

The two screenshots below show personal data found in individual files, and found in files within directories (shares and folders). You can also select the **Structured** tab to view personal data found in databases.





## View files that contain sensitive personal data

BlueXP classification automatically identifies special types of sensitive personal information, as defined by privacy regulations such as [articles 9 and 10 of the GDPR](#). For example, information regarding a person's health, ethnic origin, or sexual orientation. [See the full list](#). BlueXP classification identifies this type of information in individual files, in files within directories (shares and folders), and in database tables.

BlueXP classification uses artificial intelligence (AI), natural language processing (NLP), machine learning (ML), and cognitive computing (CC) to understand the meaning of the content that it scans in order to extract entities and categorize it accordingly.

For example, one sensitive GDPR data category is ethnic origin. Because of its NLP abilities, BlueXP classification can distinguish the difference between a sentence that reads "George is Mexican" (indicating

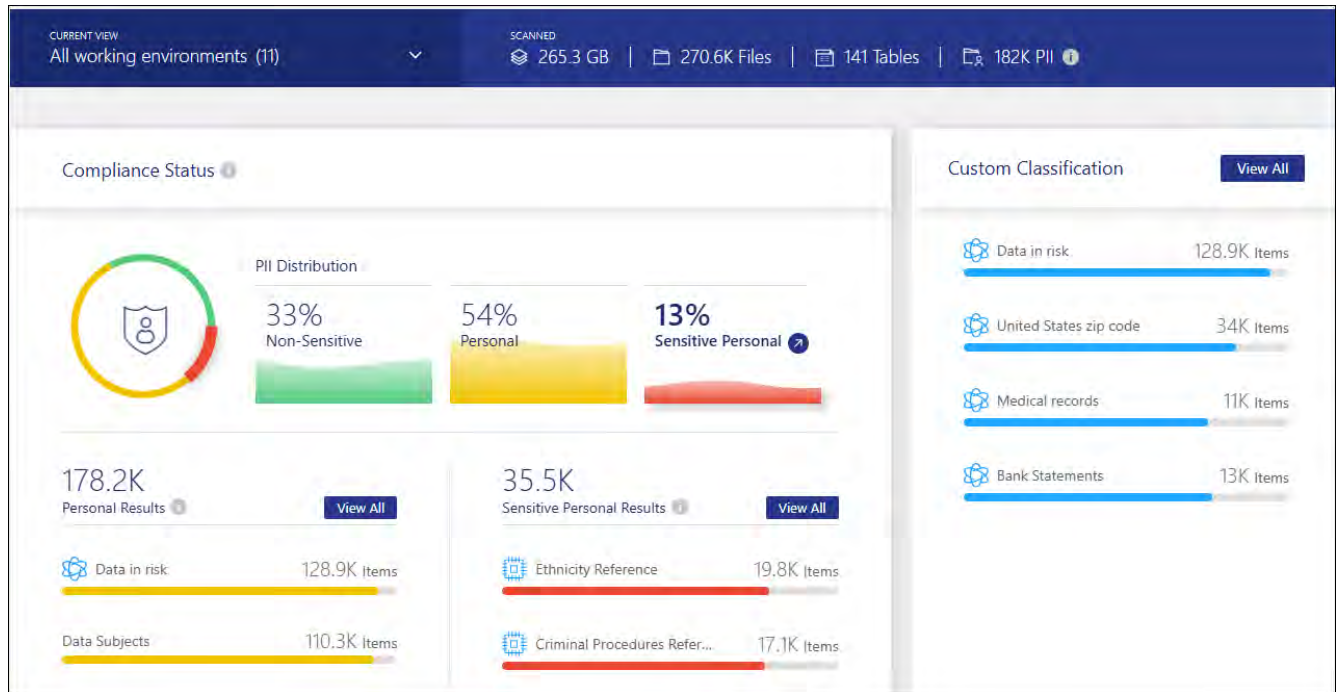
sensitive data as specified in article 9 of the GDPR), versus "George is eating Mexican food."



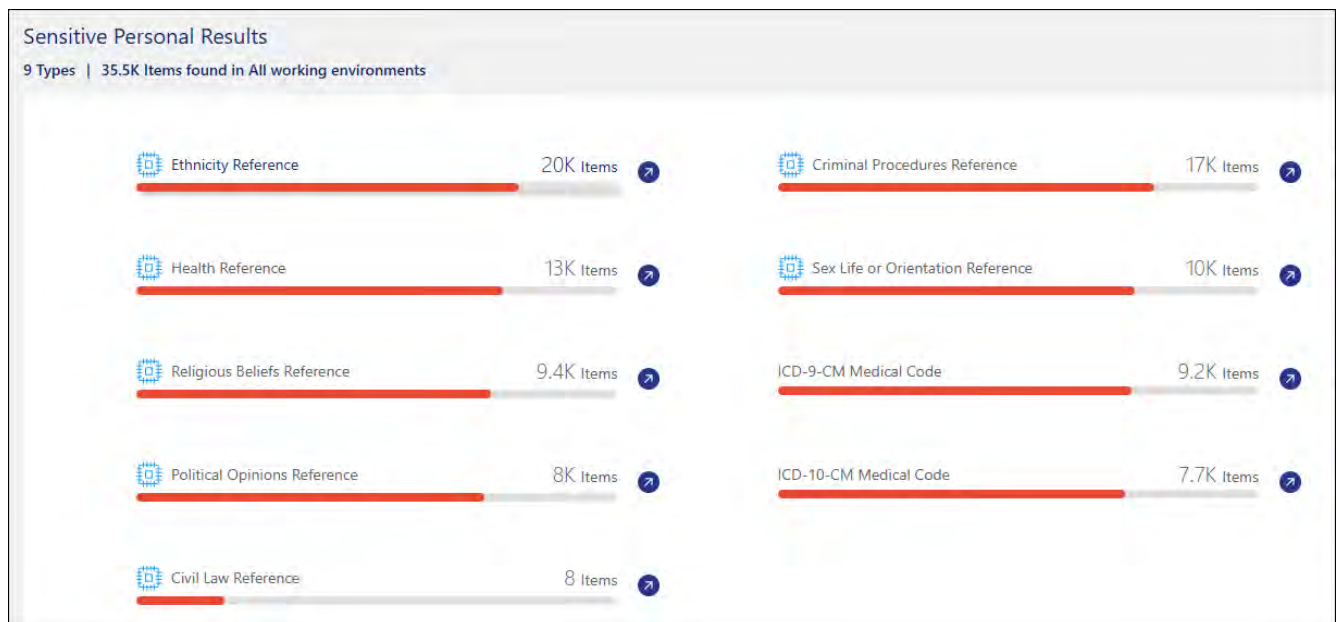
Only English is supported when scanning for sensitive personal data. Support for more languages will be added later.

### Steps

1. From the BlueXP classification menu, select **Compliance**.
2. To investigate the details for all sensitive personal data, select the icon next to the sensitive personal data percentage.



3. To investigate the details for a specific type of sensitive personal data, select **View All** and then select the **Investigate Results** arrow icon for a specific type of sensitive personal data.



- Investigate the data by searching, sorting, expanding details for a specific file, clicking **Investigate Results** to see masked information, or by downloading the file list.

## View files by categories

BlueXP classification takes the data that it scanned and divides it into different types of categories. Categories are topics based on AI analysis of the content and metadata of each file. [See the list of categories.](#)

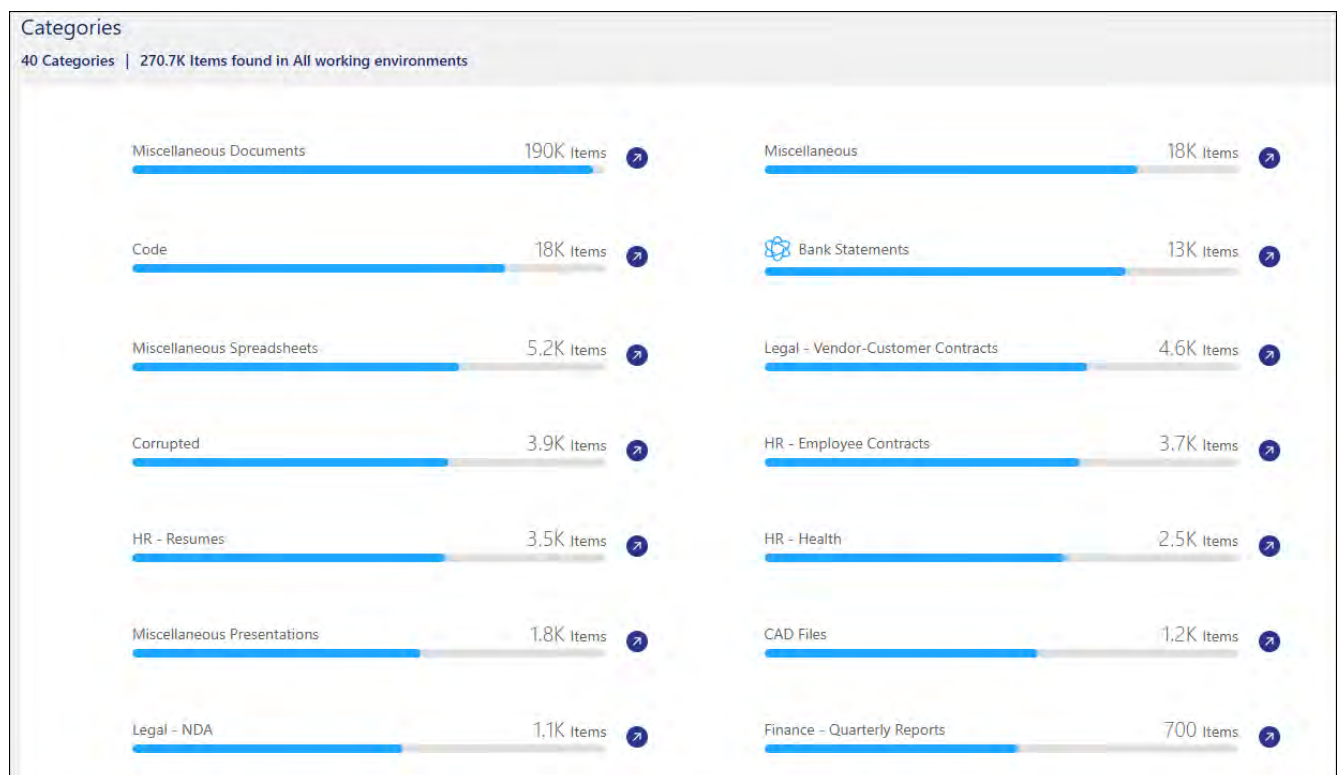
Categories can help you understand what's happening with your data by showing you the types of information that you have. For example, a category like resumes or employee contracts can include sensitive data. When you investigate the results, you might find that employee contracts are stored in a nonsecure location. You can then correct that issue.



English, German, and Spanish are supported for categories. Support for more languages will be added later.

### Steps

- From the BlueXP classification menu, select the **Compliance** tab.
- Select the **Investigate Results** arrow icon for one of the top 4 categories directly from the main screen, or select **View All** and then select the icon for any of the categories.



- Investigate the data by searching, sorting, expanding details for a specific file, selecting the **Investigate Results** arrow to see masked information, or by downloading the file list.

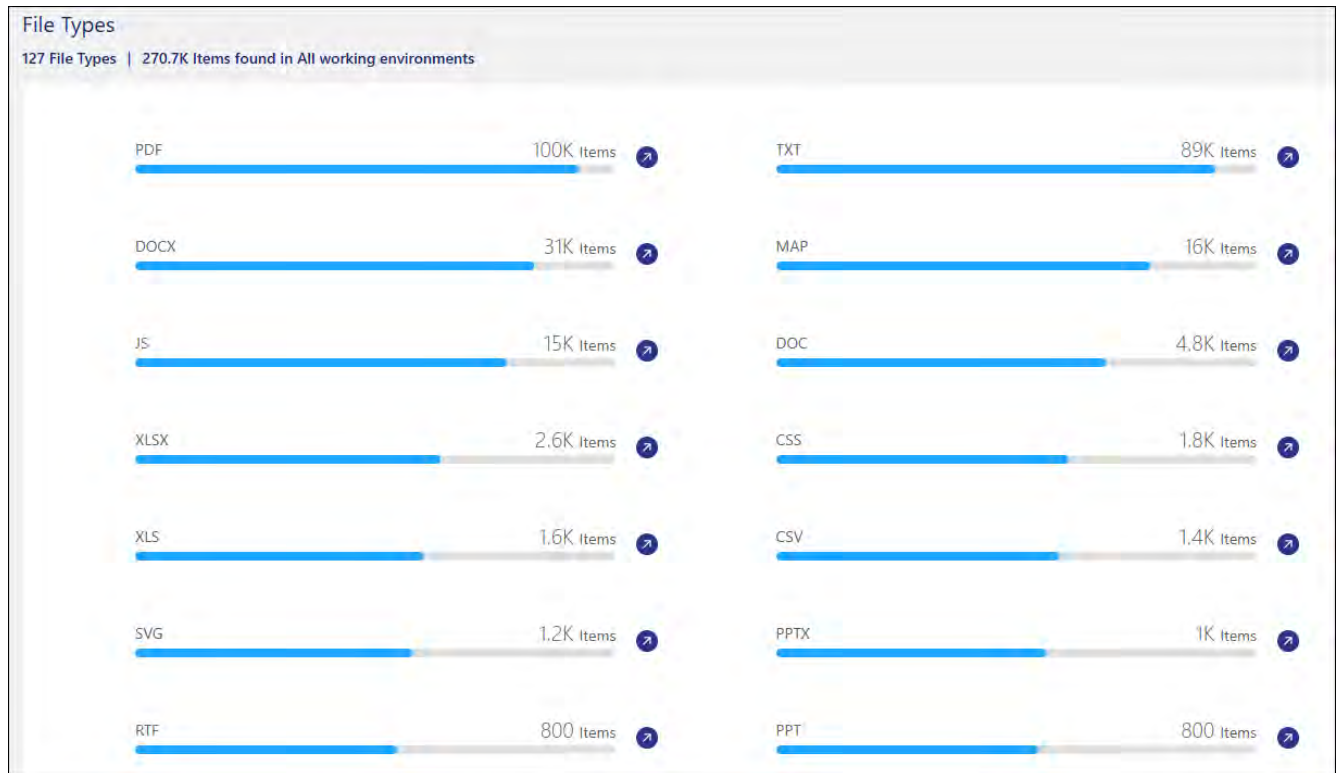
## View files by file types

BlueXP classification takes the data that it scanned and breaks it down by file type. Reviewing your file types can help you control your sensitive data because you might find that certain file types are not stored correctly. [See the list of file types.](#)

For example, you might be storing CAD files that include very sensitive information about your organization. If they are unsecured, you can take control of the sensitive data by restricting permissions or moving the files to another location.

### Steps

1. From the BlueXP classification menu, select the **Compliance** tab.
2. Select the **Investigate Results** arrow icon for one of the top 4 file types directly from the main screen, or select **View All** and then select the icon for any of the file types.



3. Investigate the data by searching, sorting, expanding details for a specific file, selecting the **Investigate Results** arrow to see masked information, or by downloading the file list.

## Categories of private data in BlueXP classification

There are many types of private data that BlueXP classification can identify in your volumes and databases.

BlueXP classification identifies two types of personal data:

- **Personally identifiable information (PII)**
- **Sensitive personal information (SPII)**



If you need BlueXP classification to identify other private data types, such as additional national ID numbers or healthcare identifiers, contact your account manager.

### Types of personal data

The personal data, or *personally identifiable information* (PII), found in files can be general personal data or

national identifiers. The third column in the table below identifies whether BlueXP classification uses [proximity validation](#) to validate its findings for the identifier.

The languages in which these items can be recognized are identified in the table.

Type	Identifier	Proximity validation?	English	German	Spanish	French	Japanese
General	Credit card number	Yes	✓	✓	✓		✓
	Data Subjects	No	✓	✓	✓		
	Email Address	No	✓	✓	✓		✓
	IBAN Number (International Bank Account Number)	No	✓	✓	✓		✓
	IP Address	No	✓	✓	✓		✓
	Password	Yes	✓	✓	✓		✓

Type	Identifier	Proximity validation?	English	German	Spanish	French	Japanese
------	------------	--------------------------	---------	--------	---------	--------	----------

National Identifiers

--

Type	Identifier	Proximity validation?	English	German	Spanish	French	Japanese
------	------------	-----------------------	---------	--------	---------	--------	----------

Type	Identifier	Proximity validation?	English	German	Spanish	French	Japanese
------	------------	-----------------------	---------	--------	---------	--------	----------



Type	Identifier	Proximity	English	German	Spanish	French	Japanese
	Corporate)						
	Latvian ID	Yes	✓	✓	✓		
	Lithuanian ID	Yes	✓	✓	✓		
	Luxembourg ID	Yes	✓	✓	✓		
	Maltese ID	Yes	✓	✓	✓		
	National Health Service (NHS) Number	Yes	✓	✓	✓		
	New Zealand Bank Account	Yes	✓	✓	✓		
	New Zealand Driver's License	Yes	✓	✓	✓		
	New Zealand IRD Number (Tax ID)	Yes	✓	✓	✓		
	New Zealand NHI (National Health Index) Number	Yes	✓	✓	✓		
	New Zealand Passport Number	Yes	✓	✓	✓		
	Polish ID (PESEL)	Yes	✓	✓	✓		
	Portuguese Tax Identification Number (NIF)	Yes	✓	✓	✓		
	Romanian ID (CNP)	Yes	✓	✓	✓		
	Singapore National Registration Identity Card (NRIC)	Yes	✓	✓	✓		
	Slovenian ID (EMSO)	Yes	✓	✓	✓		
	South African ID	Yes	✓	✓	✓		
	Spanish Tax Identification Number	Yes	✓	✓	✓		
	Swedish ID	Yes	✓	✓	✓		
	UK ID (NINO)	Yes	✓	✓	✓		
	USA California Driver's License	Yes	✓	✓	✓		
	USA Indiana Driver's License	Yes	✓	✓	✓		
	USA New York Driver's License	Yes	✓	✓	✓		
	USA Texas Driver's License	Yes	✓	✓	✓		
	USA Social Security Number (SSN)	Yes	✓	✓	✓		

## Types of sensitive personal data

BlueXP classification can find the following sensitive personal information (SPII) in files.

The items in this category can be recognized only in English at this time.

- **Criminal Procedures Reference:** Data concerning a natural person's criminal convictions and offenses.
- **Ethnicity Reference:** Data concerning a natural person's racial or ethnic origin.
- **Health Reference:** Data concerning a natural person's health.
- **ICD-9-CM Medical Codes:** Codes used in the medical and health industry.
- **ICD-10-CM Medical Codes:** Codes used in the medical and health industry.

- **Philosophical Beliefs Reference:** Data concerning a natural person’s philosophical beliefs.
- **Political Opinions Reference:** Data concerning a natural person’s political opinions.
- **Religious Beliefs Reference:** Data concerning a natural person’s religious beliefs.
- **Sex Life or Orientation Reference:** Data concerning a natural person’s sex life or sexual orientation.

## Types of categories

BlueXP classification categorizes your data as follows.

Most of these categories can be recognized in English, German, and Spanish.

Category	Type	English	German	Spanish
Finance	Balance Sheets	✓	✓	✓
	Purchase Orders	✓	✓	✓
	Invoices	✓	✓	✓
	Quarterly Reports	✓	✓	✓
HR	Background Checks	✓		✓
	Compensation Plans	✓	✓	✓
	Employee Contracts	✓		✓
	Employee Reviews	✓		✓
	Health	✓		✓
	Resumes	✓	✓	✓
Legal	NDA's	✓	✓	✓
	Vendor-Customer contracts	✓	✓	✓
Marketing	Campaigns	✓	✓	✓
	Conferences	✓	✓	✓
Operations	Audit Reports	✓	✓	✓
Sales	Sales Orders	✓	✓	
Services	RFI	✓		✓
	RFP	✓		✓
	SOW	✓	✓	✓
	Training	✓	✓	✓
Support	Complaints and Tickets	✓	✓	✓

The following metadata is also categorized and identified in the same supported languages:

- Application Data
- Archive Files

- Audio
- Breadcrumbs from BlueXP classification  
Business Application Data
- CAD Files
- Code
- Corrupted
- Database and index files
- Design Files
- Email Application Data
- Encrypted (files with a high entropy score)
- Executables
- Financial Application Data
- Health Application Data
- Images
- Logs
- Miscellaneous Documents
- Miscellaneous Presentations
- Miscellaneous Spreadsheets
- Miscellaneous "Unknown"
- Password Protected files
- Structured Data
- Videos
- Zero-Byte Files

## Types of files

BlueXP classification scans all files for category and metadata insights and displays all file types in the file types section of the dashboard. When BlueXP classification detects Personal Identifiable Information (PII) or when it performs a DSAR search, only the following file formats are supported:

.CSV, .DCM, .DOC, .DOCX, .JSON, .PDF, .PPTX, .RTF, .TXT, .XLS, .XLSX, Docs, Sheets, and Slides

## Accuracy of information found

NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that BlueXP classification identifies. You should always validate the information by reviewing the data.

Based on our testing, the table below shows the accuracy of the information that BlueXP classification finds. We break it down by *precision* and *recall*:

### Precision

The probability that what BlueXP classification finds has been identified correctly. For example, a precision rate of 90% for personal data means that 9 out of 10 files identified as containing personal information,

actually contain personal information. 1 out of 10 files would be a false positive.

## Recall

The probability for BlueXP classification to find what it should. For example, a recall rate of 70% for personal data means that BlueXP classification can identify 7 out of 10 files that actually contain personal information in your organization. BlueXP classification would miss 30% of the data and it won't appear in the dashboard.

We are constantly improving the accuracy of our results. Those improvements will be automatically available in future BlueXP classification releases.

Type	Precision	Recall
Personal data - General	90%-95%	60%-80%
Personal data - Country identifiers	30%-60%	40%-60%
Sensitive personal data	80%-95%	20%-30%
Categories	90%-97%	60%-80%

## Create a custom classification in BlueXP classification


With BlueXP classification, you can create a custom search for sensitive information. The search can be scoped to a regular expression (regex).

### Create a custom classification

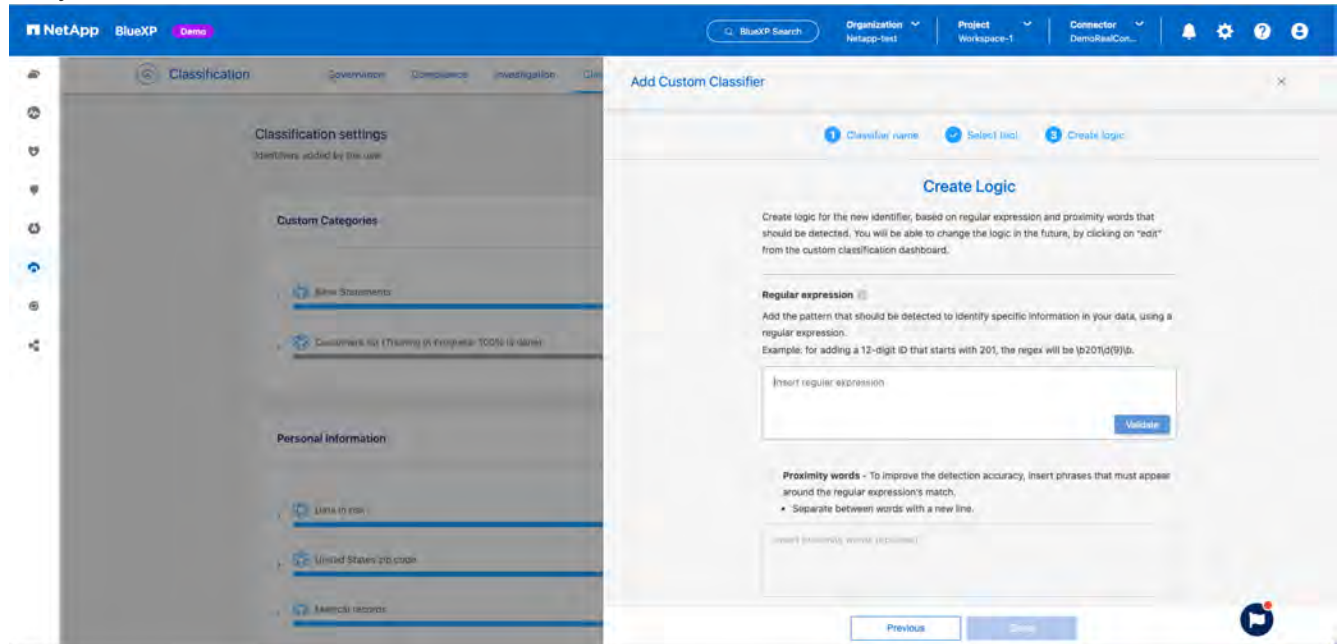
Custom classification is only available for Map & Classify scans, not mapping-only scans. This feature is currently in preview.

#### Steps

1. Select the **Custom classification** tab.



2. Select the **Add New Classifier** button.
3. Add a Name and Description for the new classifier.
4. To add the customization as a regular expression, select **Custom regular expression** then **Next**.
5. Add a pattern to detect the specific information of your data. Select **Validate** to confirm the syntax of your entry.



6. Select **Done** to create the custom classification.

The new customization is captured in the next scheduled scan. To view results, see [Generate compliance reports](#).

# Investigate the data stored in your organization with BlueXP classification

Investigate the data from your organization by viewing details in the Data Investigation page. Here is where you can continue your research after looking at the Governance dashboard. On the Investigation page, you can filter the data using one of the many filters to show only the results you want to see. You can also view file metadata, permissions for files and directories, and check for duplicate files in your storage systems.

You can navigate to this page from many areas of the BlueXP classification UI, including the Governance and Compliance dashboards with the filters selected already on those pages. You can export the data into a CSV or JSON file for further analysis or to share with others.



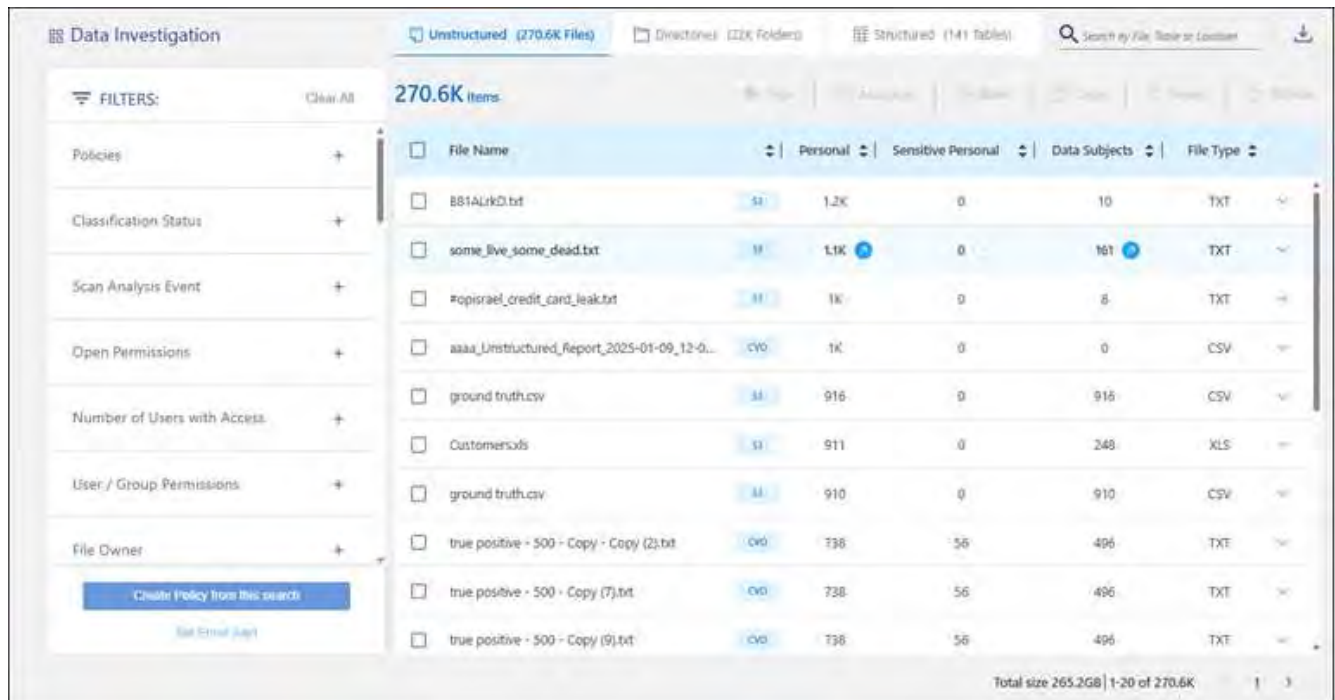
The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.

## Filter data in the Data Investigation page

You can filter the contents of the investigation page to display only the results you want to see.

### Steps

1. From the BlueXP classification menu, select **Investigation**.
2. On the Data Investigation page, do any of the following:
3. To download the contents of the page as a report after you've refined it, select the button.



4. To view the data from files (unstructured data), directories (folders and file shares), or from databases (structured data), select one of the tabs at the top.

5. To sort the results in numerical or alphabetical order, select the control at the top of each column.
6. To refine the results even more, select one of the filters in the Filter pane.



You can only view the first 10,000 results—or 500 pages—for a scan on the Data Investigation page.

### Filter data by sensitivity and content

Use the following filters to view how much sensitive information is contained in your data.

Filter	Details
Category	Select the <a href="#">types of categories</a> .
Sensitivity Level	Select the sensitivity level: Personal, Sensitive personal, or Non sensitive.
Number of identifiers	Select the range of detected sensitive identifiers per file. Includes personal data and sensitive personal data. When filtering in Directories, BlueXP classification totals the matches from all files in each folder (and sub-folders).  NOTE: The December 2023 (version 1.26.6) release removed the option to calculate the number of personal identifiable information (PII) data by Directories.
Personal Data	Select the <a href="#">types of personal data</a> .
Sensitive Personal Data	Select the <a href="#">types of sensitive personal data</a> .
Data Subject	Enter a data subject's full name or known identifier. <a href="#">Learn more about data subjects here</a> .

### Filter data by user owner and user permissions

Use the following filters to view file owners and permissions to access your data.

Filter	Details
Open Permissions	Select the type of permissions within the data and within folders/shares.
User / Group Permissions	Select one or multiple user names and/or group names, or enter a partial name.
File Owner	Enter the file owner name.
Number of users with access	Select one or multiple category ranges to show which files and folders are open to a certain number of users.

### Filter data by time

Use the following filters to view data based on time criteria.

Filter	Details
Created Time	Select a time range when the file was created. You can also specify a custom time range to further refine the search results.

Filter	Details
Discovered Time	Select a time range when BlueXP classification discovered the file. You can also specify a custom time range to further refine the search results.
Last Modified	Select a time range when the file was last modified. You can also specify a custom time range to further refine the search results.
Last Accessed	Select a time range when the file, or directory (CIFS or NFS only), was last accessed. You can also specify a custom time range to further refine the search results. For the types of files that BlueXP classification scans, this is the last time BlueXP classification scanned the file.  BlueXP classification does not extract the "last accessed time" from the following data sources: SharePoint Online, SharePoint On-premises (SharePoint Server), OneDrive, Google Drive, and Amazon S3.

### Filter data by metadata

Use the following filters to view data based on location, size, and directory or file type.

Filter	Details
File Path	Enter up to 20 partial or full paths that you want to include or exclude from the query. If you enter both include paths and exclude paths, BlueXP classification finds all files in the included paths first, then it removes files from excluded paths, and then it displays the results. Note that using "*" in this filter has no effect, and that you can't exclude specific folders from the scan - all the directories and files under a configured share will be scanned.
Directory Type	Select the directory type; either "Share" or "Folder".
File Type	Select the <a href="#">types of files</a> .
File Size	Select the file size range.
File Hash	Enter the file's hash to find a specific file, even if the name is different.

### Filter data by storage type

Use the following filters to view data by storage type.

Filter	Details
Working Environment Type	Select the type of working environment. OneDrive, SharePoint, and Google Drive are categorized under "Apps".
Working Environment name	Select specific working environments.
Storage Repository	Select the storage repository, for example, a volume or a schema.

### Filter data by saved searches

Use the following filter to view data by saved searches.



Filter	Details
Saved search	Select one saved search or multiples. Go to the <a href="#">saved searches tab</a> to view the list of existing saved searches and create new ones.

### Filter data by analysis status

Use the following filter to view data by the BlueXP classification scan status.

Filter	Details
Analysis Status	Select an option to show the list of files that are Pending First Scan, Completed being scanned, Pending Rescan, or that have Failed to be scanned.
Scan Analysis Event	Select whether you want to view files that were not classified because BlueXP classification couldn't revert last accessed time, or files that were classified even though BlueXP classification couldn't revert last accessed time.

See [details about the "last accessed time" timestamp](#) for more information about the items that appear in the Investigation page when filtering using the Scan Analysis Event.

### Filter data by duplicates

Use the following filter to view files that are duplicated in your storage.


Filter	Details
Duplicates	Select whether the file is duplicated in the repositories.

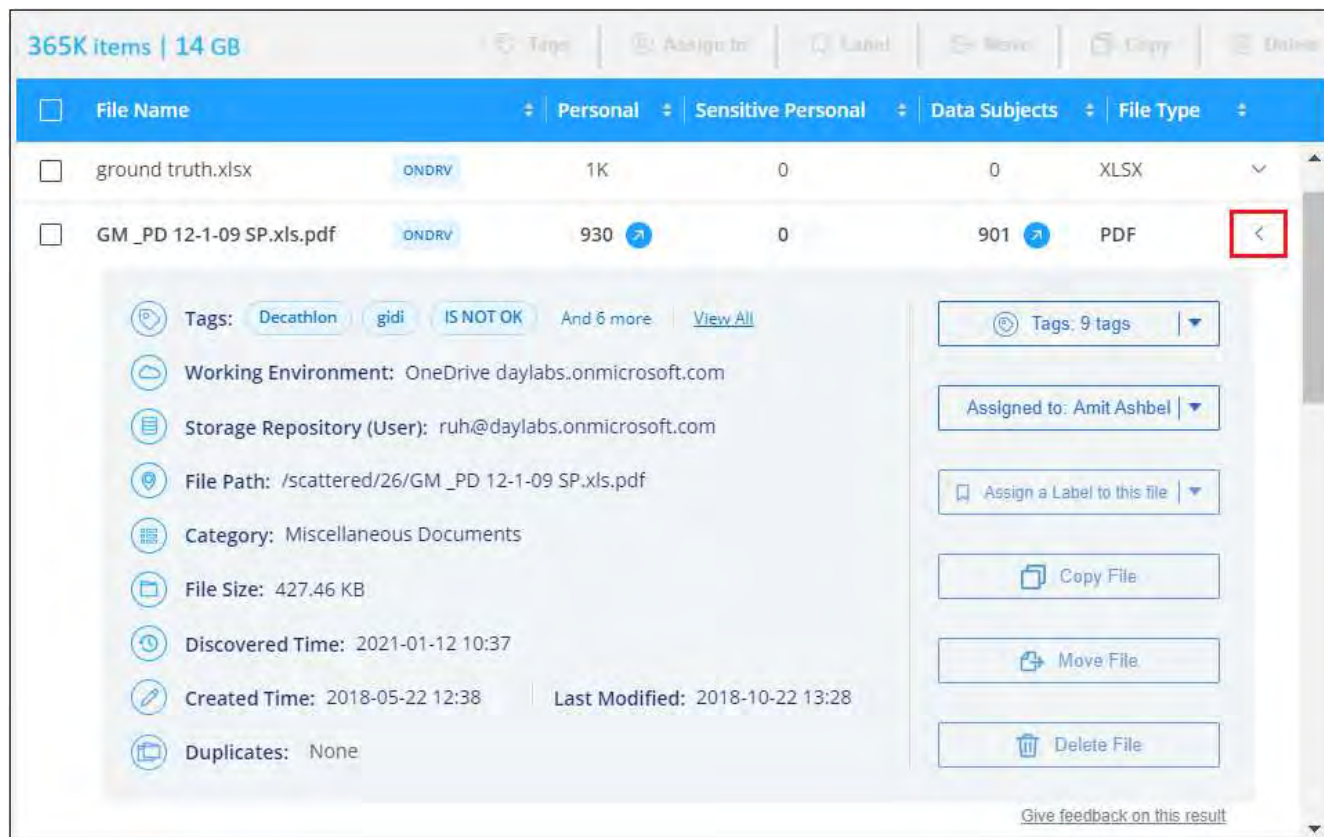
### View file metadata

In addition to showing you the working environment and volume where the file resides, the metadata shows much more information, including the file permissions, file owner, and whether there are duplicates of this file. This information is useful if you're planning to [create saved searches](#) because you can see all the information that you can use to filter your data.

The availability of information depends on the data source. For example, volume name and permissions are not shared for database files.

#### Steps

1. From the BlueXP classification menu, select **Investigation**.
2. In the Data Investigation list on the right, select the down-caret  on the right for any single file to view the file metadata.




## View users' permissions for files and directories

To view a list of all users or groups who have access to a file or to a directory and the types of permissions they have, select **View all Permissions**. This button is available only for data in CIFS shares.


Note that if you see SIDs (Security IDentifiers) instead of user and group names, you should integrate your Active Directory into BlueXP classification. [See how to do this](#).

### Steps

1. From the BlueXP classification menu, select **Investigation**.
2. In the Data Investigation list on the right, select the down-caret  on the right for any single file to view the file metadata.
3. To view a list of all users or groups who have access to a file or to a directory and the types of permissions they have, in the Open Permissions field, select **View all Permissions**.



BlueXP classification shows up to 100 users in the list.



4. Select the down-caret  button for any group to see the list of users who are part of the group.



You can expand one level of the group to see the users who are part of the group.

5. Select the name of a user or group to refresh the Investigation page so you can see all the files and directories that the user or group has access to.

## Check for duplicate files in your storage systems

You can check whether duplicate files are being stored in your storage systems. This is useful if you want to identify areas where you can save storage space. It's also good to ensure certain files that have specific permissions or sensitive information are not unnecessarily duplicated in your storage systems.

All of your files (not including databases) that are 1 MB or larger, or that contain personal or sensitive personal information, are compared to see if there are duplicates.

BlueXP classification uses hashing technology to determine duplicate files. If any file has the same hash code as another file, you can be 100% sure that the files are exact duplicates—even if the file names are different.


### Steps

1. From the BlueXP classification menu, select **Investigation**.
2. In the Investigation page Filters pane on the left, select "File Size" along with "Duplicates" ("Has duplicates") to see which files of a certain size range are duplicated in your environment.
3. Optionally, download the list of duplicate files and send it to your storage admin so they can decide which files, if any, can be deleted.
4. Optionally, [delete the file](#) yourself if you are confident that a specific version of the file is not needed.

### View if a specific file is duplicated

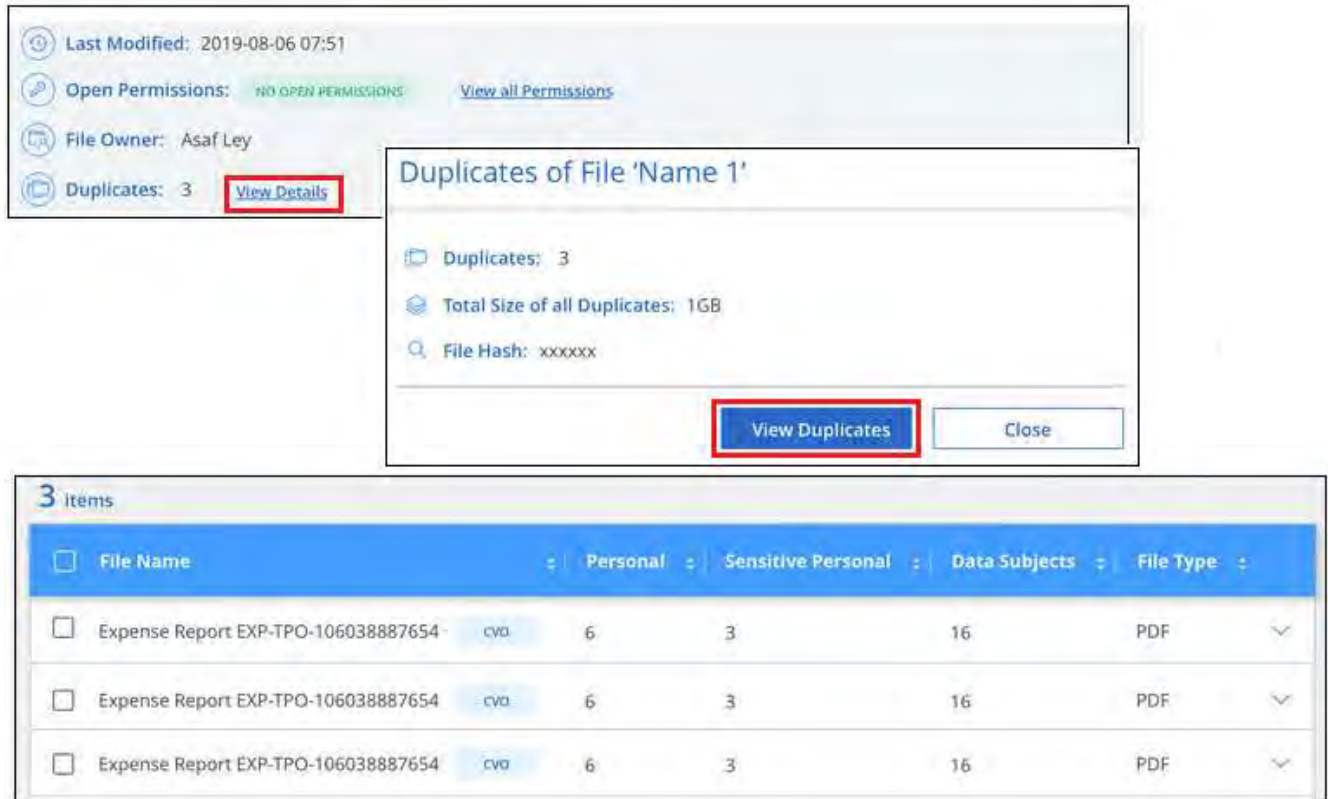
You can see if a single file has duplicates.

## Steps

1. From the BlueXP classification menu, select **Investigation**.
2. In the Data Investigation list, select  on the right for any single file to view the file metadata.

If duplicates exist for a file, this information appears next to the *Duplicates* field.

3. To view the list of duplicate files and where they are located, select **View Details**.
4. In the next page select **View Duplicates** to view the files in the Investigation page.



The screenshot shows the BlueXP interface. At the top, file metadata is displayed: Last Modified: 2019-08-06 07:51, Open Permissions: NO OPEN PERMISSIONS (with a link to View all Permissions), File Owner: Asaf Ley, and Duplicates: 3 (with a red box around the 'View Details' button). Below this, a modal window titled 'Duplicates of File 'Name 1'' is open, showing Duplicates: 3, Total Size of all Duplicates: 1GB, and File Hash: xxxxxx (with a red box around the 'View Duplicates' button). At the bottom, a table titled '3 Items' lists three duplicate files:

<input type="checkbox"/>	File Name		Personal	Sensitive Personal	Data Subjects	File Type	
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cva	6	3	16	PDF	▼
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cva	6	3	16	PDF	▼
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cva	6	3	16	PDF	▼



You can use the "file hash" value provided in this page and enter it directly in the Investigation page to search for a specific duplicate file at any time - or you can use it in a saved search.

## Create the Data Investigation Report

The Data Investigation Report is a download of the filtered contents of the Data Investigation page.

The report is available as a CSV or JSON file you can save to your local machine.

There can be up to three report files downloaded if BlueXP classification is scanning files (unstructured data), directories (folders and file shares), and databases (structured data).

The files are split into files with a fixed number of rows or records:

- JSON - 100,000 records per report that takes about 5 minutes to generate
- CSV - 200,000 records per report that takes about 4 minutes to generate



You can download a version of the CSV file to view in this browser. This version is limited to 10,000 records.

## What's included in the Data Investigation Report

The **Unstructured Files Data Report** includes the following information about your files:

- File name
- Location type
- Working environment name
- Storage repository (for example, a volume, bucket, shares)
- Repository type
- File path
- File type
- File size (in MB)
- Created time
- Last modified
- Last accessed
- File owner
  - File owner data encompasses account name, SAM account name, and e-mail address when Active Directory is configured.
- Category
- Personal information
- Sensitive personal information
- Open permissions
- Scan Analysis Error
- Deletion detection date

The deletion detection date identifies the date that the file was deleted or moved. This enables you to identify when sensitive files have been moved. Deleted files don't contribute to the file number count that appears in the dashboard or on the Investigation page. The files only appear in the CSV reports.

The **Unstructured Directories Data Report** includes the following information about your folders and file shares:


- Working environment type
- Working environment name
- Directory name
- Storage repository (for example, a folder or file shares)
- Directory owner
- Created time
- Discovered time

- Last modified
- Last accessed
- Open permissions
- Directory type

The **Structured Data Report** includes the following information about your database tables:

- DB Table name
- Location type
- Working environment name
- Storage repository (for example, a schema)
- Column count
- Row count
- Personal information
- Sensitive personal information

#### **Steps to generate the report**

1. From the Data Investigation page, select the  button on the top, right of the page.
2. Choose the report type: CSV or JSON.
3. Enter a **Report name**.
4. To download the complete report, select **Working environment** then choose the **Working Environment** and **Volume** from the respective dropdown menus. Provide a **Destination folder path**.

To download the report in the browser, select **Local** . Note this option limits the report to the first 10,000 rows and is limited to the **CSV** format. You don't need to complete any other fields if you select **Local**.

5. Select **Download Report**.

## Download Investigation Report

### Report type

CSV file       JSON file

### Report name

investigation\_report

### Export destination

Working environment       Local (limited to 10K rows)

Working environment

Volume:

Destination folder path

Download Report

Cancel

### Result

A dialog displays a message that the reports are being downloaded.

## Create a saved search based on selected filters

You can create a saved search for frequently used search filters in the Data Investigation page to easily replicate those search queries.

### Steps

1. From the BlueXP classification menu, select **Investigation**.
2. On the Data Investigation page, select the filters you want to use to create a saved search.
3. At the bottom of the Filter pane, select **Create saved search from this search**.
4. Enter a name and a description for the saved search.
5. Choose any of the following:
6. Select **Create Saved Search**.



It might take up to 15 minutes for the results to appear on the Saved Searches page.



# Manage saved searches with BlueXP classification

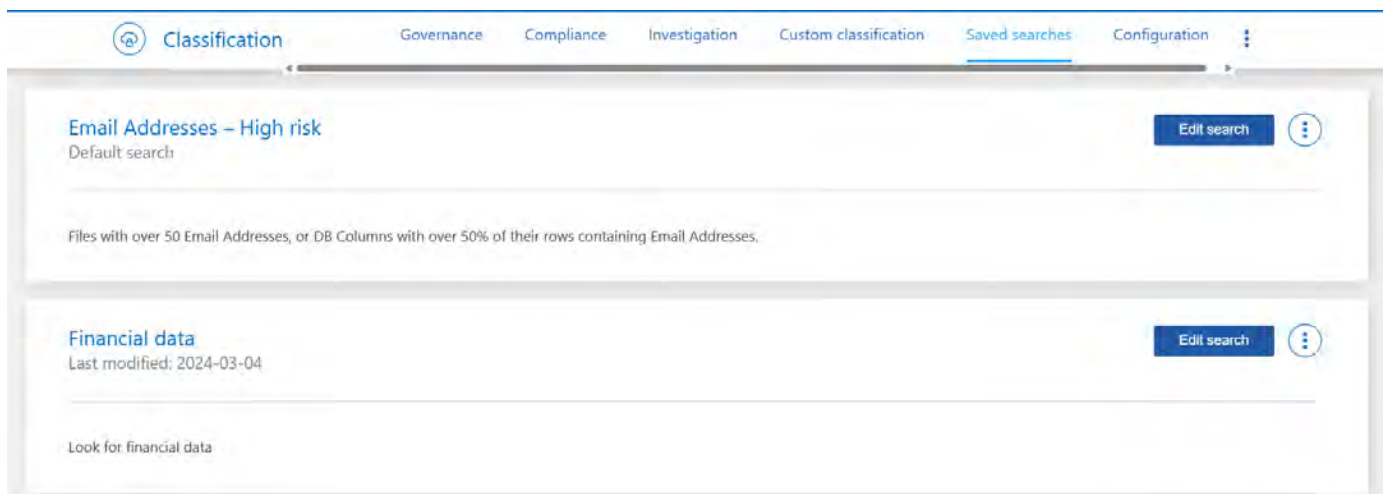
BlueXP classification supports saving your search queries. With a saved search, you can create custom filters to sort through frequent queries of your data Investigation page. BlueXP classification also includes predefined saved searches based on common requests.




In versions of BlueXP classification earlier than 1.43, saved searches were called [policies](#).


The **Saved searches** tab in the Compliance Dashboard lists all the predefined and custom saved searches available on this instance of BlueXP classification.

Saved searches also appear in the list of filters in the Investigation page.



## View saved searches results in the Investigation page

To display the results for a saved search in the Investigation page, select the  button for a specific search then select **Investigate Results**.



## Create custom saved searches

You can create your own custom saved searches that provide results for queries specific to your organization. Results are returned for all files and directories (shares and folders) that match the search criteria.

### Steps

1. In the Investigation tab, define a search by selecting the filters you want to use. See [Filtering data in the Investigation page](#) for details.
2. Once you have all the filter characteristics set to your liking, select **Create saved search**.



## ☰ Data Investigation

☰ FILTERS: Clear All

---

Storage Repository 3 +

---

File / Directory Path +

---

Category +

---

Sensitivity Level +

---

Save this search

3. Name the saved search and add a description. The name must be unique.
4. Select **Create Saved Search**.

## Create search

---

This will save the current selected filters and search term as a saved search. You can view or delete this later from the "Saved searches" tab.

Note it may take up to 15 minutes for results to be displayed for a new saved search.

Name this search

Give it a detailed description that explains what it searches for

---

Create search

Cancel

Once you've created the search, you can view it in the **Saved searches** tab.

### Edit saved searches

You can modify the query criteria for a saved search (that is, the defined filters) to add or remove certain parameters.

You cannot modify default saved searches.

#### Steps

1. From the Saved searches page, select **Edit Search** for the search that you want to change.

Sensitive data

Last modified: 2024-03-04

Edit search

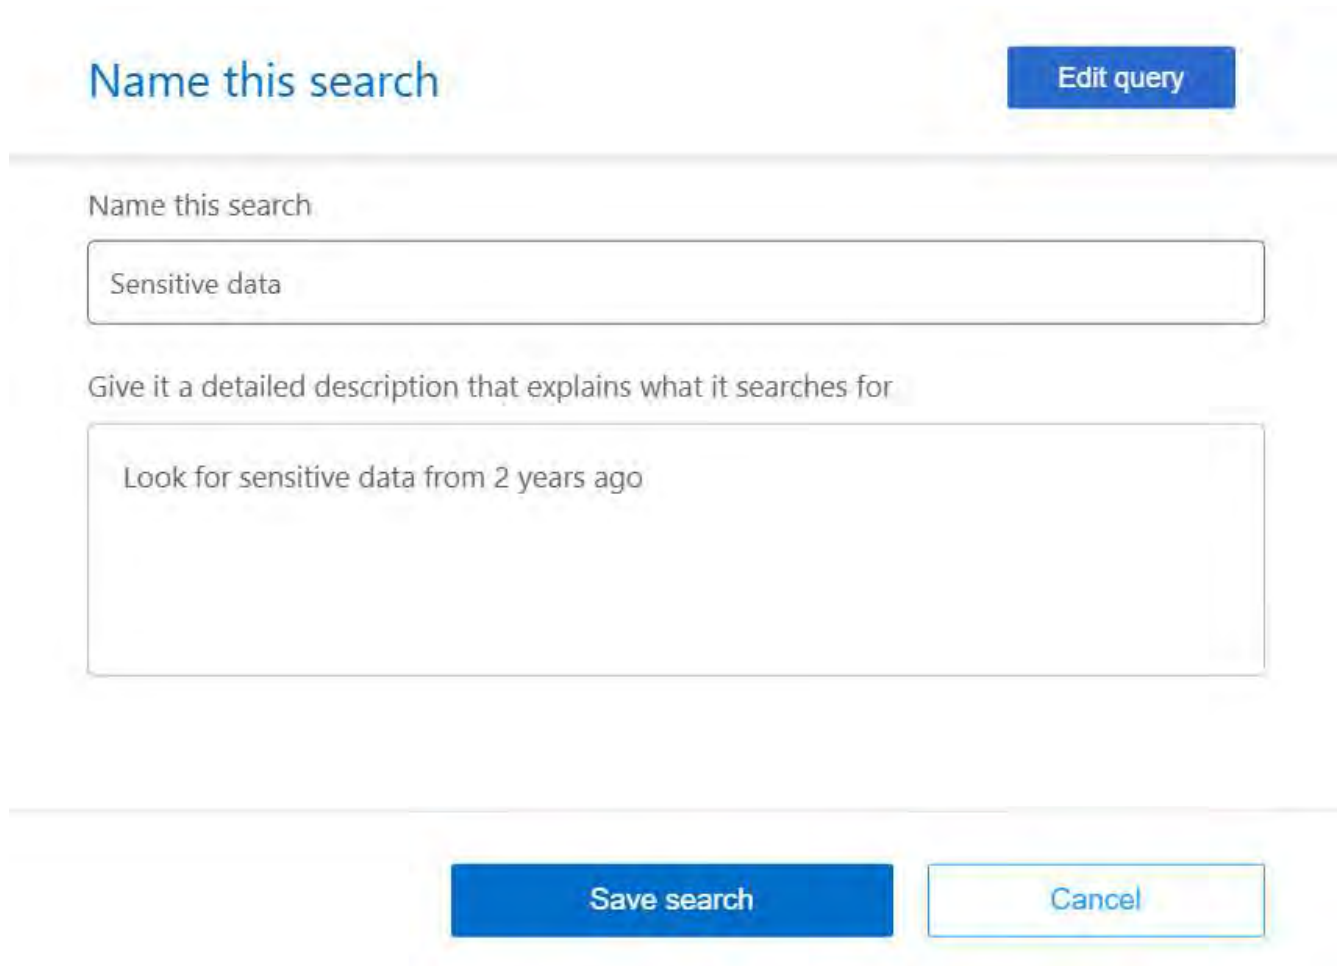


---

Look for sensitive data from 2 years ago

2. Make the changes to the name and description fields. To only change the name and description fields, select **Save search**.

To change the filters for the saved search, select **Edit query**.



Name this search

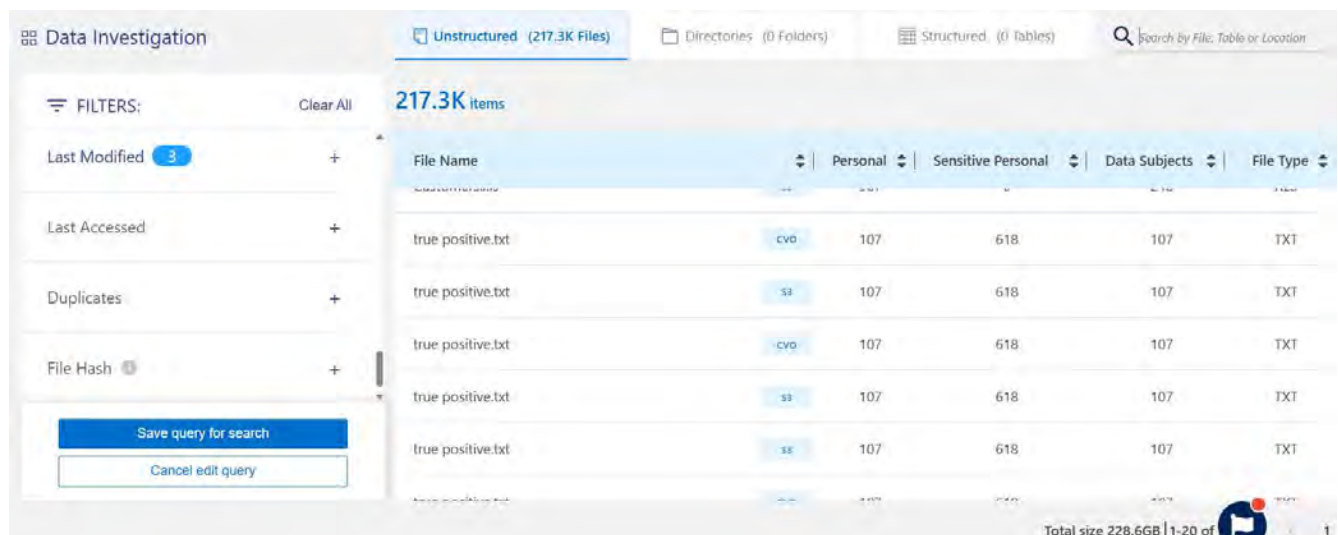
Sensitive data

Give it a detailed description that explains what it searches for

Look for sensitive data from 2 years ago

Save search Cancel

3. In the Investigation page, edit the query. You can add, remove, or modify filters. To complete your changes, select **Save query for this search**.



Data Investigation

Unstructured (217.3K Files) Directories (0 Folders) Structured (0 Tables) Search by File, Table or Location

FILTERS: Clear All 217.3K items


File Name	Personal	Sensitive Personal	Data Subjects	File Type	
true positive.txt	cvo	107	618	107	TXT
true positive.txt	ss	107	618	107	TXT
true positive.txt	cvo	107	618	107	TXT
true positive.txt	ss	107	618	107	TXT
true positive.txt	ss	107	618	107	TXT
true positive.txt	ss	107	618	107	TXT

Save query for search Cancel edit query

Total size 228.6GB | 1-20 of

## Delete saved searches

You can delete any custom saved search if you no longer need it. You can't delete default saved searches.

To delete a saved search, select the  button for a specific search, select **Delete search**, then select **Delete search** again in the confirmation dialog.

## Default searches

BlueXP classification provides the following system-defined search queries:

- **Data Subject names - High risk**

Files with more than 50 data subject names

- **Email Addresses - High risk**

Files with more than 50 email addresses or database columns with more than 50% of their rows containing email addresses

- **Personal data - High risk**

Files with more than 20 personal data identifiers or database columns with more than 50% of their rows containing personal data identifiers

- **Private data - Stale over 7 years**

Files containing personal or sensitive personal information, last modified more than 7 years ago

- **Protect - High**

Files or database columns that contain a password, credit card information, IBAN number, or social security number

- **Protect - Low**

Files that have not been accessed for more than 3 years

- **Protect - Medium**

Files that contain files or database columns with personal data identifiers including ID numbers, tax identification numbers, drivers license numbers, medicinal IDs, or passport numbers

- **Sensitive Personal data - High risk**

Files with more than 20 sensitive personal data identifiers or database columns with greater than 50% of their rows containing sensitive personal data

## Change the BlueXP classification scan settings for your repositories

You can manage how your data is being scanned in each of your working environments and data sources. You can make the changes on a "repository" basis; meaning you can

make changes for each volume, schema, user, etc. depending on the type of data source you are scanning.

Some of the things you can change are whether a repository is scanned or not, and whether BlueXP classification is performing a [mapping scan](#) or a [mapping & classification scan](#). You can also pause and resume scanning, for example, if you need to stop scanning a volume for a period of time.

## View the scan status for your repositories

You can view the individual repositories that BlueXP classification is scanning (volumes, buckets, etc.) for each working environment and data source. Additionally, you can see how many have been "Mapped", and how many have been "Classified". Classification takes a longer time as the full AI identification is being performed on all data.

You can view the scanning status of each work environment on the Configuration page:

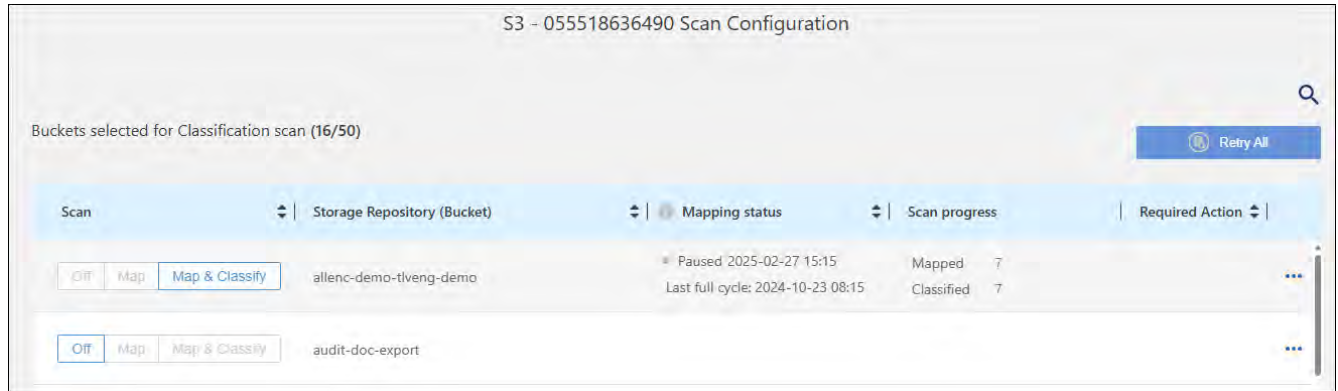
- **Initializing** (light blue dot): The map or classify configuration is activated. This appears for few seconds before starting the "pending queue" status.
- **Pending queue** (orange dot): The scan task is waiting to be listed in the scanning queue.
- **Queued** (orange dot): The task was successfully added to the scanning queue. The system will start mapping or classifying the volume when its turn in the queue arrives.
- **Running** (green dot): The scan task, which was in the queue, is actively in progress on the selected storage repository.
- **Finished** (green dot): The scan of the storage repository is complete.
- **Paused** (gray dot): You selected the "Pause" option to pause scanning. While the changes in the volume are not displayed in the system, the scanned insights are still shown.
- **Error** (red dot): The scan cannot complete because it has encountered issues. If you need to complete an action, the error appears in the tooltip under the "Required action" column. Otherwise, the system shows an "error" status and tries to recover. When it finishes, the status changes.
- **Not scanning**: The volume configuration of "Off" was selected and the system is not scanning the volume.

## Steps

1. From the BlueXP classification menu, select **Configuration**.

The screenshot shows the 'Identity Services' configuration page in the BlueXP interface. The page is titled 'Identity Services' and includes a 'Quick Navigation' sidebar with options for 'Identity Services', 'Working Environments', and 'Scanner Groups'. The main content area displays the configuration for a working environment named 'share2scan.netapp.com'. Below this, there are 11 'Working Environments' listed, with filters for 'S3', 'CVO', 'DB', and 'SHARES'. The selected environment is 'S3 - 055518636490 | 50 Buckets', with a 'Scanner Group name: default' and 'Working Environment ID: S3'. A 'Scan Mode' section shows a progress bar and a status of '16 Classified', '16 Mapped', and '34 Not Scanned'. A note indicates 'Continuously scanning all selected Buckets'. The page also features 'Active Directory Integrated' and 'Add Working Environment' buttons at the top right.

2. From the Configuration tab, select the **Configuration** button for the working environment.
3. In the Scan Configuration page, view the scan settings for all repositories.



4. Hover your cursor over the chart in the *Mapping Status* column to see the number of files that remain to be mapped or classified in each repository (bucket in this example).

## Change the type of scanning for a repository

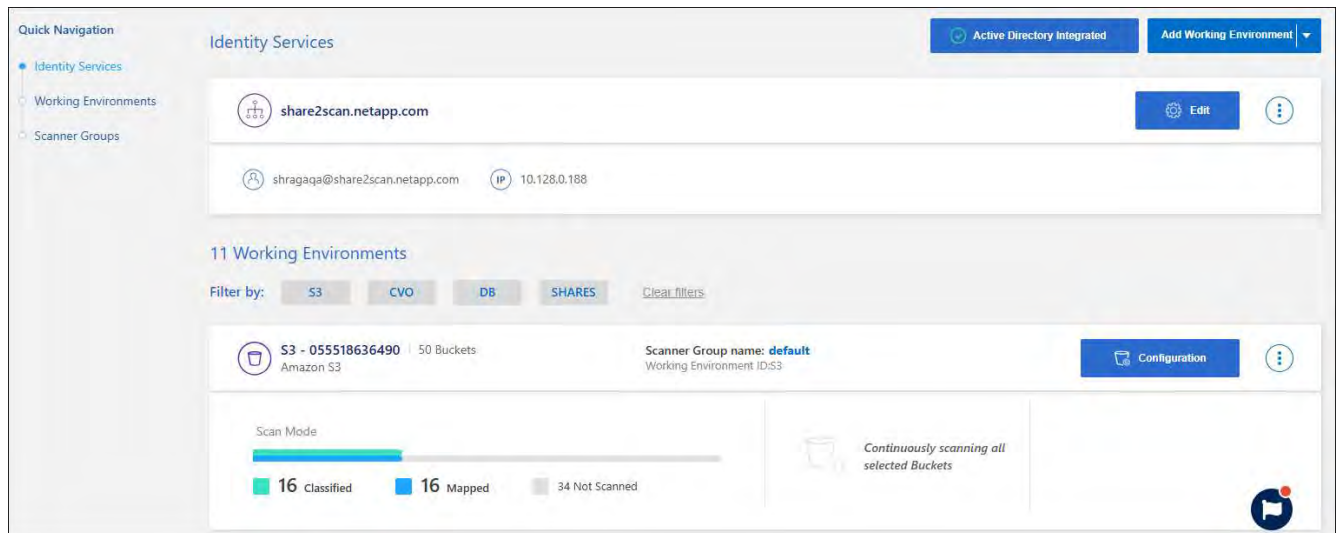
You can start or stop mapping-only scans, or mapping and classification scans, in a working environment at any time from the Configuration page. You can also change from mapping-only scans to mapping and classification scans, and vice-versa.




Databases can't be set to mapping-only scans. Database scanning can be Off or On; where On is equivalent to Map & Classify.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration tab, select the **Configuration** button for the working environment.




3. In the Scan Configuration page, change any of the repositories (buckets in this example) to perform **Map** or **Map & Classify** scans.



Certain types of working environments enable you to change the type of scanning globally for all repositories using a button bar at the top of the page. This is valid for Cloud Volumes ONTAP, on-premises ONTAP, Azure NetApp Files, and Amazon FSx for ONTAP systems.

The example below shows this button bar for an Azure NetApp Files system.



## Prioritize scans

You can prioritize the most important mapping-only scans or map & classify scans to ensure high priority scans are completed first.

By default, scans are queued based on the order in which they are initiated. With the ability to prioritize scans, you can move scans to the front of the queue. Multiple scans can be prioritized. Priority is designated in a first-in, first-out order, meaning the first scan you prioritize moves to the front of the queue; the second scan you prioritize becomes second in the queue, and so forth.

Priority is granted on a one-time basis. Automatic rescans of mapping data occur in the default order.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. Select the resources you want to prioritize.
3. From the Actions ... option, select **Prioritize scan**.


## Stop scanning for a repository

You can stop scanning a repository (for example, a volume) if you no longer need to monitor it for compliance. You do this by turning scanning "off". When scanning is turned off, all the indexing and information about that volume is removed from the system, and charging for scanning the data is stopped.


### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration tab, select the **Configuration** button for the working environment.





3. In the Scan Configuration page select **Off** to stop scanning for a particular bucket.



## Pause and resume scanning for a repository


You can "pause" scanning on a repository if you want to temporarily stop scanning certain content. Pausing scanning means that BlueXP classification won't perform any future scans for changes or additions to the repository, but that all the current results will still be displayed in the system. Pausing scanning does not stop charging for the scanned the data because the data still exists.

You can "resume" scanning at any time.

### Steps

1. From the BlueXP classification menu, select **Configuration**.
2. From the Configuration tab, select the **Configuration** button for the working environment.





3. In the Scan Configuration page, select the Actions **...** icon.
4. Select **Pause** to pause scanning for a volume, or select **Resume** to resume scanning for a volume that had been previously paused.

## View BlueXP classification compliance reports

BlueXP classification provides reports that you can use to better understand the status of your organization's data privacy program.

By default, the BlueXP classification dashboards display compliance and governance data for all working environments, databases, and data sources. If you want to view reports that contain data for only some of the working environments, you can filter to see just them.



- The reports described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan can only generate the Data Mapping Report.
- NetApp can't guarantee 100% accuracy of the personal data and sensitive personal data that BlueXP classification identifies. You should always validate the information by reviewing the data.

The following reports are available for BlueXP classification:

- **Data Discovery Assessment report:** Provides a high-level analysis of the scanned environment to highlight the system's findings and to show areas of concern and potential remediation steps.
- **Data Mapping report:** Provides information about the size and number of files in your working environments. This includes usage capacity, age of data, size of data, and file types.
- **Data Subject Access Request report:** Enables you to extract a report of all files that contain information regarding a data subject's specific name or personal identifier.
- **HIPAA report:** Helps you identify the distribution of health information across your files.
- **PCI DSS report:** Helps you identify the distribution of credit card information across your files.
- **Privacy Risk Assessment report:** Provides privacy insights from your data and a privacy risk score.
- **Reports on a specific information type:** Reports are available that include details about the identified files that contain personal data and sensitive personal data. You can also see files broken down by

category and file type.


## Select the working environments for reports

You can filter the contents of the BlueXP classification Compliance dashboard to see compliance data for all working environments and databases, or for just specific working environments.

When you filter the dashboard, BlueXP classification scopes the compliance data and reports to just those working environments that you selected.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. Select the Working environments filter drop-down and select the working environments.
3. Select **View**.



## Data Subject Access Request Report

Privacy regulations such as the European GDPR grant data subjects (such as customers or employees) the right to access their personal data. When a data subject requests this information, this is known as a DSAR (data subject access request). Organizations are required to respond to these requests "without undue delay", and at the latest within one month of receipt.

You can respond to a DSAR by searching for a subject's full name or known identifier (such as an email

address) and then downloading a report. The report is designed to aid in your organization's requirement to comply with GDPR or similar data privacy laws.

### How can BlueXP classification help you respond to a DSAR?

When you perform a data subject search, BlueXP classification finds all of the files, buckets, OneDrive, and SharePoint accounts that have that person's name or identifier in it. BlueXP classification checks the latest pre-indexed data for the name or identifier. It doesn't initiate a new scan.

After the search is complete, you can then download the list of files for a Data Subject Access Request report. The report aggregates insights from the data and puts it into legal terms that you can send back to the person.



Data subject search is not supported within databases at this time.

### Search for data subjects and download reports

Search for the data subject's full name or known identifier and then download a file list report or DSAR report. You can search by [any personal information type](#).




English, German, Japanese, and Spanish are supported when searching for the names of data subjects. Support for more languages will be added later.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. From the Compliance page, scroll down and select **Data Subjects**.
3. Search for the data subject's full name or known identifier.

Here's an example that shows a search for the name *john doe*:



4. Choose one of the available options:

- **Download DSAR Report:** A formal response to the access request that you can send to the data subject. This report contains automatically-generated information based on data that BlueXP classification found on the data subject and is designed to be used as a template. You should complete the form and review it internally before sending it to the data subject.
- **Investigate Results:** A page that enables you to investigate the data by searching, sorting, expanding details for a specific file, and by downloading the file list.



If there are more than 10,000 results, only the top 10,000 appear in the file list.

## Health Insurance Portability and Accountability Act (HIPAA) Report

The Health Insurance Portability and Accountability Act (HIPAA) Report can help you identify files containing health information. It is designed to aid in your organization's requirement to comply with HIPAA data privacy laws. The information BlueXP classification looks for includes:

- Health reference pattern
- ICD-10-CM Medical code
- ICD-9-CM Medical code
- HR - Health category
- Health Application Data category

The report includes the following information:

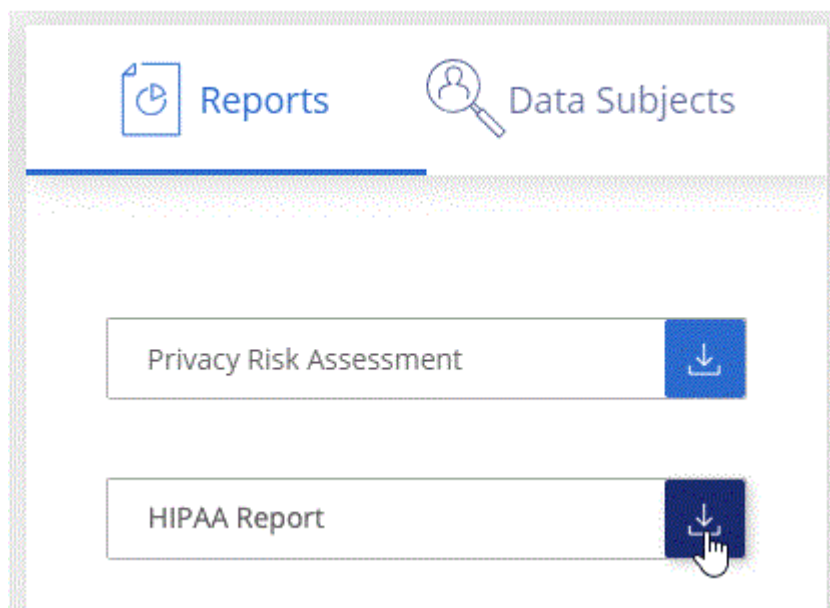
- Overview: How many files contain health information and in which working environments.
- Encryption: The percentage of files containing health information that are on encrypted or unencrypted working environments. This information is specific to Cloud Volumes ONTAP.
- Ransomware Protection: The percentage of files containing health information that are on working environments that do or don't have ransomware protection enabled. This information is specific to Cloud Volumes ONTAP.
- Retention: The timeframe in which the files were last modified. This is helpful because you shouldn't keep health information for longer than you need to process it.
- Distribution of Health Information: The working environments where the health information was found and whether encryption and ransomware protection are enabled.

## Generate the HIPAA Report

Go to the Compliance tab to generate the report.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. Scroll down and locate the **Reports** pane.
3. Select the download icon next to **HIPAA Report.**



### Result

BlueXP classification generates a PDF report that you can review and send to other groups as needed.

## Payment Card Industry Data Security Standard (PCI DSS) Report

The Payment Card Industry Data Security Standard (PCI DSS) Report can help you identify the distribution of credit card information across your files.

The report includes the following information:

- Overview: How many files contain credit card information and in which working environments.

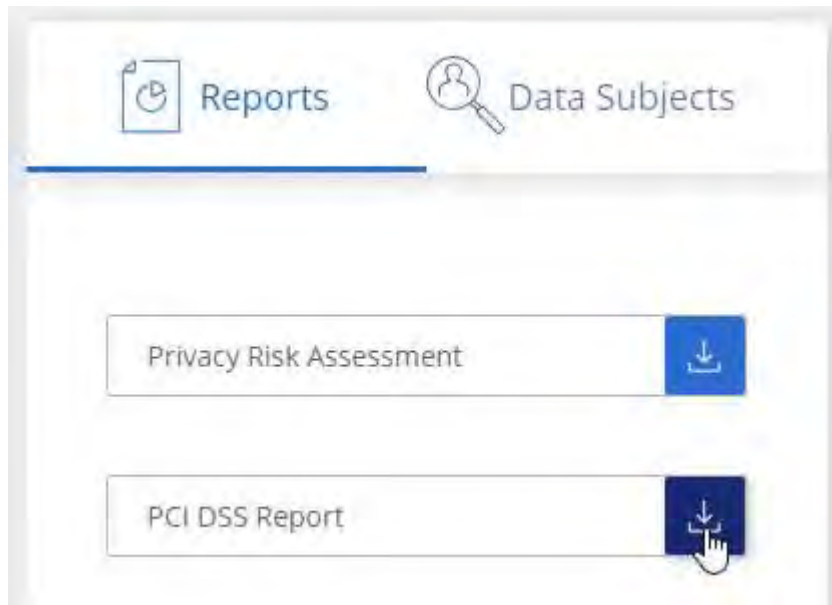
- **Encryption:** The percentage of files containing credit card information that are on encrypted or unencrypted working environments. This information is specific to Cloud Volumes ONTAP.
- **Ransomware Protection:** The percentage of files containing credit card information that are on working environments that do or don't have ransomware protection enabled. This information is specific to Cloud Volumes ONTAP.
- **Retention:** The timeframe in which the files were last modified. This is helpful because you shouldn't keep credit card information for longer than you need to process it.
- **Distribution of Credit Card Information:** The working environments where the credit card information was found and whether encryption and ransomware protection are enabled.

## Generate the PCI DSS Report

Go to the Compliance tab to generate the report.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. Scroll down and locate the **Reports** pane.
3. Select the download icon next to **PCI DSS Report**.



### Result

BlueXP classification generates a PDF report that you can review and send to other groups as needed.

## Privacy Risk Assessment Report

The Privacy Risk Assessment Report provides an overview of your organization's privacy risk status, as required by privacy regulations such as GDPR and CCPA.

The report includes the following information:

- **Compliance status:** A severity score and the distribution of data, whether it's non-sensitive, personal, or sensitive personal.
- **Assessment overview:** A breakdown of the types of personal data found, as well as the categories of data.


- Data subjects in this assessment: The number of people, by location, for which national identifiers were found.

## Generate the Privacy Risk Assessment Report

Go to the Compliance tab to generate the report.

### Steps

1. From the BlueXP classification menu, select "**Compliance.**"
2. Scroll down and locate the **Reports** pane.
3. Select the download icon next to **Privacy Risk Assessment.**



### Result

BlueXP classification generates a PDF report that you can review and send to other groups as needed.

### Severity score

BlueXP classification calculates the severity score for the Privacy Risk Assessment Report on the basis of three variables:

- The percentage of personal data out of all data.
- The percentage of sensitive personal data out of all data.
- The percentage of files that include data subjects, determined by national identifiers such as national IDs, Social Security numbers, and tax ID numbers.

The logic used to determine the score is as follows:

Severity score	Logic
0	All three variables are exactly 0%
1	One of the variables are larger than 0%
2	One of the variables are larger than 3%

<b>Severity score</b>	<b>Logic</b>
3	Two of the variables are larger than 3%
4	Three of the variables are larger than 3%
5	One of the variables are larger than 6%
6	Two of the variables are larger than 6%
7	Three of the variables are larger than 6%
8	One of the variables are larger than 15%
9	Two of the variables are larger than 15%
10	Three of the variables are larger than 15%



# Manage BlueXP classification

## Exclude specific directories from BlueXP classification scans

If you want BlueXP classification to exclude scanning data that resides in certain data source directories, you can add these directory names to a configuration file. After you apply this change, the BlueXP classification engine will exclude scanning data in those directories.

Note that BlueXP classification is configured by default to exclude scanning volume snapshot data because that content is identical to the content in the volume.

This functionality is available in BlueXP classification version 1.29 and greater (starting in March 2024).

### Supported data sources

Excluding specific directories from BlueXP classification scans is supported for NFS and CIFS shares in the following data sources:

- On-premises ONTAP
- Cloud Volumes ONTAP
- Amazon FSx for NetApp ONTAP
- Azure NetApp Files
- General file shares

### Define the directories to exclude from scanning

Before you can exclude directories from classification scanning, you need to log into the BlueXP classification system so you can edit a configuration file and run a script. See how to [log in to the BlueXP classification system](#) depending on whether you manually installed the software on a Linux machine or if you deployed the instance in the cloud.



- You can exclude a maximum of 50 directory paths per BlueXP classification system.
- Excluding directory paths may affect scanning times.

### Steps

1. On the BlueXP classification system, go to `/opt/netapp/config/custom_configuration` and open the file `data_provider.yaml`.
2. In the `"data_providers"` section, under the line `"exclude:"`, enter the directory paths to exclude. For example:

```
exclude:  
- "folder1"  
- "folder2"
```

Do not change anything else in this file.

3. Save the changes to the file.

4. Go to `/opt/netapp/Datasense/tools/customer_configuration/data_providers` and run the following script:

```
update_data_providers_from_config_file.sh
```

This command commits the directories to be excluded from scanning to the classification engine.

## Result

All subsequent scans of your data will exclude scanning of those specified directories.

You can add, edit, or delete items from the exclude list using these same steps. The revised exclude list will be updated after you run the script to commit your changes.

## Examples

### Configuration 1:

Every folder that contains "folder1" anywhere in the name will be excluded from all data sources.

```
data_providers:  
  exclude:  
    - "folder1"
```

### Expected results for paths that will be excluded:

- /CVO1/folder1
- /CVO1/folder1name
- /CVO1/folder10
- /CVO1/\*folder1
- /CVO1/+folder1name
- /CVO1/notfolder10
- /CVO22/folder1
- /CVO22/folder1name
- /CVO22/folder10

### Examples for paths that will not be excluded:

- /CVO1/\*folder
- /CVO1/foldername
- /CVO22/\*folder20

### Configuration 2:

Every folder that contains "folder1" only at the start of the name will be excluded.

```
data_providers:
  exclude:
    - "\\*folder1"
```

**Expected results for paths that will be excluded:**

- /CVO/\*folder1
- /CVO/\*folder1name
- /CVO/\*folder10

**Examples for paths that will not be excluded:**

- /CVO/folder1
- /CVO/folder1name
- /CVO/not\*folder10

**Configuration 3:**

Every folder in data source "CVO22" that contains "folder1" anywhere in the name will be excluded.

```
data_providers:
  exclude:
    - "CVO22/folder1"
```

**Expected results for paths that will be excluded:**

- /CVO22/folder1
- /CVO22/folder1name
- /CVO22/folder10

**Examples for paths that will not be excluded:**

- /CVO1/folder1
- /CVO1/folder1name
- /CVO1/folder10

## Escaping special characters in folder names

If you have a folder name that contains one of the following special characters and you want to exclude data in that folder from being scanned, you'll need to use the escape sequence `\\` before the folder name.

```
., +, *, ?, ^, $, (, ), [, ], {, }, |
```

For example:

Path in source: `/project/*not_to_scan`

Syntax in exclude file: `"\\*not_to_scan"`

## View the current exclusion list

It's possible for the contents of the `data_provider.yaml` configuration file to be different than what has actually been committed after running the `update_data_providers_from_config_file.sh` script. To view the current list of directories that you've excluded from BlueXP classification scanning, run the following command from `/opt/netapp/Datasense/tools/customer_configuration/data_providers`:

```
get_data_providers_configuration.sh
```

## Define additional group IDs as open to organization in BlueXP classification

When group IDs (GIDs) are attached to files or folders in NFS file shares they define the permissions for the file or folder; such as whether they are "open to the organization". If some group IDs (GIDs) are not initially set up with the "Open to Organization" permission level, you can add that permission to the GID so that any files and folders that have that GID attached will be deemed "open to the organization".

After you make this change and BlueXP classification rescans your files and folders, any files and folders that have these group IDs attached will show this permission in the Investigation Details page, and they'll also appear in reports where you are displaying file permissions.

To activate this functionality, you need to log into the BlueXP classification system so you can edit a configuration file and run a script. See how to [log in to the BlueXP classification system](#) depending on whether you manually installed the software on a Linux machine or if you deployed the instance in the cloud.

### Add the "open to organization" permission to group IDs

You need to have the group ID numbers (GIDs) before starting this task.

#### Steps

1. On the BlueXP classification system, go to `/opt/netapp/config/custom_configuration` and open the file `data_provider.yaml`.
2. In the line `organization_group_ids: []` add the group IDs. For example:

```
organization_group_ids: [1014, 1015, 21, 2021, 1013, 2020, 1018, 1019]
```

Do not change anything else in this file.

3. Save the changes to the file.
4. Go to `/opt/netapp/Datasense/tools/customer_configuration/data_providers` and run the following script:

```
update_data_providers_from_config_file.sh
```

This command commits the revised group ID permissions to the classification engine.

## Result

All subsequent scans of your data will identify files or folders that have these group IDs attached as "open to organization".

You can edit the list of group IDs and delete any group IDs you added in the past using these same steps. The revised list of group IDs will be updated after you run the script to commit your changes.

## View the current list of group IDs

It's possible for the contents of the `data_provider.yaml` configuration file to differ from what has actually been committed after running the `update_data_providers_from_config_file.sh` script. To view the current list of group IDs that you've added to BlueXP classification, run the following command from `/opt/netapp/Datasense/tools/customer_configuration/data_providers`:


```
get_data_providers_configuration.sh
```


## Remove data sources from BlueXP classification

If you need to, you can stop BlueXP classification from scanning one or more working environments, databases, or file share groups.

### Deactivate compliance scans for a working environment

When you deactivate scans, BlueXP classification no longer scans the data on the working environment and it removes the indexed compliance insights from the BlueXP classification instance (the data from the working environment itself isn't deleted).


1. From the *Configuration* page, select the  button in the row for the working environment then **Deactivate Data Sense**.



You can also disable compliance scans for a working environment from the Services panel when you select the working environment.

## Remove a database from BlueXP classification

If you no longer want to scan a certain database, you can delete it from the BlueXP classification interface and stop all scans.


1. From the *Configuration* page, select the  button in the row for the database then **Remove DB Server**.



## Remove a group of file shares from BlueXP classification

If you no longer want to scan user files from a file shares group, you can delete the File Shares Group from the BlueXP classification interface and stop all scans.

### Steps

1. From the *Configuration* page, select the  button in the row for the File Shares Group then **Remove File Shares Group**.



2. Click **Delete Group of Shares** from the confirmation dialog.


## Uninstall BlueXP classification

You can uninstall BlueXP classification to troubleshoot issues or to permanently remove the software from the host. Deleting the instance also deletes the associated disks where the indexed data resides, meaning all the information BlueXP classification has scanned will be permanently deleted.

The steps that you need to use depend on whether you deployed BlueXP classification in the cloud or on an on-premises host.

### Uninstall BlueXP classification from a cloud deployment

You can uninstall and delete the BlueXP classification instance from the cloud provider environment if you no longer want to use BlueXP classification.

1. At the top of the BlueXP classification page, select  then **Uninstall Classification**.



2. In the *Uninstall Classification* dialog, type **uninstall** to confirm that you want to disconnect the BlueXP classification instance from the BlueXP Connector then select **Uninstall**.
3. Go to your cloud provider's console and delete the BlueXP classification instance. The instance is named *CloudCompliance* with a generated hash (UUID) concatenated to it. For example: *CloudCompliance-16bb6564-38ad-4080-9a92-36f5fd2f71c7*

This deletes the instance and all associated data that had been collected by BlueXP classification.

## Uninstall BlueXP classification from an on-premises deployment

You can uninstall BlueXP classification from a host if you no longer want to use BlueXP classification, or if you had an issue that requires reinstallation.

1. At the top of the BlueXP classification page, select  then **Uninstall Classification**.



2. In the *Uninstall Classification* dialog, type **uninstall** to confirm that you want to disconnect the BlueXP classification instance from the BlueXP Connector then select **Uninstall**.
3. To uninstall the software from the host, run the `cleanup.sh` script on the BlueXP classification host machine, for example:

```
cleanup.sh
```

The script is located in the `/install/light_probe/onprem_installer/cleanup.sh` directory.

See how to [log in to the BlueXP classification host machine](#).

# Deprecated features

## BlueXP classification deprecated features

BlueXP classification is available as a core capability within BlueXP at no additional charge. By including BlueXP classification as a core BlueXP capability available to all customers, NetApp is enabling you to access tailored data management with core features.

There are some features and functionality that are deprecated in the BlueXP core version starting with version 1.31 and later and are still supported in legacy versions 1.30 and earlier.

### Supported data sources

Data source	Legacy versions 1.30 and earlier	BlueXP core versions 1.31 and later
Cloud Volumes ONTAP (deployed in AWS, Azure, or GCP)	Yes	Yes
On-premises ONTAP clusters	Yes	Yes
StorageGRID	Yes	Yes
Azure NetApp Files	Yes	Yes
Amazon FSx for ONTAP	Yes	Yes
Google Cloud NetApp Volumes	Yes	Yes
Cloud Volumes Service for Google Cloud	Yes	Yes
Databases	Yes	Yes
Amazon S3	Yes	No
Google Cloud Storage	Yes	No
OneDrive	Yes	No
SharePoint Online	Yes	No
SharePoint On-premises (SharePoint Server)	Yes	No
Google Drive	Yes	No

### Compliance features

Feature	Legacy versions 1.30 and earlier	BlueXP core versions 1.31 and later
Identify Personal Identifiable Information (PII)	Yes	Yes
Identify sensitive personal information	Yes	Yes
Respond to Data Subject Access Requests (DSAR)	Yes	Yes



<b>Feature</b>	<b>Legacy versions 1.30 and earlier</b>	<b>BlueXP core versions 1.31 and later</b>
Create a custom list of "personal data" that is identified	Yes	No
Notify users through email when files contain certain PII. (You define this criteria using <a href="#">Policies</a> .)	Yes	No
Use directory-level filters	Yes	Yes
Use directory-level PII analysis	Yes	No

## Features to manage your data

<b>Feature</b>	<b>Legacy versions 1.30 and earlier</b>	<b>BlueXP core versions 1.31 and later</b>
Move, copy, and delete source files	Yes	No
Categorize data using Status tags	Yes	No
Categorize data using AIP labels	Yes	No
Assign files to users	Yes	No
Rescan data on demand	Yes	No
Create custom classifiers	Yes	No
Exclude directories from scanning	Yes	Yes
Search for names within files	Yes	Yes
Export data to NFS/CIFS from investigation	Yes	Yes
Export data to CSV from investigation	Yes	Yes
Support multiple scanners	Yes	No
Integrate Active Directory	Yes	Yes
Use permission analysis and filters	Yes	Yes
Use the file card	Yes	Yes
Use the heatmap	Yes	Yes
Use actions on Dashboard and file card	Yes	No
Use file access audit logging	Yes	No
Enable file access from the Configuration page	Yes	No
Use certain predefined policies	Yes	No

## Deploy BlueXP classification deprecations

## Install BlueXP classification on multiple hosts for large configurations with no internet access

Complete a few steps to install BlueXP classification on multiple hosts in an on-premises site that doesn't have internet access - also known as *private mode*. This type of installation is perfect for your secure sites.

For very large configurations where you'll be scanning petabytes of data in sites without internet access, you can include multiple hosts to provide additional processing power. When using multiple host systems, the primary system is called the *Manager node* and the additional systems that provide extra processing power are called *Scanner nodes*.

Follow these steps when installing BlueXP classification software on multiple on-premises hosts in an offline environment.



The following information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

### Before you begin

- Verify that all your Linux systems for the Manager and Scanner nodes meet the host requirements.
- Verify that you have installed the two prerequisite software packages (Docker Engine or Podman, and Python 3).
- Make sure you have root privileges on the Linux systems.
- Verify that your offline environment meets the required permissions and connectivity.
- You must have the IP addresses of the scanner node hosts that you plan to use.
- The following ports and protocols must be enabled on all hosts:

Port	Protocols	Description
2377	TCP	Cluster management communications
7946	TCP, UDP	Inter-node communication
4789	UDP	Overlay network traffic
50	ESP	Encrypted IPsec overlay network (ESP) traffic
111	TCP, UDP	NFS Server for sharing files between the hosts (needed from each scanner node to manager node)
2049	TCP, UDP	NFS Server for sharing files between the hosts (needed from each scanner node to manager node)

### Steps

1. Follow steps 1 through 8 from the [Single-host installation](#) on the manager node.
2. As shown in step 9, when prompted by the installer, you can enter the required values in a series of prompts, or you can provide the required parameters as command line arguments to the installer.

In addition to the variables available for a single-host installation, a new option `-n <node_ip>` is used to specify the IP addresses of the scanner nodes. Multiple node IPs are separated by a comma.

For example, this command adds 3 scanner nodes:

```
sudo ./install.sh -a <account_id> -c <client_id> -t <user_token> --host  
<ds_host> --manager-host <cm_host> -n <node_ip1>,<node_ip2>,<node_ip3> --no  
-proxy --darksite
```

3. Before the manager node installation completes, a dialog displays the installation command needed for the scanner nodes. Copy the command (for example: `sudo ./node_install.sh -m 10.11.12.13 -t ABCDEF-1-3u69m1-1s35212`) and save it in a text file.
4. On **each** scanner node host:
  - a. Copy the Data Sense installer file (`cc_onprem_installer.tar.gz`) to the host machine.
  - b. Unzip the installer file.
  - c. Paste and run the command that you copied in step 3.

When the installation finishes on all scanner nodes and they have been joined to the manager node, the manager node installation finishes as well.

## Result

The BlueXP classification installer finishes installing packages, and registers the installation. Installation can take 15 to 25 minutes.

## What's Next

From the Configuration page you can select the local [on-prem ONTAP clusters](#) and local [databases](#) that you want to scan.

# Scan data deprecations

## Scan Amazon S3 buckets with BlueXP classification

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

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

## Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

### Set up the S3 requirements in your cloud environment

Ensure that your cloud environment can meet the requirements for BlueXP classification, including preparing an IAM role and setting up connectivity from BlueXP classification to S3. [See the complete list.](#)

2

### Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

**3**

### Activate BlueXP classification on your S3 working environment

Select the Amazon S3 working environment, click **Enable**, and select an IAM role that includes the required permissions.

**4**

### Select the buckets to scan

Select the buckets that you'd like to scan and BlueXP classification will start scanning them.

## Reviewing S3 prerequisites

The following requirements are specific to scanning S3 buckets.

### Set up an IAM role for the BlueXP classification instance

BlueXP classification needs permissions to connect to the S3 buckets in your account and to scan them. Set up an IAM role that includes the permissions listed below. BlueXP prompts you to select an IAM role when you enable BlueXP classification on the Amazon S3 working environment.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:Get*",
        "s3:List*",
        "s3:PutObject"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedRolePolicies"
      ],
      "Resource": [
        "arn:aws:iam::*:policy/*",
        "arn:aws:iam::*:role/*"
      ]
    }
  ]
}
```

## Provide connectivity from BlueXP classification to Amazon S3

BlueXP classification needs a connection to Amazon S3. The best way to provide that connection is through a VPC Endpoint to the S3 service. For instructions, see [AWS Documentation: Creating a Gateway Endpoint](#).

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

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

An alternative is to provide the connection by using a NAT Gateway.



You can't use a proxy to get to S3 over the internet.

## Deploying the BlueXP classification instance

[Deploy BlueXP classification in BlueXP](#) if there isn't already an instance deployed.

You need to deploy the instance using a Connector deployed in AWS so that BlueXP automatically discovers the S3 buckets in this AWS account and displays them in an Amazon S3 working environment.

**Note:** Deploying BlueXP classification in an on-premises location is not currently supported when scanning S3 buckets.

Upgrades to BlueXP classification software are automated as long as the instance has internet connectivity.

## Activating BlueXP classification on your S3 working environment

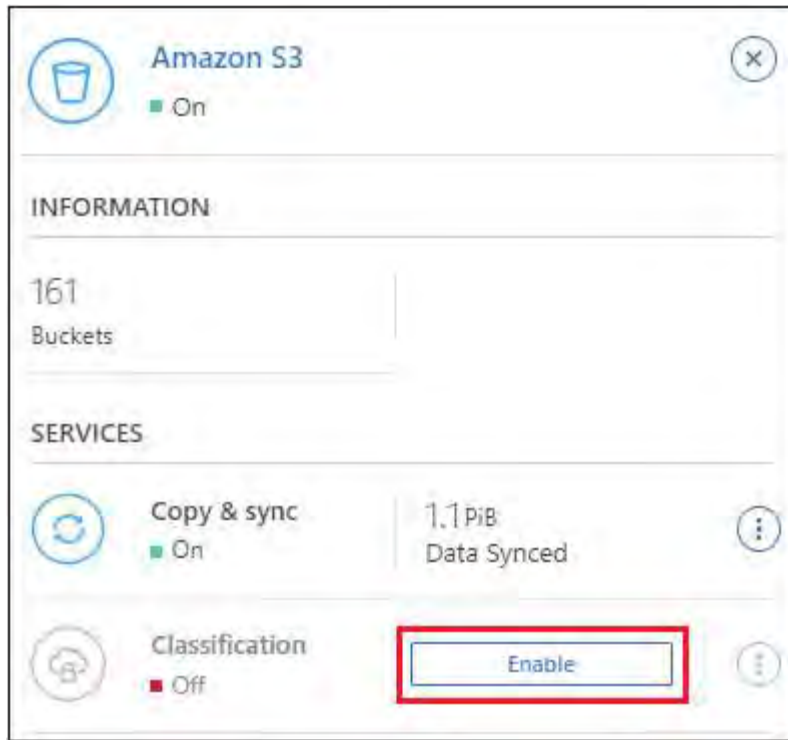
Enable BlueXP classification on Amazon S3 after you verify the prerequisites.

### Steps

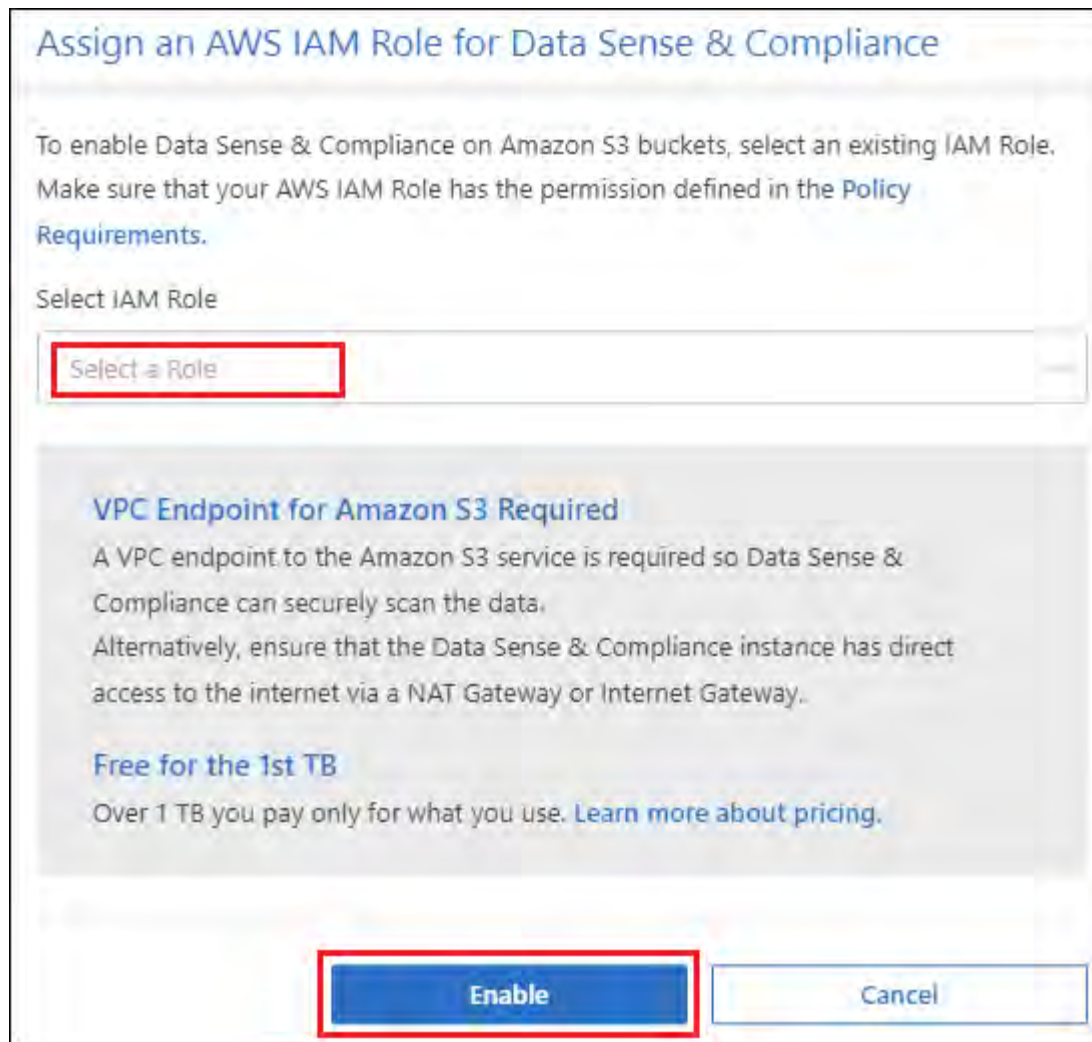
1. From the BlueXP left navigation menu, click **Storage > Canvas**.
2. Select the Amazon S3 working environment.



3. In the Services pane on the right, click **Enable** next to **Classification**.



4. When prompted, assign an IAM role to the BlueXP classification instance that has [the required permissions](#).



5. Select **Enable**.



You can also enable compliance scans for a working environment from the Configuration page by selecting the  button then **Activate BlueXP classification**.

### Result

BlueXP assigns the IAM role to the instance.

### Enabling and disabling compliance scans on S3 buckets

After BlueXP enables BlueXP classification on Amazon S3, the next step is to configure the buckets that you want to scan.

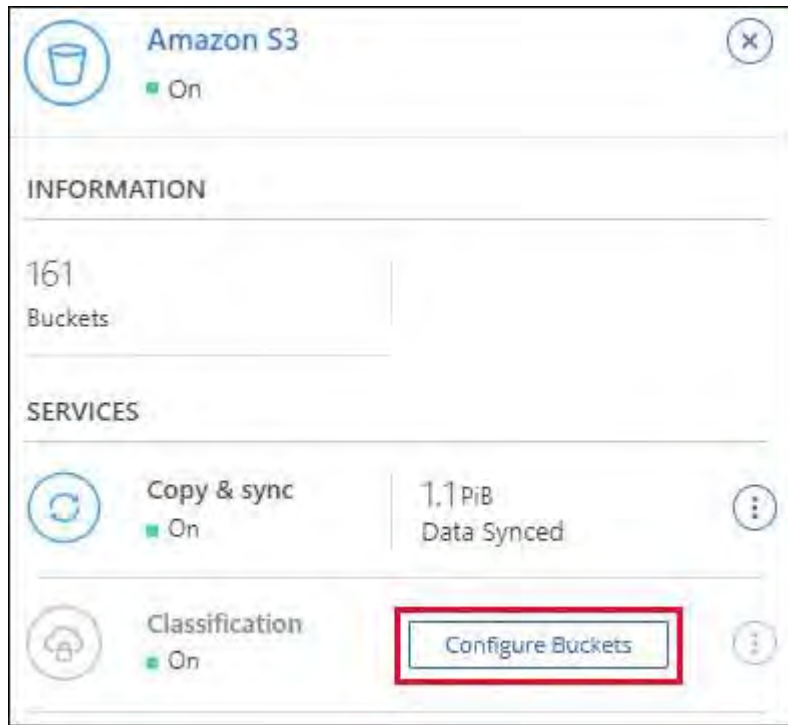
When BlueXP is running in the AWS account that has the S3 buckets you want to scan, it discovers those buckets and displays them in an Amazon S3 working environment.

BlueXP classification can also [scan S3 buckets that are in different AWS accounts](#).

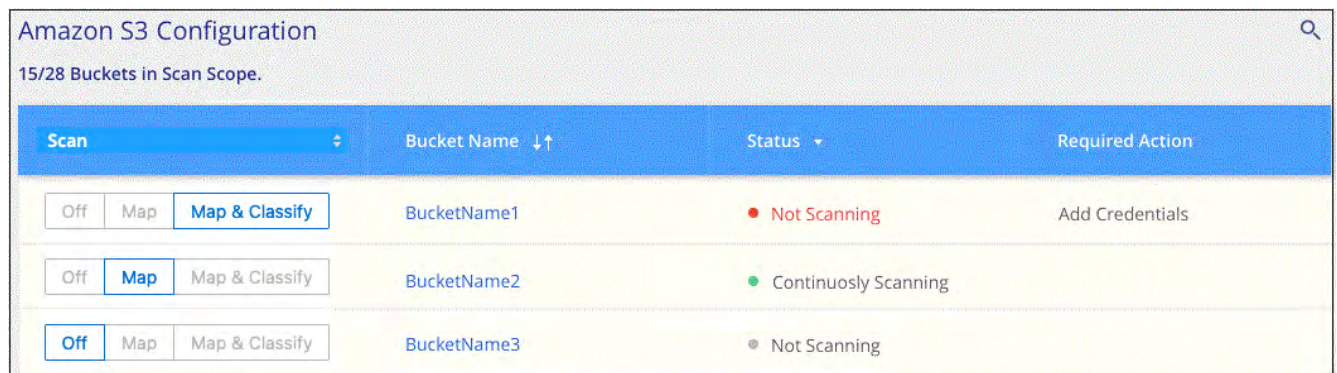
### Steps

1. Select the Amazon S3 working environment.
2. In the Services pane on the right, click **Configure Buckets**.





3. Enable mapping-only scans, or mapping and classification scans, on your buckets.



To:	Do this:
Enable mapping-only scans on a bucket	Click <b>Map</b>
Enable full scans on a bucket	Click <b>Map &amp; Classify</b>
Disable scanning on a bucket	Click <b>Off</b>

### Result

BlueXP classification starts scanning the S3 buckets that you enabled. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

### Scanning buckets from additional AWS accounts

You can scan S3 buckets that are under a different AWS account by assigning a role from that account to access the existing BlueXP classification instance.

### Steps



1. Go to the target AWS account where you want to scan S3 buckets and create an IAM role by selecting **Another AWS account**.

## Create role



### Select type of trusted entity

 <b>AWS service</b> EC2, Lambda and others	 <b>Another AWS account</b> Belonging to you or 3rd party	 <b>Web identity</b> Cognito or any OpenID provider	 <b>SAML 2.0 federation</b> Your corporate directory
--	---	---	--

Allows entities in other accounts to perform actions in this account. [Learn more](#)

### Specify accounts that can use this role

Account ID\*  ⓘ

- Options
- Require external ID (Best practice when a third party will assume this role)
  - Require MFA ⓘ

Be sure to do the following:

- Enter the ID of the account where the BlueXP classification instance resides.
- Change the **Maximum CLI/API session duration** from 1 hour to 12 hours and save that change.
- Attach the BlueXP classification IAM policy. Make sure it has the required permissions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:Get*",
        "s3:List*",
        "s3:PutObject"
      ],
      "Resource": "*"
    }
  ]
}
```

2. Go to the source AWS account where the BlueXP classification instance resides and select the IAM role that is attached to the instance.
  - a. Change the **Maximum CLI/API session duration** from 1 hour to 12 hours. Save the change.
  - b. Select **Attach policies** then **Create policy**.
  - c. Create a policy that includes the "sts:AssumeRole" action and specify the ARN of the role that you created in the target account.


```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::<ADDITIONAL-ACCOUNT-
ID>:role/<ADDITIONAL_ROLE_NAME>"
    },
    {
      "Effect": "Allow",
      "Action": [
        "iam:GetPolicyVersion",
        "iam:GetPolicy",
        "iam:ListAttachedRolePolicies"
      ],
      "Resource": [
        "arn:aws:iam::*:policy/*",
        "arn:aws:iam::*:role/*"
      ]
    }
  ]
}

```

The BlueXP classification instance profile account receives access to the additional AWS account.

3. Navigate to the **Amazon S3 Configuration** page and the new AWS account is displayed. Note that it can take a few minutes for BlueXP classification to sync the new account's working environment and show this information.



4. Click **Activate BlueXP classification & Select Buckets** and select the buckets you want to scan.

### Result

BlueXP classification starts scanning the new S3 buckets that you enabled.

## Scan OneDrive accounts with BlueXP classification

Complete a few steps to start scanning files in your user's OneDrive folders with BlueXP classification.

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

## Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

### Review OneDrive prerequisites

Ensure that you have the Admin credentials to log into the OneDrive account.

2

### Deploy the BlueXP classification instance

Deploy [BlueXP classification](#) if there isn't already an instance deployed.

3

### Add the OneDrive account

Using Admin user credentials, log into the OneDrive account that you want to access so that it is added as a new working environment.

4

### Add the users and select the type of scanning

Add the list of users from the OneDrive account that you want to scan and select the type of scanning. You can add up to 100 users at time.

## Reviewing OneDrive requirements

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.

- You must have the Admin login credentials for the OneDrive for Business account that provides read access to the user's files.
- You'll need a line-separated list of the email addresses for all the users whose OneDrive folders you want to scan.

## Deploying the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

BlueXP classification can be [deployed in the cloud](#) or [in an on-premises location that has internet access](#).

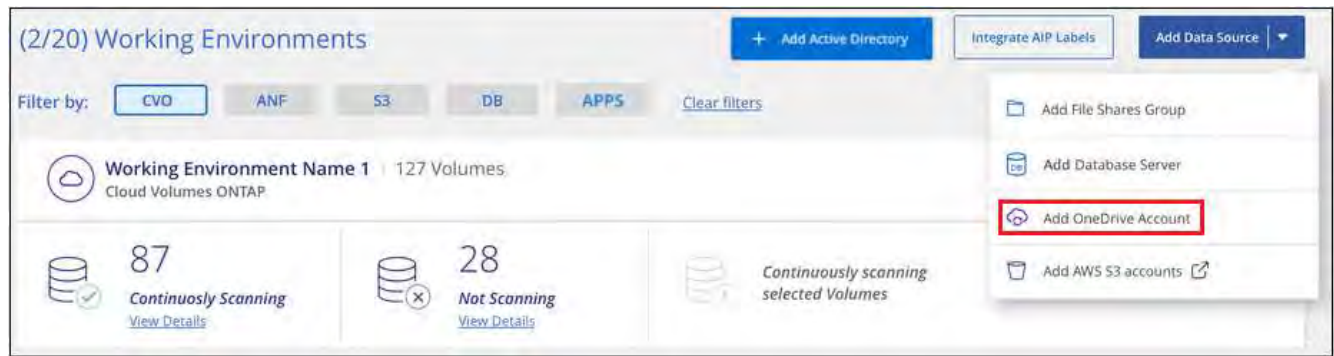
Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

## Adding the OneDrive account

Add the OneDrive account where the user files reside.

### Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add OneDrive Account**.



2. In the Add a OneDrive account dialog, click **Sign in to OneDrive**.
3. In the Microsoft page that appears, select the OneDrive account and enter the required Admin user and password, then click **Accept** to allow BlueXP classification to read data from this account.

The OneDrive account is added to the list of working environments.

### Adding OneDrive users to compliance scans


You can add individual OneDrive users, or all of your OneDrive users, so that their files will be scanned by BlueXP classification.

#### Steps

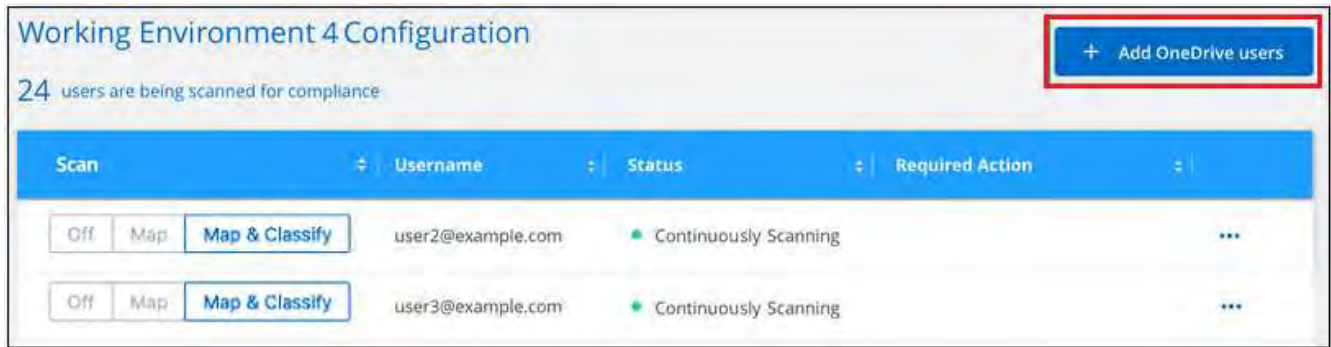
1. From the *Configuration* page, click the **Configuration** button for the OneDrive account.



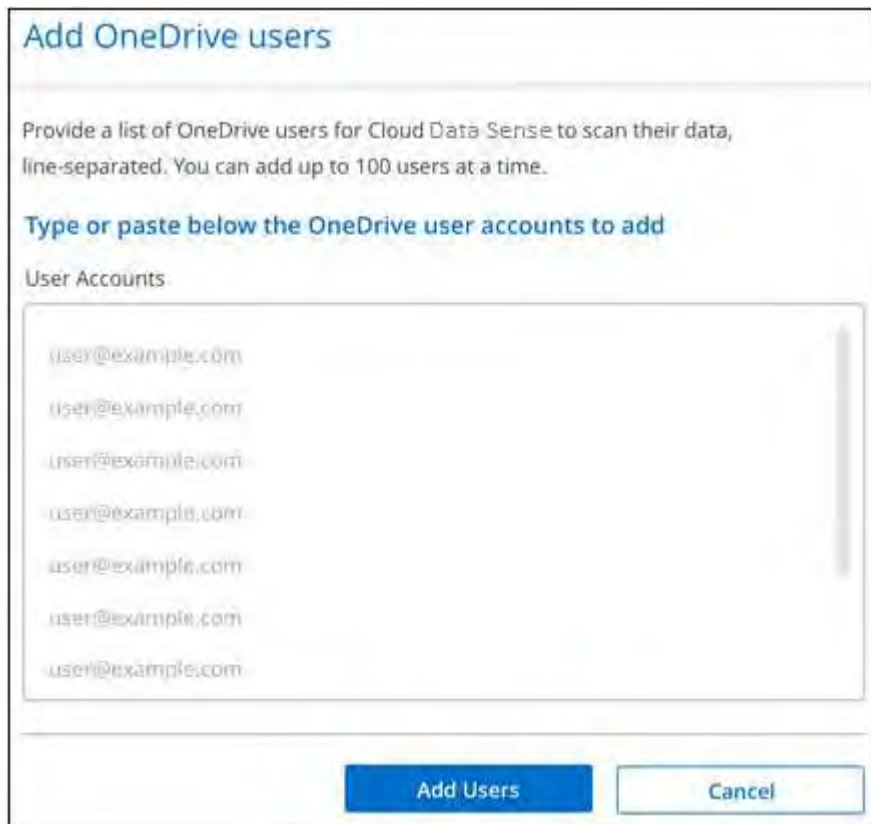
2. If this is the first time adding users for this OneDrive account, click **Add your first OneDrive users**.



If you are adding additional users from a OneDrive account, click **Add OneDrive users**.



3. Add the email addresses for the users whose files you want to scan - one email address per line (up to 100 maximum per session) - and click **Add Users**.



A confirmation dialog displays the number of users who were added.

If the dialog lists any users who could not be added, capture this information so that you can resolve the issue. In some cases you can re-add the user with a corrected email address.

4. Enable mapping-only scans, or mapping and classification scans, on user files.


To:	Do this:
Enable mapping-only scans on user files	Click <b>Map</b>
Enable full scans on user files	Click <b>Map &amp; Classify</b>
Disable scanning on user files	Click <b>Off</b>

## Result

BlueXP classification starts scanning the files for the users you added, and the results are displayed in the Dashboard and in other locations.

## Removing a OneDrive user from compliance scans

If users leave the company or if their email address changes, you can remove individual OneDrive users from having their files scanned at any time. Just click **Remove OneDrive User** from the Configuration page.



## Scan SharePoint accounts with BlueXP classification

Complete a few steps to start scanning files in your SharePoint Online and SharePoint On-Premise accounts with BlueXP classification.

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

## Review SharePoint requirements

Review the following prerequisites to make sure you are ready to activate BlueXP classification on a SharePoint account.

- You must have the Admin user login credentials for the SharePoint account that provides read access to all SharePoint sites.
  - For SharePoint Online you can use a non-Admin account, but that user must have permission to access all the SharePoint sites that you want to scan.
- For SharePoint On-Premise, you'll also need the URL of the SharePoint Server.
- You will need a line-separated list of the SharePoint site URLs for all the data you want to scan.

## Deploy the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

- For SharePoint Online, BlueXP classification can be [deployed in the cloud](#).
- For SharePoint On-Premises, BlueXP classification can be installed [in an on-premises location that has internet access](#) or [in an on-premises location that does not have internet access](#).

When BlueXP classification is installed in a site without internet access, the BlueXP Connector also must be installed in that same site without internet access. [Learn more](#).

Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

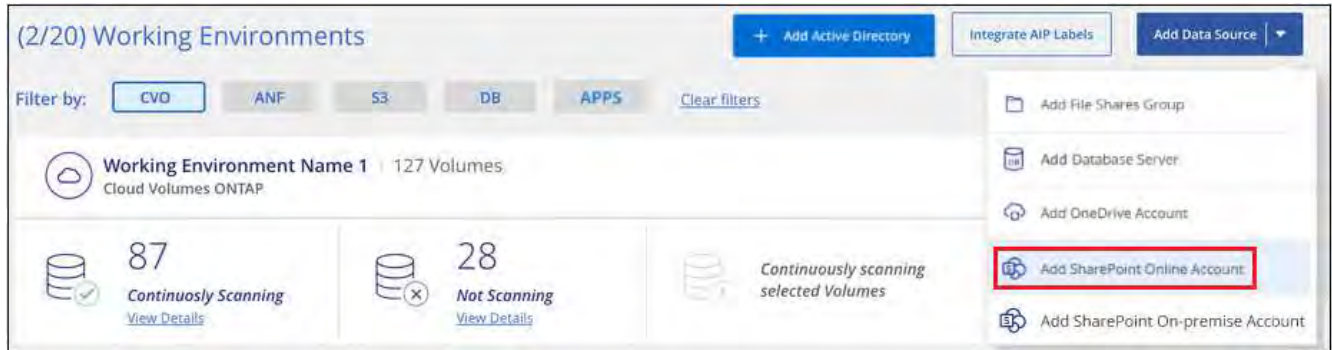


## Add a SharePoint Online account

Add the SharePoint Online account where the user files reside.

### Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add SharePoint Online Account**.



2. In the Add a SharePoint Online Account dialog, click **Sign in to SharePoint**.
3. In the Microsoft page that appears, select the SharePoint account and enter the user and password (Admin user or other user with access to the SharePoint sites), then click **Accept** to allow BlueXP classification to read data from this account.

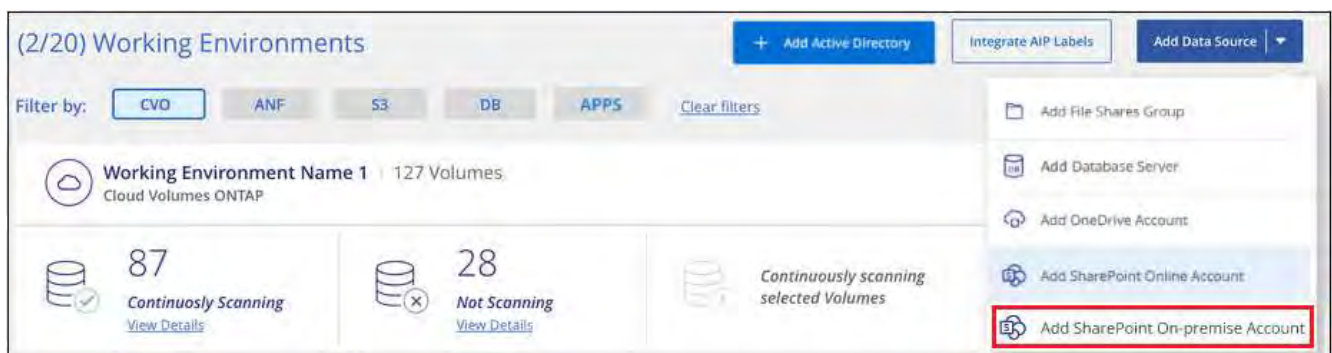
The SharePoint Online account is added to the list of working environments.

## Add a SharePoint On-premise account

Add the SharePoint On-premise account where the user files reside.

### Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add SharePoint On-premise Account**.



2. In the Log into the SharePoint On-Premise Server dialog, enter the following information:
  - Admin user in the format "domain/user" or "user@domain", and admin password
  - URL of the SharePoint Server

Log into the SharePoint On-Premises Server

To activate Data Sense on your SharePoint business account, sign in to SharePoint with an Admin user.

Username  Password

URL

3. Click **Connect**.

The SharePoint On-premise account is added to the list of working environments.

### Add SharePoint sites to compliance scans

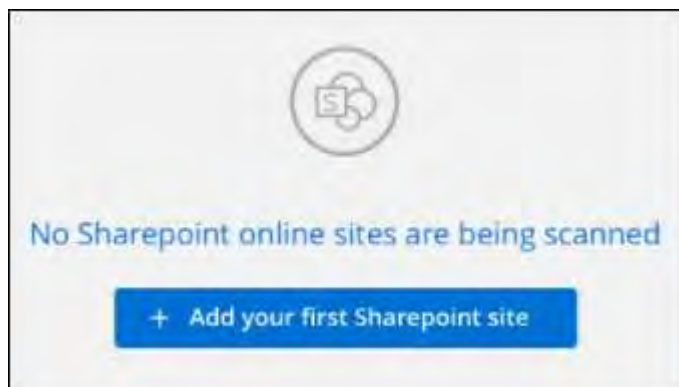
You can add individual SharePoint sites, or up to 1,000 SharePoint sites in the account, so that the associated files will be scanned by BlueXP classification. The steps are the same whether you are adding SharePoint Online or SharePoint On-premise sites.

#### Steps

1. From the *Configuration* page, click the **Configuration** button for the SharePoint account.

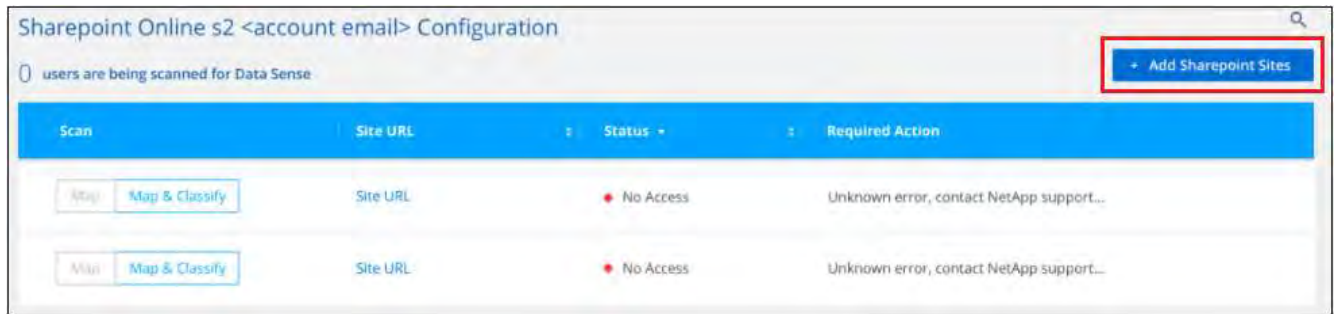


2. If this is the first time adding sites for this SharePoint account, click **Add your first SharePoint site**.

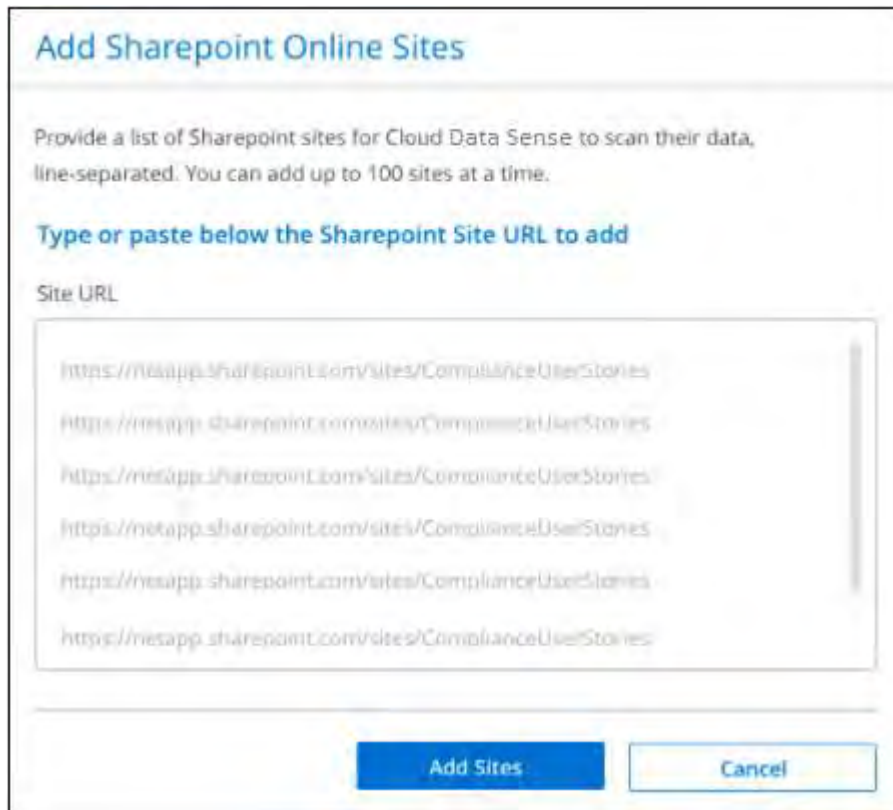


If you are adding additional users from a SharePoint account, click **Add SharePoint Sites**.





3. Add the URLs for the sites whose files you want to scan - one URL per line (up to 100 maximum per session) - and click **Add Sites**.



A confirmation dialog displays the number of sites that were added.

If the dialog lists any sites that could not be added, capture this information so that you can resolve the issue. In some cases you can re-add the site with a corrected URL.

4. If you need to add more than 100 sites for this account, just click **Add SharePoint Sites** again until you have added all your sites for this account (up to 1,000 sites total for each account).
5. Enable mapping-only scans, or mapping and classification scans, on the files in the SharePoint sites.

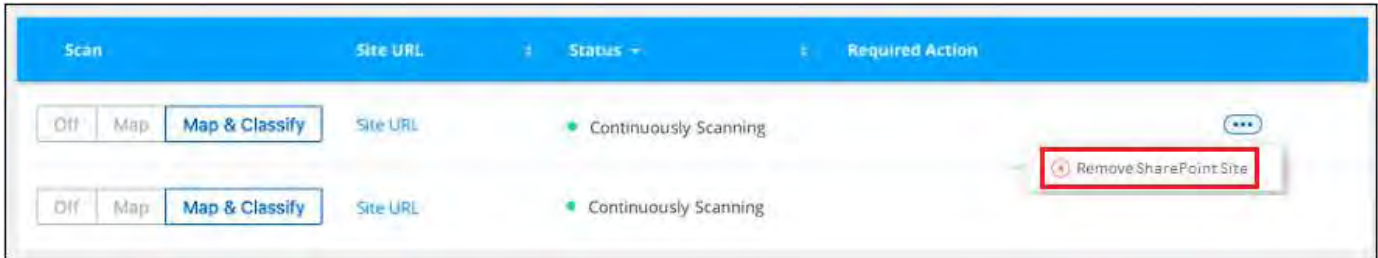
To:	Do this:
Enable mapping-only scans on files	Click <b>Map</b>
Enable full scans on files	Click <b>Map &amp; Classify</b>
Disable scanning on files	Click <b>Off</b>

## Result

BlueXP classification starts scanning the files in the SharePoint sites you added, and the results are displayed in the Dashboard and in other locations.

## Remove a SharePoint site from compliance scans

If you remove a SharePoint site in the future, or decide not to scan files in a SharePoint site, you can remove individual SharePoint sites from having their files scanned at any time. Just click **Remove SharePoint Site** from the Configuration page.



Note that you can [delete the entire SharePoint account from BlueXP classification](#) if you no longer want to scan any user data from the SharePoint account.

## Scan Google Drive accounts with BlueXP classification

Complete a few steps to start scanning user files in your Google Drive accounts with BlueXP classification.

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

### Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

#### 1 Review Google Drive prerequisites

Ensure that you have the Admin credentials to log into the Google Drive account.

#### 2 Deploy BlueXP classification

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

#### 3 Log into the Google Drive account

Using Admin user credentials, log into the Google Drive account that you want to access so that it is added as a new data source.

#### 4 Select the type of scanning for the user files

Select the type of scanning you want to perform on the user files; mapping or mapping and classifying.

## Review Google Drive requirements

Review the following prerequisites to make sure you are ready to enable BlueXP classification on a Google Drive account.

- You must have the Admin login credentials for the Google Drive account that provides read access to the user's files

## Current restrictions

The following BlueXP classification features are not currently supported with Google Drive files:

- When viewing files in the Data Investigation page, the actions in the button bar aren't active. You can't copy, move, delete, etc. any files.
- Permissions can't be identified within files in Google Drive, so no permission information is displayed in the Investigation page.

## Deploy BlueXP classification

Deploy BlueXP classification if there isn't already an instance deployed.

BlueXP classification can be [deployed in the cloud](#) or [in an on-premises location that has internet access](#).

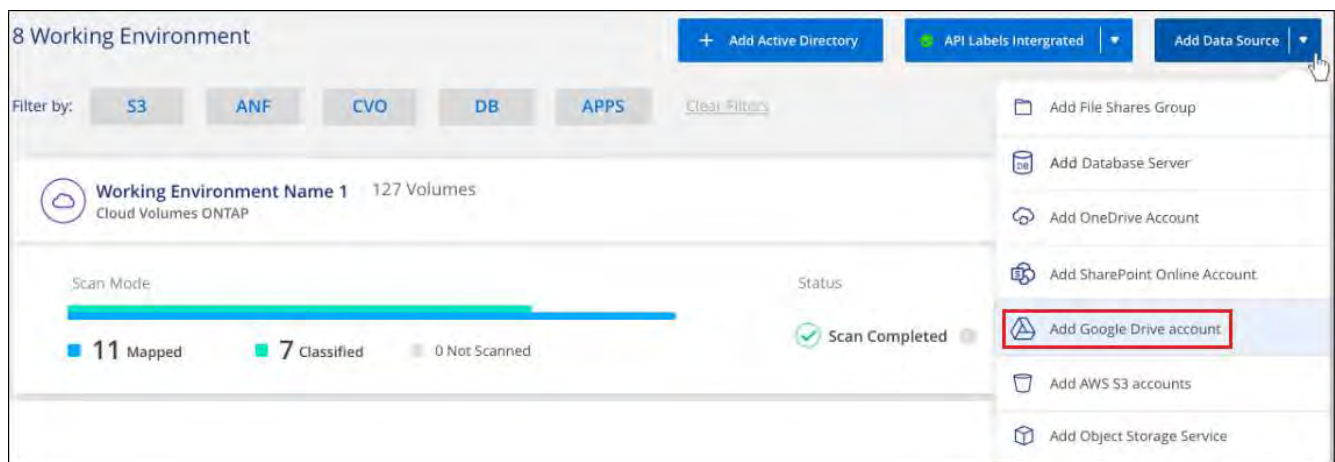
Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

## Add the Google Drive account

Add the Google Drive account where the user files reside. If you want to scan files from multiple users, you'll need to run through this step for each user.

## Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add Google Drive Account**.



2. In the Add a Google Drive Account dialog, click **Sign in to Google Drive**.
3. In the Google page that appears, select the Google Drive account and enter the required Admin user and password, then click **Accept** to allow BlueXP classification to read data from this account.

The Google Drive account is added to the list of working environments.

### Select the type of scanning for user data

Select the type of scanning that BlueXP classification will perform on the user's data.

#### Steps

1. From the *Configuration* page, click the **Configuration** button for the Google Drive account.
1. Enable mapping-only scans, or mapping and classification scans, on the files in the Google Drive account.



To:	Do this:
Enable mapping-only scans on files	Click <b>Map</b>
Enable full scans on files	Click <b>Map &amp; Classify</b>
Disable scanning on files	Click <b>Off</b>

#### Result

BlueXP classification starts scanning the files in the Google Drive account you added, and the results are displayed in the Dashboard and in other locations.

### Remove a Google Drive account from compliance scans

Since only a single user's Google Drive files are part of a single Google Drive account, if you want to stop scanning files from a user's Google Drive account, then you should [delete the Google Drive account from BlueXP classification](#).

### Scan StorageGRID data with BlueXP classification

Complete a few steps to start scanning data within object storage directly with BlueXP classification. BlueXP classification can scan data from any Object Storage service which uses the Simple Storage Service (S3) protocol. This includes NetApp StorageGRID, IBM Cloud Object Store, Linode, B2 Cloud Storage, Amazon S3, and more.



The following information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

Beginning with version 1.31, BlueXP classification is part of the core BlueXP offering. For more information, see [Scan StorageGRID data](#).

## Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

### Review object storage prerequisites

You need to have the endpoint URL to connect with the object storage service.

You need to have the Access Key and Secret Key from the object storage provider so that BlueXP classification can access the buckets.

2

### Deploy the BlueXP classification instance

Deploy [BlueXP classification](#) if there isn't already an instance deployed.

3

### Add the Object Storage Service

Add the object storage service to BlueXP classification.

4

### Select the buckets to scan

Select the buckets that you'd like to scan and BlueXP classification will start scanning them.

## Reviewing object storage requirements

Review the following prerequisites to make sure that you have a supported configuration before you enable BlueXP classification.

- You need to have the endpoint URL to connect with the object storage service.
- You need to have the Access Key and Secret Key from the object storage provider so that BlueXP classification can access the buckets.

## Deploying the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

If you are scanning data from S3 object storage that is accessible over the internet, you can [deploy BlueXP classification in the cloud](#) or [deploy BlueXP classification in an on-premises location that has internet access](#).

If you are scanning data from S3 object storage that has been installed in a dark site that has no internet access, you need to [deploy BlueXP classification in the same on-premises location that has no internet access](#). This also requires that the BlueXP Connector is deployed in that same on-premises location.


Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

## Adding the object storage service to BlueXP classification

Add the object storage service.

## Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add Object Storage Service**.



2. In the Add Object Storage Service dialog, enter the details for the object storage service and click **Continue**.
  - a. Enter the name you want to use for the Working Environment. This name should reflect the name of the object storage service to which you are connecting.
  - b. Enter the Endpoint URL to access the object storage service.
  - c. Enter the Access Key and Secret Key so that BlueXP classification can access the buckets in the object storage.

The screenshot shows the 'Add Object Storage Service' dialog box. It contains the following text: 'Cloud Data Sense can scan data from any Object Storage service which uses the S3 protocol. This includes NetApp StorageGRID, IBM Object Store, and more. To continue, enter the following information. In the next steps you'll need to select the buckets you want to scan.' Below this are four input fields: 'Name the Working Environment' (containing 'object\_myIBM'), 'Endpoint URL' (containing 'http://my.endpoint.com'), 'Access Key' (containing 'AIJKDO574NDJG86795'), and 'Secret Key' (containing a series of dots). At the bottom are 'Continue' and 'Cancel' buttons.

### Result

The new Object Storage Service is added to the list of working environments.


### Enabling and disabling compliance scans on object storage buckets

After you enable BlueXP classification on your Object Storage Service, the next step is to configure the buckets that you want to scan. BlueXP classification discovers those buckets and displays them in the working environment you created.


### Steps

1. In the Configuration page, click **Configuration** from the Object Storage Service working environment.





2. Enable mapping-only scans, or mapping and classification scans, on your buckets.



To:	Do this:
Enable mapping-only scans on a bucket	Click <b>Map</b>
Enable full scans on a bucket	Click <b>Map &amp; Classify</b>
Disable scanning on a bucket	Click <b>Off</b>

## Result

BlueXP classification starts scanning the buckets that you enabled. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

## Manage data deprecations

### View governance details about your data using the BlueXP classification Governance dashboard

Gain control of the costs related to the data on your organizations' storage resources. BlueXP classification identifies the amount of stale data, non-business data, duplicate files, and very large files in your systems so you can decide whether you want to remove or tier some files to less expensive object storage.

Additionally, if you're planning to migrate data from on-premises locations to the cloud, you can view the size of the data and whether any of the data contains sensitive information before moving it.

**NOTE** This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

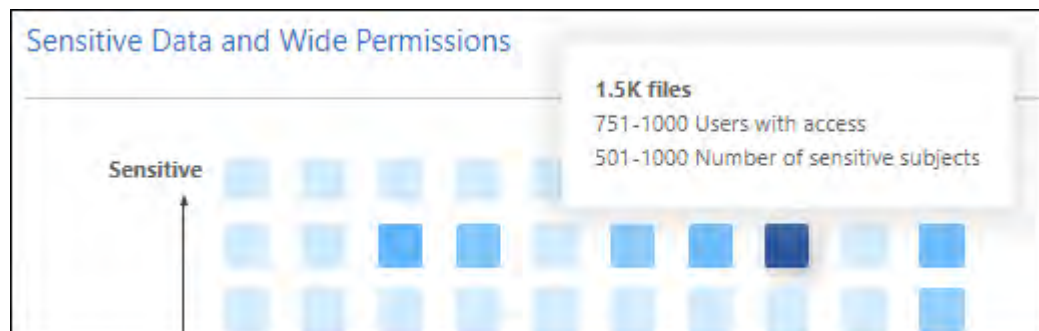
### Data listed by sensitivity and wide permissions on the Governance dashboard

The *Sensitive Data and Wide Permissions* area on the Governance dashboard provides a heatmap of files that contain sensitive data (including both sensitive and sensitive personal data) and that are overly permissive. This can help you to see where you may have some risks with sensitive data.



This applies to BlueXP classification versions 1.30 and earlier.

Files are rated based on the number of users with permission to access the files on the X axis (lowest to highest), and the number of sensitive identifiers within the files on the Y axis (lowest to highest). The blocks represent the number of files that match the items from the X and Y axes. The lighter colored blocks are good; with fewer users able to access the files, and with fewer sensitive identifiers per file. The darker blocks are the items you may want to investigate. For example, the screen below shows the tooltip text for the dark blue block. It shows that you have 1,500 files where 751-1000 users have access, and where there are 501-1000 sensitive identifiers per file.



You can click the block you are interested in to view the filtered results of the affected files in the Investigation page so that you can investigate further.

No data is displayed in this panel if you haven't integrated an identity service with BlueXP classification. [See how to integrate your Active Directory service with BlueXP classification.](#)



This panel supports files in CIFS shares, OneDrive, and SharePoint data sources. There is currently no support for databases, Google Drive, Amazon S3, and generic object storage.

### Classification area on the dashboard showing AIP labels

The *Classification* area on the dashboard provides a list of the most identified Azure Information Protection (AIP) Labels in your scanned data.

If you have subscribed to Azure Information Protection (AIP), you can classify and protect documents and files by applying labels to content. Reviewing the most used AIP labels that are assigned to files enables you to see which labels are most used in your files.

See [AIP Labels](#) for more information.



## Organize your private data with BlueXP classification

BlueXP classification provides many ways for you to manage and organize your private data. This makes it easier to see the data that is most important to you.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

- If you are subscribed to [Azure Information Protection \(AIP\)](#) to classify and protect your files, you can use BlueXP classification to manage AIP labels.
- You can add Tags to files that you want to mark for organization or for some type of follow-up.
- You can assign a BlueXP user to a specific file, or to multiple files, so that person can be responsible for managing the file.
- With the saved search functionality, you can create your own custom search queries so that you can easily see the results by clicking one button.
- You can send email alerts to BlueXP users, or any other email address, when certain critical Policies return results.



The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.

### Should I use tags or labels?

Below is a comparison of BlueXP classification tagging and Azure Information Protection labeling.

Tags	Labels
File tags are an integrated part of BlueXP classification.	Requires that you have subscribed to Azure Information Protection (AIP).
The tag is only kept in the BlueXP classification database - it is not written to the file. It does not change the file, or the file accessed or modified times.	The label is part of the file and when the label changes, the file changes. This change also changes the file accessed and modified times.
You can have multiple tags on a single file.	You can have one label on a single file.
The tag can be used for internal BlueXP classification action, such as copy, move, delete, run a policy, etc.	Other systems that can read the file can see the label - which can be used for additional automation.
Only a single API call is used to see if a file has a tag.	

### Categorize your data using AIP labels

You can manage AIP labels in the files that BlueXP classification is scanning if you have subscribed to [Azure Information Protection \(AIP\)](#). AIP enables you to classify and protect documents and files by applying labels to content. BlueXP classification enables you to view the labels that are already assigned to files, add labels to files, and change labels when a label already exists.

BlueXP classification supports AIP labels within the following file types: .DOC, .DOCX, .PDF, .PPTX, .XLS, .XLSX.



- You can't currently change labels in files larger than 30 MB. For OneDrive, SharePoint, and Google Drive accounts the maximum file size is 4 MB.
- If a file has a label which doesn't exist anymore in AIP, BlueXP classification considers it as a file without a label.
- If you've deployed BlueXP classification in a Government region, or in an on-prem location that has no internet access (also known as a dark site), then the AIP label functionality is unavailable.

### Integrate AIP labels in your project or workspace


Before you can manage AIP labels, you need to integrate the AIP label functionality into BlueXP classification by signing into your existing Azure account. Once enabled, you can manage AIP labels within files for all [data sources](#) in your BlueXP project or workspace.

### Requirements

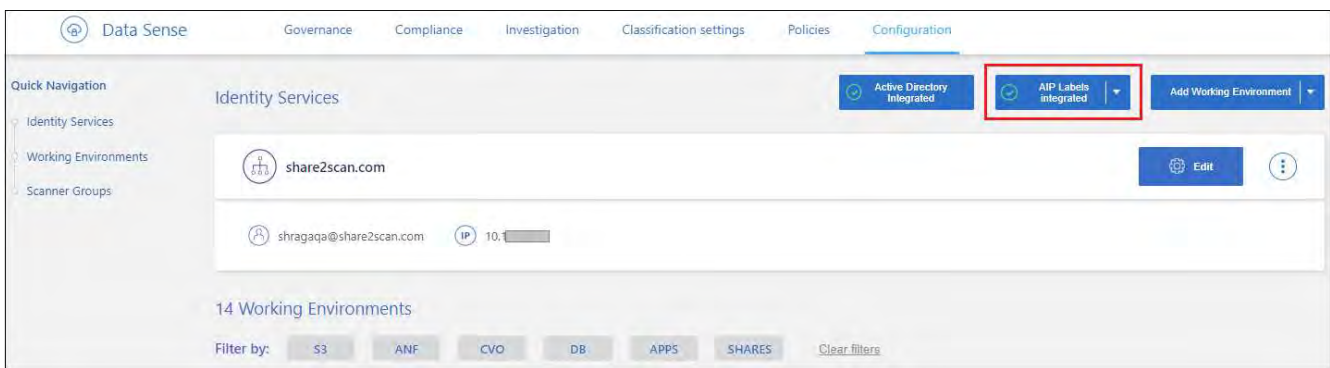
- You must have an account and an Azure Information Protection license.
- You must have the login credentials for the Azure account.
- If you plan to change labels in files that reside in Amazon S3 buckets, ensure that the permission `s3:PutObject` is included in the IAM role. See [setting up the IAM role](#).

### Steps

1. From the BlueXP classification Configuration page, click **Integrate AIP Labels**.



2. In the Integrate AIP Labels dialog, click **Sign in to Azure**.
3. In the Microsoft page that appears, select the account and enter the required credentials.
4. Return to the BlueXP classification tab and you'll see the message "AIP Labels were integrated successfully with the account <account\_name>".
5. Click **Close** and you'll see the text *AIP Labels integrated* at the top of the page.



### Result

You can view and assign AIP labels from the results pane of the Investigation page. You can also assign AIP

labels to files using Policies.

### View AIP labels in your files

You can view the current AIP label that is assigned to a file.

In the Data Investigation results pane, click  for the file to expand the file metadata details.



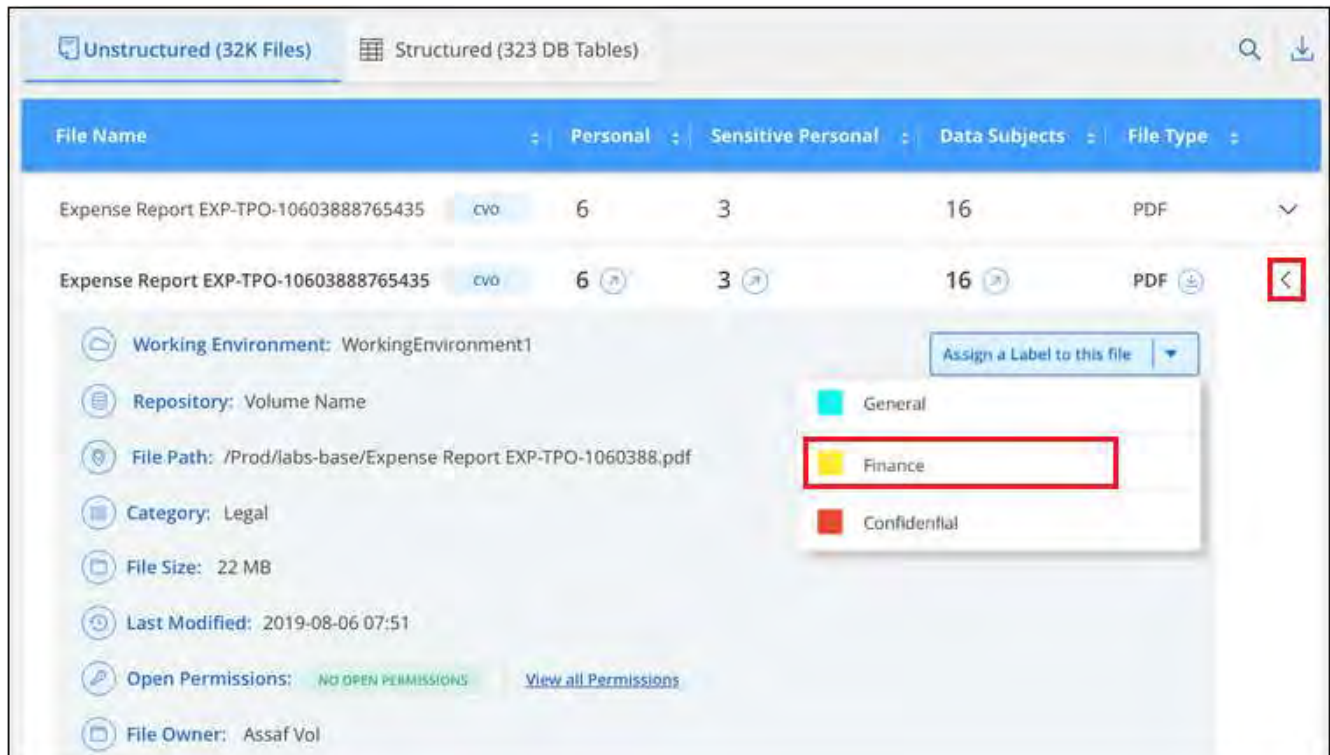
### Assign AIP labels manually

You can add, change, and remove AIP labels from your files using BlueXP classification.

Follow these steps to assign an AIP label to a single file.

### Steps

1. In the Data Investigation results pane, click  for the file to expand the file metadata details.



2. Click **Assign a Label to this file** and then select the label.

The label appears in the file metadata.

Follow these steps to assign an AIP label to multiple files. Note that you can assign an AIP label to a maximum of 20 files at a time (one page in the UI).

### Steps


1. In the Data Investigation results pane, select the file, or files, that you want to label.



◦ To select individual files, check the box for each file ( Volume\_1).

◦ To select all files on the current page, check the box in the title row ( File Name).

2. From the button bar, click **Label** and select the AIP label:



The AIP label is added to the metadata for all selected files.

### Remove the AIP integration

If you no longer want the ability to manage AIP labels in files, you can remove the AIP account from the BlueXP classification interface.

Note that no changes are made to the labels you have added using BlueXP classification. The labels that exist in files will stay as they currently exist.

### Steps

1. From the *Configuration* page, click **AIP Labels integrated > Remove Integration**.



2. Click **Remove Integration** from the confirmation dialog.

### Apply tags to manage your scanned files

You can add a tag to files that you want to mark for some type of follow-up. For example, you may have found some duplicate files and you want to delete one of them, but you need to check to see which one should be deleted. You could add a tag of "Check to delete" to the file so you know this file requires some research and some type of future action.

BlueXP classification enables you to view the tags that are assigned to files, add or remove tags from files, and change the name or delete an existing tag.

Note that the tag is not added to the file in the same way as AIP Labels are part of the file metadata. The tag is just seen by BlueXP users using BlueXP classification so you can see if a file needs to be deleted or checked for some type of follow-up.

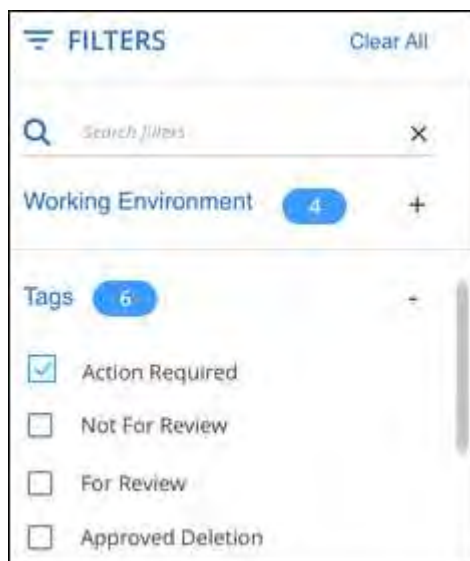


Tags assigned to files in BlueXP classification are not related to the tags you can add to resources, such as volumes or virtual machine instances. BlueXP classification tags are applied at the file level.

### View files that have certain tags applied

You can view all the files that have specific tags assigned.

1. Click the **Investigation** tab from BlueXP classification.
2. In the Data Investigation page, click **Tags** in the Filters pane and then select the required tags.






The Investigation Results pane displays all the files that have those tags assigned.

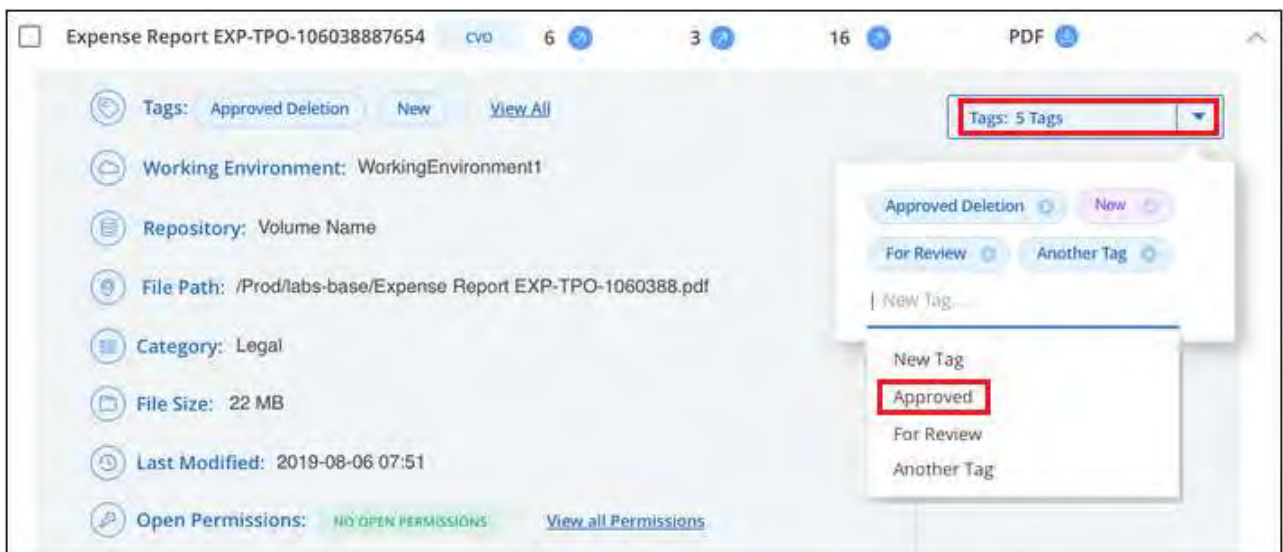
### Assign tags to files

You can add tags to a single file or to a group of files.

To add a tag to a single file:

#### Steps

1. In the Data Investigation results pane, click  for the file to expand the file metadata details.
2. Click the **Tags** field and the currently assigned tags are displayed.
3. Add the tag or tags:
  - To assign an existing tag, click in the **New Tag...** field and start typing the name of the tag. When the tag you are looking for appears, select it and press **Enter**.
  - To create a new tag and assign it to the file, click in the **New Tag...** field, enter the name of the new tag, and press **Enter**.



The tag appears in the file metadata.

To add a tag to multiple files:

#### Steps

1. In the Data Investigation results pane, select the file, or files, that you want to tag.


255 items 1.2 GB | 2 Selected 3 MB Tags Assign to Label Copy Move Delete

<input type="checkbox"/>	File Name	Personal	Sensitive Personal	Data Subjects	File Type	
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	16	PDF
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF

- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).
- To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 items on this page selected Select all items in list (63K items)**, click **Select all items in list (xxx items)**.

You can apply tags to a maximum of 100,000 files at a time.

- From the button bar, click **Tags** and the currently assigned tags are displayed.
- Add the tag or tags:
  - To assign an existing tag, click in the **New Tag...** field and start typing the name of the tag. When the tag you are looking for appears, select it and press **Enter**.
  - To create a new tag and assign it to the file, click in the **New Tag...** field, enter the name of the new tag, and press **Enter**.



- Approve adding the tags in the confirmation dialog and the tags are added to the metadata for all selected files.

#### Delete tags from files

You can delete a tag if you don't need to use it anymore.

Just click the **x** for an existing tag.

If you had selected multiple files, the tag is removed from all the files.

### Assign users to manage certain files

You can assign a BlueXP user to a specific file, or to multiple files, so that person can be responsible for any follow-up actions that need to be done on the file. This capability is often used with the feature to add custom Status tags to a file.


For example, you might have a file that contains certain personal data that allows too many users read and write access (open permissions). So you could assign the Status tag "Change permissions" and assign this file to user "Joan Smith" so they can decide how to fix the issue. When they have fixed the issue they could change the Status tag to "Completed".


Note that the user name is not added to the file as part of the file metadata - it is just seen by BlueXP users when using BlueXP classification.

A new Filter in the Investigation page enables you to easily view all files that have the same person in the "Assigned To" field.

Follow these steps to assign a user to a single file.

#### Steps

1. In the Data Investigation results pane, click  for the file to expand the file metadata details.
2. Click the **Assigned to** field and select the user name.



The User name appears in the file metadata.

Follow these steps to assign a user to multiple files. Note that you can assign a user to a maximum of 20 files at a time (one page in the UI).

#### Steps

1. In the Data Investigation results pane, select the file, or files, that you want to assign to a user.




255 items 1.2 GB | 2 Selected 3 MB

Tags Assign to Label Copy Move Delete

<input type="checkbox"/>	File Name	Personal	Sensitive Personal	Data Subjects	File Type	
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	16	PDF
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF

- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).

2. From the button bar, click **Assign to** and select the user name:



The user is added to the metadata for all selected files.

## Manage your private data with BlueXP classification

BlueXP classification provides many ways for you to manage your private data. Some functionality makes it easier to prepare for migrating your data, while other functionality allows you to make changes to the data.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

- You can copy files to a destination NFS share if you want to make a copy of certain data and move it to a different NFS location.
- You can clone an ONTAP volume to a new volume, while including only selected files from the source volume in the new cloned volume. This is useful for situations where you're migrating data and you want to exclude certain files from the original volume.
- You can copy and synchronize files from a source repository to a directory in a specific destination location. This is useful for situations where you're migrating data from one source system to another while there is

still some final activity on the source files.

- You can move source files that BlueXP classification is scanning to any NFS share.
- You can delete files that seem insecure or too risky to leave in your storage system, or that you have identified as duplicate.



- The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.
- Data from Google Drive accounts can't use any of these capabilities at this time.

## Copy source files

You can copy any source files that BlueXP classification is scanning. There are three types of copy operations depending on what you're trying to accomplish:

- **Copy files** from the same, or different, volumes or data sources to a destination NFS share.

This is useful if you want to make a copy of certain data and move it to a different NFS location.

- **Clone an ONTAP volume** to a new volume in the same aggregate, but include only selected files from the source volume in the new cloned volume.

This is useful for situations where you're migrating data and you want to exclude certain files from the original volume. This action uses the [NetApp FlexClone](#) functionality to quickly duplicate the volume and then remove the files that you **didn't** select.

- **Copy and synchronize files** from a single source repository (ONTAP volume, S3 bucket, NFS share, etc.) to a directory in a specific destination (target) location.

This is useful for situations where you're migrating data from one source system to another. After the initial copy, the service syncs any changed data based on the schedule that you set. This action uses the [NetApp BlueXP copy and sync](#) functionality to copy and sync data from a source to a target.

## Copy source files to an NFS share

You can copy source files that BlueXP classification is scanning to any NFS share. The NFS share doesn't need to be integrated with BlueXP classification, you just need to know the name of the NFS share where all selected files will be copied in the format `<host_name>:/<share_path>`.



You can't copy files that reside in databases.

## Requirements

- You must have permissions to copy files. [Learn about user access to compliance information.](#)
- Copying files requires that the destination NFS share allows access from the BlueXP classification instance.
- You can copy between 1 and 100,000 files at a time.


## Steps

1. In the Data Investigation results pane, select the file, or files, that you want to copy, and click **Copy**.



- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).
- To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 items on this page selected** [Select all items in list \(63K items\)](#), click **Select all items in list (xxx items)**.

2. In the *Copy Files* dialog, select the **Regular Copy** tab.



3. Enter the name of the NFS share where all selected files will be copied in the format `<host_name>:/<share_path>`, and click **Copy**.

A dialog appears with the status of the copy operation.

You can view the progress of the copy operation in the [Actions Status](#) pane.

Note that you can also copy an individual file when viewing the metadata details for a file. Just click **Copy File**.



### Clone volume data to a new volume

You can clone an existing ONTAP volume that BlueXP classification is scanning using NetApp *FlexClone* functionality. This allows you to quickly duplicate the volume while including only those files you selected. This is useful if you're migrating data and you want to exclude certain files from the original volume, or if you want to create a copy of a volume for testing.

The new volume is created in the same aggregate as the source volume. Ensure that you have enough space for this new volume in the aggregate before you start this task. Contact your storage administrator if necessary.


**Note:** FlexGroup volumes can't be cloned because they're not supported by FlexClone.

### Requirements

- You must have permissions to copy files. [Learn about user access to compliance information.](#)
- You must select a minimum of 20 files.
- All selected files must be from the same volume, and the volume must be online.
- The volume must be from a Cloud Volumes ONTAP or on-premises ONTAP system. No other data sources are currently supported.
- The FlexClone license must be installed on the cluster. This license is installed by default on Cloud Volumes ONTAP systems.

### Steps

1. In the Data Investigation pane, create a filter by selecting a single **Working Environment** and a single **Storage Repository** to make sure all the files are from the same ONTAP volume.



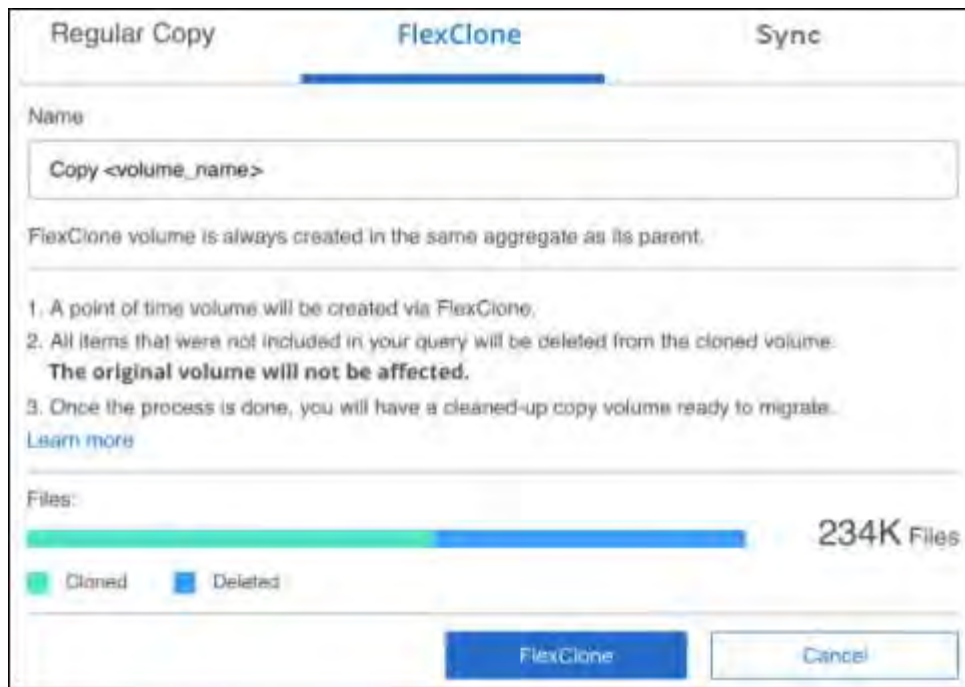
Apply any other filters so that you're seeing only the files that you want to clone to the new volume.

2. In the Investigation results pane, select the files that you want to clone and click **Copy**.



- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).
- To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 items on this page selected** [Select all items in list \(63K items\)](#), click **Select all items in list (xxx items)**.

3. In the *Copy Files* dialog, select the **FlexClone** tab. This page shows the total number of files that will be cloned from the volume (the files you selected), and the number of files that are not included/deleted (the files you didn't select) from the cloned volume.



4. Enter the name of the new volume, and click **FlexClone**.

A dialog appears with the status of the clone operation.

## Result

The new, cloned volume is created in the same aggregate as the source volume.

You can view the progress of the clone operation in the [Actions Status pane](#).

If you initially selected **Map all volumes** or **Map & Classify all volumes** when you enabled BlueXP classification for the working environment where the source volume resides, then BlueXP classification will scan the new cloned volume automatically. If you didn't use either of these selections initially, then if you want to scan this new volume, you'll need to [enable scanning on the volume manually](#).

## Copy and synchronize source files to a target system

You can copy source files that BlueXP classification is scanning from any supported unstructured data source to a directory in a specific target destination location ([target locations that are supported by BlueXP copy and sync](#)). After the initial copy, any data changed in the files are synchronized based on the schedule that you configure.

This is useful for situations where you're migrating data from one source system to another. This action uses the [NetApp BlueXP copy and sync](#) functionality to copy and sync data from a source to a target.



You can't copy and sync files that reside in databases, OneDrive accounts, or SharePoint accounts.

## Requirements

- You must have permissions to copy and sync files. [Learn about user access to compliance information](#).
- You must select a minimum of 20 files.
- All selected files must be from the same source repository (ONTAP volume, S3 bucket, NFS or CIFS share, etc.).

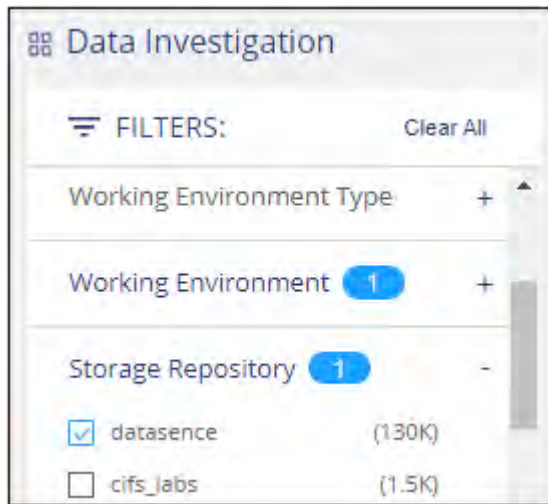


- You'll need to activate the BlueXP copy and sync service and configure a minimum of one data broker that can be used to transfer files between the source and target systems. Review the BlueXP copy and sync requirements beginning with the [Quick Start description](#).

Note that the BlueXP copy and sync service has separate service charges for your sync relationships, and will incur resource charges if you deploy the data broker in the cloud.

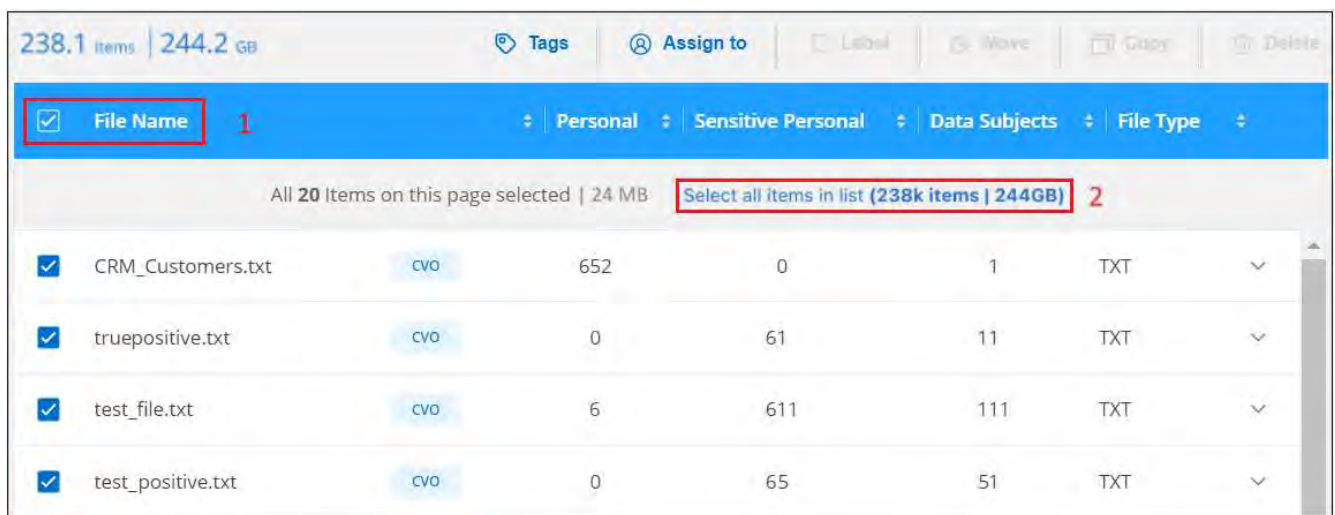
## Steps

- In the Data Investigation pane, create a filter by selecting a single **Working Environment** and a single **Storage Repository** to make sure all the files are from the same repository.

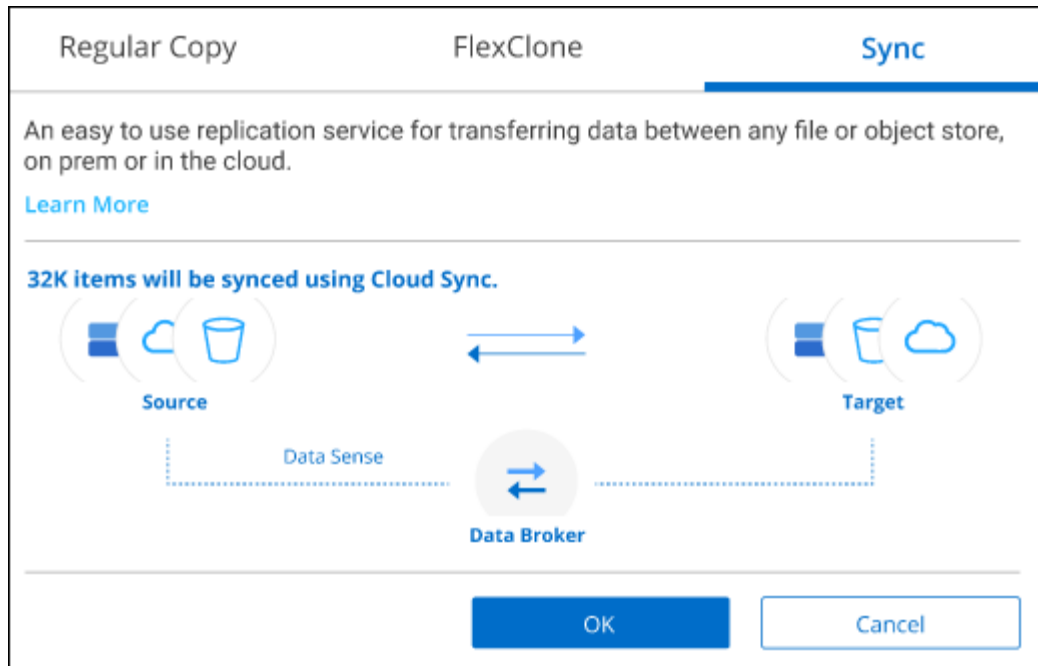


Apply any other filters so that you're seeing only the files that you want to copy and sync to the destination system.

- In the Investigation results pane, select all files on all pages by checking the box in the title row ( **File Name**), then in the pop-up message **All 20 Items on this page selected** [Select all Items in list \(63K Items\)](#) click **Select all items in list (xxx items)**, and then click **Copy**.



- In the *Copy Files* dialog, select the **Sync** tab.



- If you are sure that you want to sync the selected files to a destination location, click **OK**.

The BlueXP copy and sync UI is opened in BlueXP.

You are prompted to define the sync relationship. The Source system is pre-populated based on the repository and files you already selected in BlueXP classification.

- You'll need to select the Target system and then select (or create) the Data Broker you plan to use. Review the BlueXP copy and sync requirements beginning with the [Quick Start description](#).

### Result

The files are copied to the target system and they'll be synchronized based on the schedule you define. If you select a one-time sync then the files are copied and synchronized one time only. If you choose a periodic sync, then the files are synchronized based on the schedule. Note that if the source system adds new files that match the query you created using filters, those *new* files will be copied to the destination and synchronized in the future.

Note that some of the usual BlueXP copy and sync operations are disabled when it is invoked from BlueXP classification:

- You can't use the **Delete Files on Source** or **Delete Files on Target** buttons.
- Running a report is disabled.

### Move source files to an NFS share

You can move source files that BlueXP classification is scanning to any NFS share. The NFS share doesn't need to be integrated with BlueXP classification.

Optionally, you can leave a breadcrumb file in the location of the moved file. A breadcrumb file helps your users understand why a file was moved from its original location. For each moved file, the system creates a breadcrumb file in the source location named `<filename>-breadcrumb-<date>.txt`. You can add text in the dialog box that will be added to the breadcrumb file to indicate the location where the file was moved and the user who moved the file.



Note that the subdirectory structure from the source file is recreated on the destination share when the file is moved so it is easier to understand where the file was moved from. If a file with the same name exists in the destination location, the file will not be moved.



You can't move files that reside in databases.

## Requirements

- You must have permissions to move files. [Learn about user access to compliance information.](#)
- The source files can be located in the following data sources: On-premises ONTAP, Cloud Volumes ONTAP, Azure NetApp Files, File Shares, and SharePoint Online.
- You can move a maximum of 15 million files at a time.
- Only files which are 50 MB or smaller are moved.
- The destination NFS share must allow access from the BlueXP classification instance IP address.

## Steps

1. In the Data Investigation results pane, select the file, or files, that you want to move.

<input type="checkbox"/>	File Name	Personal	Sensitive Personal	Data Subjects	File Type	
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	16	PDF
<input checked="" type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF
<input type="checkbox"/>	Expense Report EXP-TPO-106038887654	cvo	6	3	6	PDF

- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).
- To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 Items on this page selected** [Select all Items in list \(63K Items\)](#), click **Select all items in list (xxx items)**.

2. From the button bar, click **Move**.

### Move Files (63)

The files will be moved to the destination folder you provide and will no longer be available at their current location.

Moving files is supported only to destination folders in NFS Shares. Any NFS Share is supported, no matter where it is hosted, as long as the share's export policy allows access from the data connector instance IP address.

---

The status of this action will appear in the Action Status:

---


**Enter the NFS destination folder path to continue**

---

**Leave breadcrumb**

A breadcrumb file helps your users understand why a file was moved from its original location. For each moved file, the system creates a breadcrumb file in the source location named `<filename>-breadcrumb-<date>.txt`.

---

 Max length should be maximum 400 characters

---

3. In the *Move Files* dialog, enter the name of the NFS share where all selected files will be moved in the format `<host_name>:/<share_path>`.
4. If you want to leave a breadcrumb file, check the *Leave breadcrumb* box. You can enter text in the dialog box to indicate the location where the file was moved and the user who moved the file, and any other information, such as the reason the file was moved.
5. Click **Move Files**.

Note that you can also move an individual file when viewing the metadata details for a file. Just click **Move File**.



### Delete source files

You can permanently remove source files that seem insecure or too risky to leave in your storage system, or that you've identified as a duplicate. This action is permanent and there is no undo or restore.



You can't delete files that reside in databases. All other data sources are supported.

Deleting files requires the following permissions:

- For NFS data - the export policy needs to be defined with write permissions.
- For CIFS data - the CIFS credentials need to have write permissions.
- For S3 data - the IAM role must include the following permission: `s3:DeleteObject`.

### Delete source files manually

#### Requirements

- You must have permissions to delete files. [Learn about user access to compliance information.](#)
- You can delete a maximum of 100,000 files at a time.

#### Steps

1. In the Data Investigation results pane, select the file, or files, that you want to delete.



- To select individual files, check the box for each file ( Volume\_1).
- To select all files on the current page, check the box in the title row ( File Name).

To select all files on all pages, check the box in the title row ( File Name), and then in the pop-up message **All 20 items on this page selected** [Select all items in list \(63K items\)](#), click **Select all items in list (xxx items)**.

- From the button bar, click **Delete**.
- Because the delete operation is permanent, you must type "**permanently delete**" in the subsequent *Delete File* dialog and click **Delete File**.

You can view the progress of the delete operation in the [Actions Status](#) pane.

Note that you can also delete an individual file when viewing the metadata details for a file. Just click **Delete file**.



## Add personal data identifiers to your BlueXP classification scans


BlueXP classification provides many ways for you to add a custom list of "personal data" that BlueXP classification will identify in future scans, giving you the full picture about where potentially sensitive data resides in *all* your organizations' files.

 To create a custom classification in version 1.43 and later, see [Create a custom classification](#).

 This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

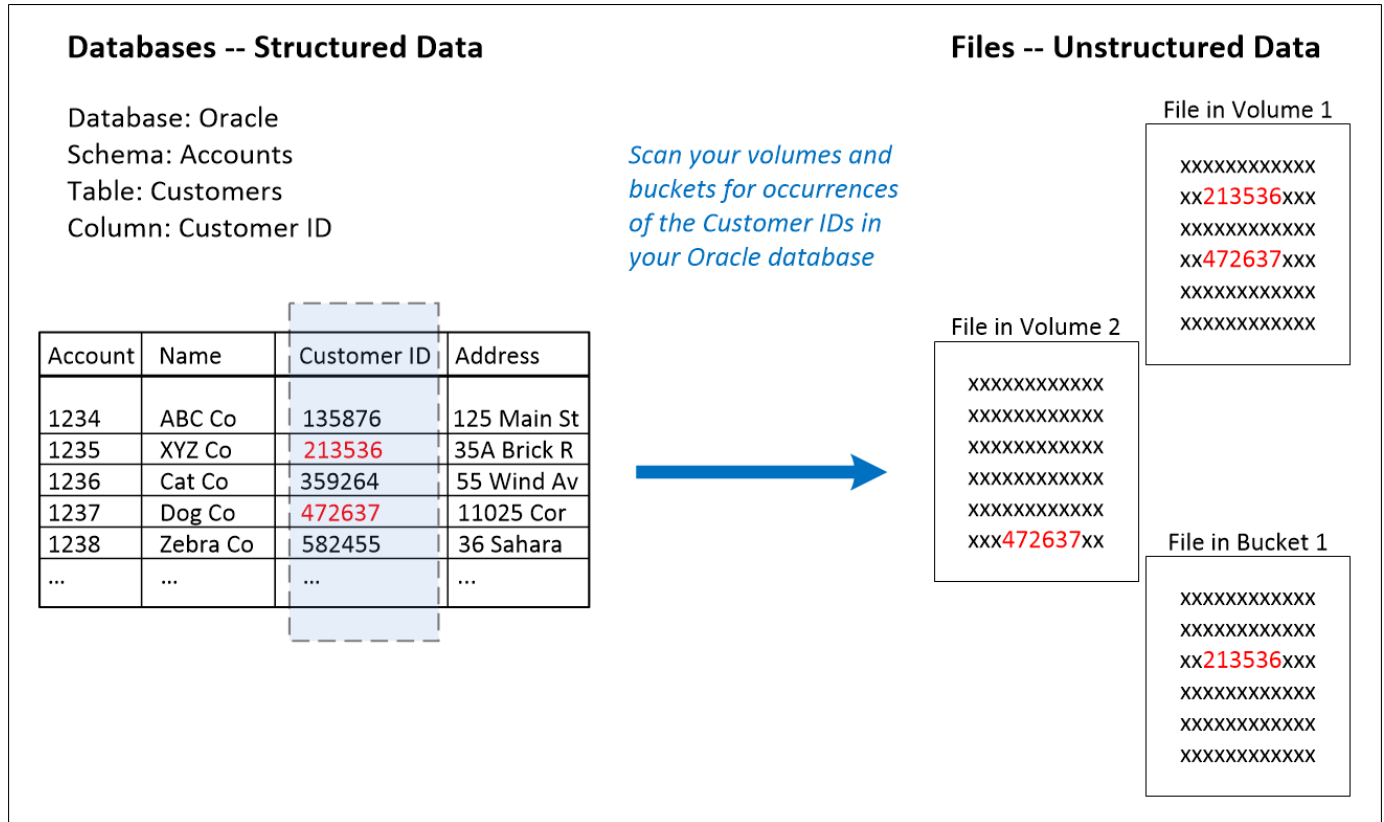
- You can add unique identifiers based on specific columns in databases you are scanning.
- You can add custom keywords from a text file — these words are identified within your data.
- You can add a personal pattern using a regular expression (regex) — the regex is added to the existing predefined patterns.
- You can add custom categories to identify where specific categories of information are found in your data.

All of these mechanisms to add custom scanning criteria are supported in all languages.

 The capabilities described in this section are available only if you have chosen to perform a full classification scan on your data sources. Data sources that have had a mapping-only scan do not show file-level details.

## Add custom personal data identifiers from your databases

*Data Fusion* allows you to scan your organization's data to identify whether unique identifiers from your databases are found in any of your other data sources. You can choose the additional identifiers that BlueXP classification will look for in its scans by selecting a specific column or columns in a database table. For example, the diagram below shows how data fusion is used to scan your volumes, buckets, and databases for occurrences of all your Customer IDs from your Oracle database.



As you can see, two unique Customer IDs have been found in two volumes and in one S3 bucket. Any matches in database tables will also be identified.


Note that since you're scanning your own databases, whatever language your data is stored in will be used to identify data in future BlueXP classification scans.

### Steps

You must have [added at least one database server](#) to BlueXP classification before you can add data fusion sources.

1. In the Configuration page, click **Manage Data Fusion** in the database where the source data resides.

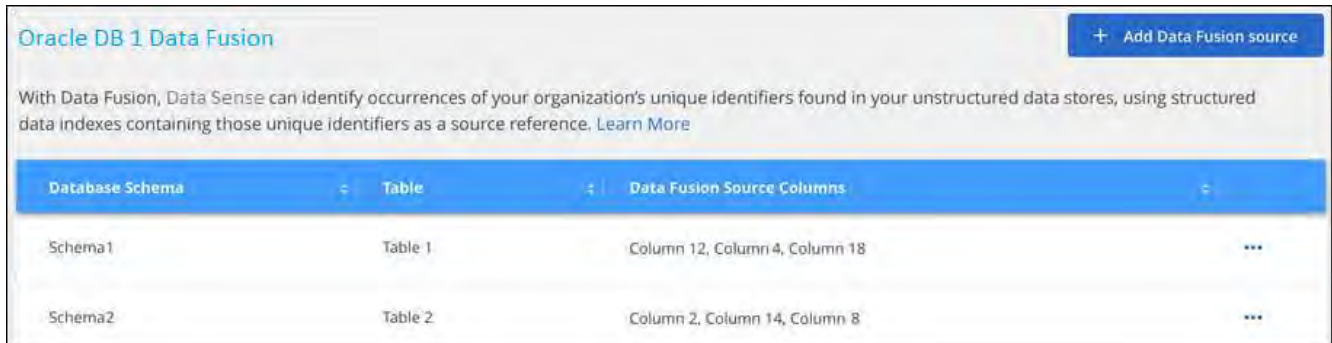




2. Click **Add Data Fusion source** on the next page.
3. In the *Add Data Fusion Source* page:
  - a. Select the Database Schema from the drop-down menu.
  - b. Enter the Table name in that schema.
  - c. Enter the Column, or Columns, that contain the unique identifiers you want to use.


When adding multiple columns, enter each column name, or table view name, on a separate line.

4. Click **Add Data Fusion Source**.



## Results

After the next scan, the results will include this new information in the Compliance Dashboard under the "Personal Results" section, and in the Investigation page in the "Personal Data" filter. The name you used for the classifier appears in the filter list, for example `Customers.CustomerID`.



## Delete a Data Fusion source

If at some point you decide not to scan your files using a certain Data Fusion source, you can select the source row from the Data Fusion inventory page and click **Delete Data Fusion Source**.



## Add custom keywords from a list of words

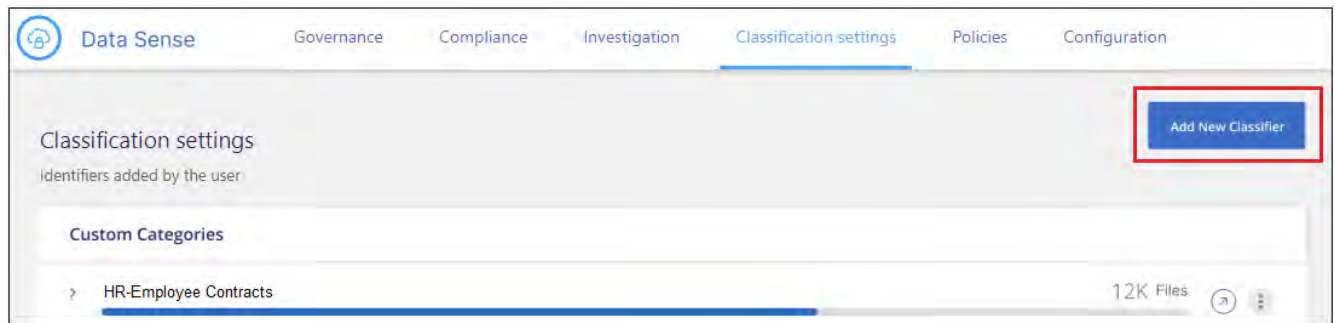
You can add custom keywords to BlueXP classification so that it will identify where that information is found in your data. You add the keywords just by entering each word you want BlueXP classification to recognize. The keywords are added to the existing predefined keywords that BlueXP classification already uses, and the results will be visible under the personal patterns section.

For example, you may want to see where internal Product Names are mentioned in all of your files to make sure these names are not accessible in locations that are not secure.

After updating the custom keywords, BlueXP classification will restart scanning all data sources. After the scan has completed, the new results will appear in the BlueXP classification Compliance Dashboard under the "Personal Results" section, and in the Investigation page in the "Personal Data" filter.

## Steps

1. From the *Classification settings* tab, click **Add New Classifier** to launch the *Add Custom Classifier* wizard.



2. In the *Select type* page, enter the name of the classifier, provide a brief description, select **Personal identifier**, and then click **Next**.

The name you enter will appear in the BlueXP classification UI as the heading for scanned files that match the classifier requirements, and as the name of the filter in the Investigation page.

You can also check the box to "Mask detected results in the system" so the full result won't appear in the UI. For example, you may want to do this to hide full credit card numbers or similar personal data (the mask would appear in the UI like this: "\*\*\*\* \* 3434").

1 Select type      2 Select tool      3 Create Logic

## Select type

Select the type of classifier that you want to add to the system, and provide the name and description. Data Sense re-scans all your data sources after you add a new classifier. When the scan is complete, all matching results are displayed in the "Classification Settings" dashboard and in other Data Sense pages.

Classifier name

Internal Product Names

Description

Identify internal product names found in all files

**Personal identifier**

The classifier will be added to the system as a new personal identifier. Any matches are considered "personal data", and they are added to the results that are displayed in the Personal Results page and in the Investigation page. [See the list of personal data that Data Sense identifies by default.](#)

Mask detected results in the system

**Category**

The classifier will be added to the system as a new Category. Any matches are added to the results that are displayed in the Categories page and in the Investigation page. [See the list of categories that Data Sense identifies by default.](#)

Previous      **Next**

3. In the *Select Data Analysis Tool* page, select **Custom keywords** as the method you want to use to define the classifier, and then click **Next**.



## Select Data Analysis Tool

Select the tool that will be used to build the list of words, or patterns, that Data Sense will attempt to match in your data sources.

**Custom keywords** ⓘ  
Create a custom personal pattern based on a list of keywords that you provide.


**Custom regular expression** ⓘ  
Create a custom personal pattern based on a regular expression that you define.

**DB fusion** ⓘ  
Create a custom personal pattern based on the values found in specific columns in your scanned database tables. This allows you to identify whether unique identifiers from your databases are found in any of your other data sources.

[Previous](#) [Next](#)

4. In the *Create Logic* page, enter the keywords you want to recognize - each word on a separate line - and click **Validate**.

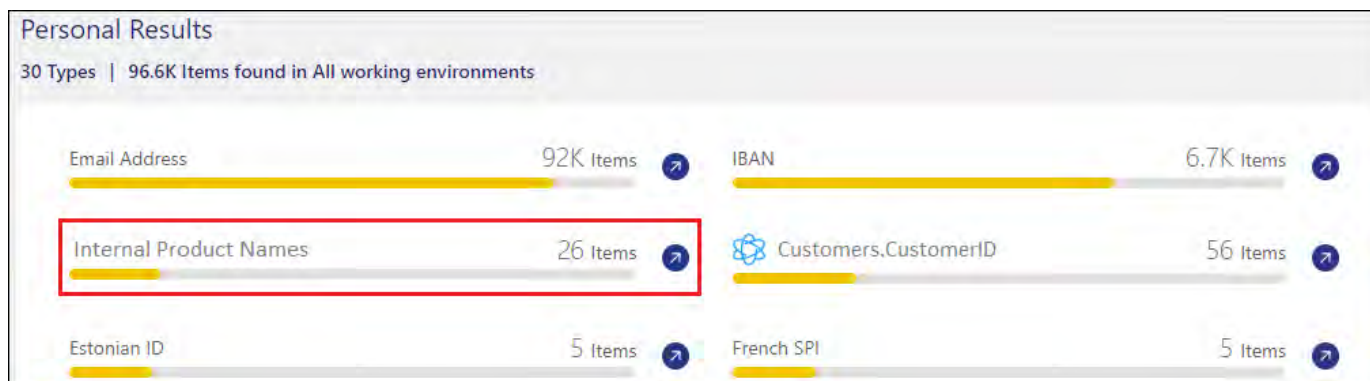
The screenshot below shows internal Product Names (different types of owls). The BlueXP classification search for these items is not case sensitive.



5. Click **Done** and BlueXP classification starts to rescan your data.

## Results

After the scan is complete, the results will include this new information in the Compliance Dashboard under the "Personal Results" section, and in the Investigation page in the "Personal Data" filter.



As you can see, the name of the classifier is used as the name in the Personal Results panel. In this manner you can activate many different groups of keywords and see the results for each group.

## Add custom personal data identifiers using a regex

You can add a personal pattern to identify specific information in your data using a custom regular expression (regex). This allows you to create a new custom regex to identify new personal information elements that don't yet exist in the system. The regex is added to the existing predefined patterns that BlueXP classification

already uses, and the results will be visible under the personal patterns section.

For example, you may want to see where your internal Product IDs are mentioned in all of your files. If the Product ID has a clear structure, for example, it is a 12-digit number that starts with 201, you can use the custom regex feature to search for it in your files. The regular expression for this example is `\b201\d{9}\b`.

After adding the regex, BlueXP classification will restart scanning all data sources. After the scan has completed, the new results will appear in the BlueXP classification Compliance Dashboard under the "Personal Results" section, and in the Investigation page in the "Personal Data" filter.

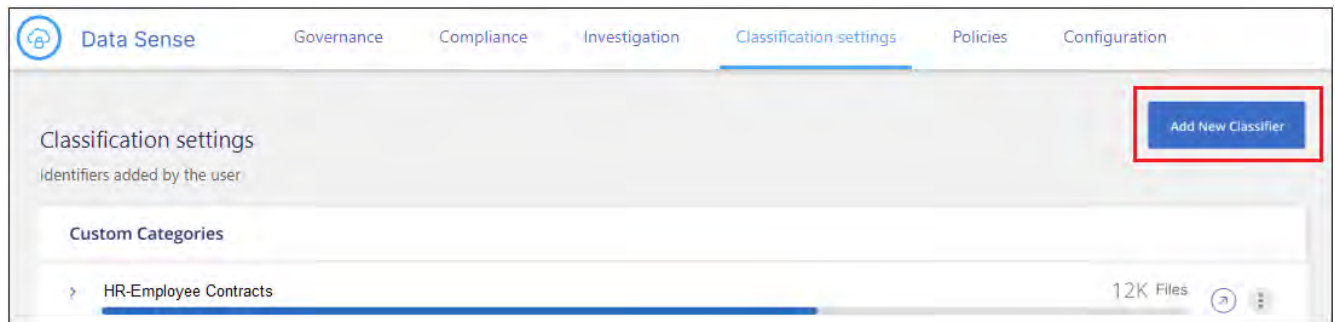
If you need assistance in building the regular expression, refer to [Regular expressions 101](#). Choose **Python** for the Flavor to see the types of results BlueXP classification will match from the regular expression. The [Python Regex Tester page](#) is also useful by displaying a graphical representation of your patterns.



BlueXP classification doesn't support pattern flags when creating a regex. This means you should not use "/".

## Steps

1. From the *Classification settings* tab, click **Add New Classifier** to launch the *Add Custom Classifier* wizard.



2. In the *Select type* page, enter the name of the classifier, provide a brief description, select **Personal identifier**, and then click **Next**.

The name you enter will appear in the BlueXP classification UI as the heading for scanned files that match the classifier requirements, and as the name of the filter in the Investigation page. You can also check the box to "Mask detected results in the system" so the full result won't appear in the UI. For example, you may want to do this to hide full credit card numbers or similar personal data.

1 Select type      2 Select tool      3 Create Logic

## Select type

Select the type of classifier that you want to add to the system, and provide the name and description. Data Sense re-scans all your data sources after you add a new classifier. When the scan is complete, all matching results are displayed in the "Classification Settings" dashboard and in other Data Sense pages.

Classifier name

Internal Product ID

Description

Identify internal product IDs found in all files

**Personal identifier**

The classifier will be added to the system as a new personal identifier. Any matches are considered "personal data", and they are added to the results that are displayed in the Personal Results page and in the Investigation page. [See the list of personal data that Data Sense identifies by default.](#)

Mask detected results in the system

**Category**

The classifier will be added to the system as a new Category. Any matches are added to the results that are displayed in the Categories page and in the Investigation page. [See the list of categories that Data Sense identifies by default.](#)

Previous      **Next**

3. In the *Select Data Analysis Tool* page, select **Custom regular expression** as the method you want to use to define the classifier, and then click **Next**.

## Select Data Analysis Tool

Select the tool that will be used to build the list of words, or patterns, that Data Sense will attempt to match in your data sources.

**Custom keywords** ⓘ  
Create a custom personal pattern based on a list of keywords that you provide.

**Custom regular expression** ⓘ  
Create a custom personal pattern based on a regular expression that you define.

**DB fusion** ⓘ  
Create a custom personal pattern based on the values found in specific columns in your scanned database tables. This allows you to identify whether unique identifiers from your databases are found in any of your other data sources.

4. In the *Create Logic* page, enter the regular expression and any proximity words, and click **Done**.
  - a. You can enter any legal regular expression. Click the **Validate** button to have BlueXP classification verify that the regular expression is valid, and that it is not too broad — meaning it will return too many results.
  - b. Optionally, you can enter some proximity words to help refine the accuracy of the results. These are words that will typically be found within 300 characters of the pattern you are searching for (either before or after the found pattern). Enter each word, or phrase, on a separate line.



## Create Logic

Create logic for the new identifier, based on regular expression and keywords that should be detected.

---

**Regular expression** ⓘ

Add the pattern that should be detected to identify specific information in your data, using a custom regular expression.

Validate

✔ **Success:** Regular expression is valid.

**Proximity words** - To improve the detection accuracy, insert phrases that must appear near by the regular expression's match.

Previous

Done

### Results

The classifier is added and BlueXP classification starts to rescan all your data sources. You are returned to the Custom Classifiers page where you can view the number of files that have matched your new classifier. Results from scanning all of your data sources will take some time depending on the number of files that need to be scanned.

Data Sense
Governance
Compliance
Investigation
Classification settings
Policies
Configuration

### Classification settings

Add New Classifier

Identifiers added by the user

**Custom Categories**

> HR - Employee Contracts
7.5K Files

**Personal information**

> Internal Product ID
12K Files

### Add custom categories

BlueXP classification takes the data that it scans and divides it into different types of categories. Categories are topics based on artificial intelligence analysis of the content and metadata of each file. [See the list of](#)

[predefined categories](#).

Categories can help you understand what's happening with your data by showing you the types of information that you have. For example, a category like *resumes* or *employee contracts* may include sensitive data. When you investigate the results, you might find that employee contracts are stored in an insecure location. You can then correct that issue.

You can add custom categories to BlueXP classification so you can identify where categories of information that are unique for your data estate are found in your data. You add each category by creating "training" files that contain the categories of data that you want to identify, and then have BlueXP classification scan those files to "learn" through AI so that it can identify that data in your data sources. The categories are added to the existing predefined categories that BlueXP classification already identifies, and the results are visible under the Categories section.

For example, you may want to see where compressed installation files in .gz format are located in your files so that you can remove them, if necessary.

After updating the custom categories, BlueXP classification will restart scanning all data sources. After the scan has completed, the new results will appear in the BlueXP classification Compliance Dashboard under the "Categories" section, and in the Investigation page in the "Category" filter. [See how to view files by categories](#).

### Before you begin


You'll need to create a minimum of 25 training files that contain samples of the categories of data that you want BlueXP classification to recognize. The following file types are supported:

.CSV, .DOC, .DOCX, .GZ, .JSON, .PDF, .PPTX, .RTF, .TXT, .XLS, .XLSX, Docs, Sheets, and Slides

The files must be a minimum of 100 bytes, and they must be located in a folder that is accessible by BlueXP classification.

### Steps

1. From the *Classification settings* tab, click **Add New Classifier** to launch the *Add Custom Classifier* wizard.



2. In the *Select type* page, enter the name of the classifier, provide a brief description, select **Category**, and then click **Next**.

The name you enter will appear in the BlueXP classification UI as the heading for scanned files that match the category of data you are defining, and as the name of the filter in the Investigation page.

1 Select type      2 Select tool      3 Create Logic

## Select type

Select the type of classifier that you want to add to the system, and provide the name and description. Data Sense re-scans all your data sources after you add a new classifier. When the scan is complete, all matching results are displayed in the "Classification Settings" dashboard and in other Data Sense pages.

Classifier name

Compressed installer files

Description

Installation files in .GZ format

Personal identifier

The classifier will be added to the system as a new personal identifier. Any matches are considered "personal data", and they are added to the results that are displayed in the Personal Results page and in the Investigation page. [See the list of personal data that Data Sense identifies by default.](#)

Mask detected results in the system

**Category**

The classifier will be added to the system as a new Category. Any matches are added to the results that are displayed in the Categories page and in the Investigation page. [See the list of categories that Data Sense identifies by default.](#)

Previous      **Next**

3. In the *Create Logic* page, make sure you have the learning files prepared, and then click **Select files**.

## Create Logic

**AI-based similarity training** ⓘ

- Insert NFS folder path
- The folder should contain minimum 25 files and maximum 1000 files that will be used for the AI training.
- Supported file types: pdf, docx, doc, pptx, xls, xlsx, csv, txt, gz, rtf, docs, sheets, slides, json
- The keywords are not case sensitive
- Minimum file size: 100B

Compressed Installer files

**Select Files**

4. Enter the IP address of the volume, and the path where the training files are located, and click **Add**.



Insert folder path that contains at least 25 files for the training

Enter the IP address and volume name, along with the path to the location of the training files.

IP: XXX.XXX.XXX.XXX:/VolumeName

Training Data - Folder path: folder/path/

**Add** Cancel

5. Verify that the training files were recognized by BlueXP classification. Click the **x** to remove any training files that do not meet the requirements. Then click **Done**.

### Create Logic

**AI-based similarity training**

- Insert NFS folder path
- The folder should contain minimum 25 files and maximum 1000 files that will be used for the AI training.
- Supported file types: pdf, docx, doc, pptx, xls, xlsx, csv, txt, gz, rtf, docs, sheets, slides, json
- The keywords are not case sensitive.
- Minimum file size: 100B

Select Files

**Compressed Installer files**

Total uploaded files: 54

File name	File Size	File Type	Reliability	included in training
File1	56	File type	Sufficient	x
File2	22	File type	Sufficient	x
File3	43	File type	Sufficient	x
File4	11	File type	Sufficient	x

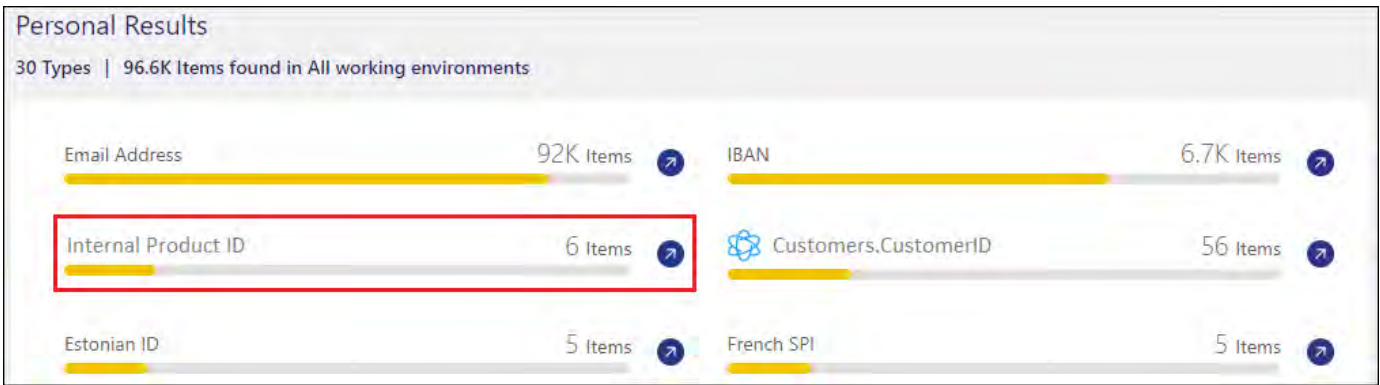
Previous Done

## Results

The new category is created as defined by the training files and added to BlueXP classification. Then BlueXP classification starts to rescan all your data sources to identify files that fit into this new category. You are returned to the Custom Classifiers page where you can view the number of files that have matched your new category. Results from scanning all of your data sources will take some time depending on the number of files that need to be scanned.

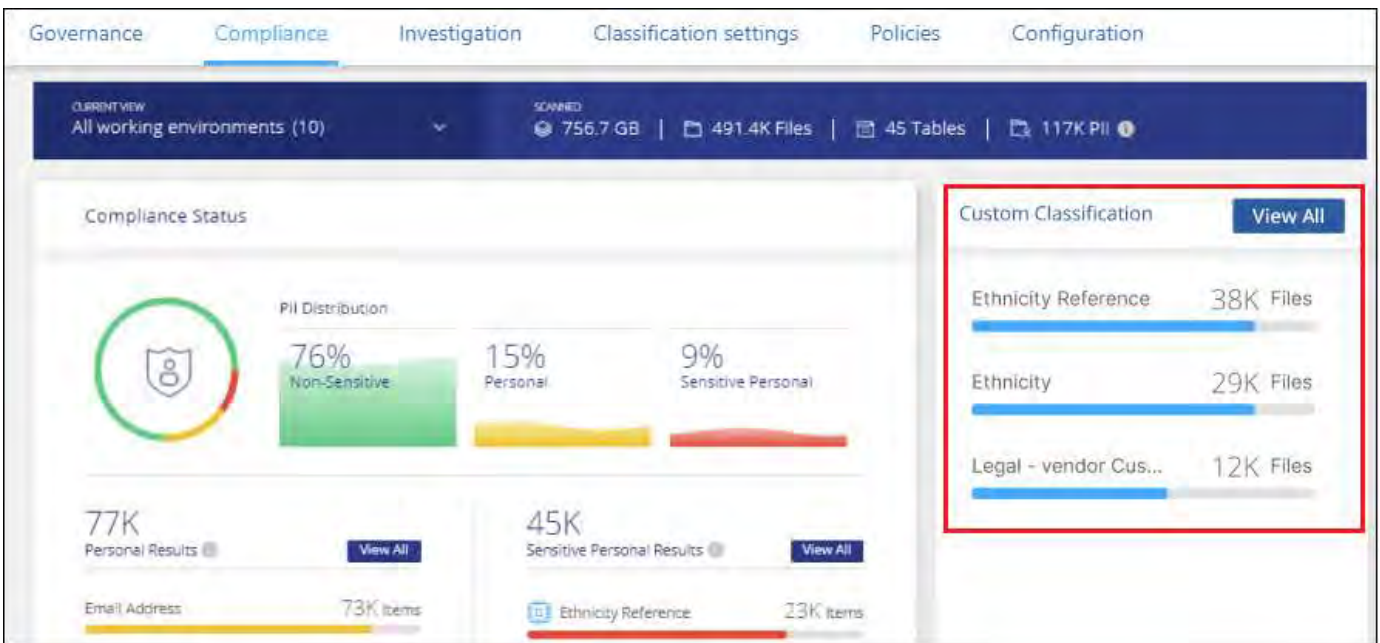
## View results from your custom classifiers

You can view the results from any of your custom classifiers in the Compliance Dashboard and in the Investigation page. For example, this screenshot shows the matched information in the Compliance Dashboard under the "Personal Results" section.



Click the  button to see the detailed results in the Investigation page.

Additionally, all of your custom classifier results appear in the Custom Classifiers tab, and the top 6 custom classifier results are displayed in the Compliance Dashboard, as shown below.

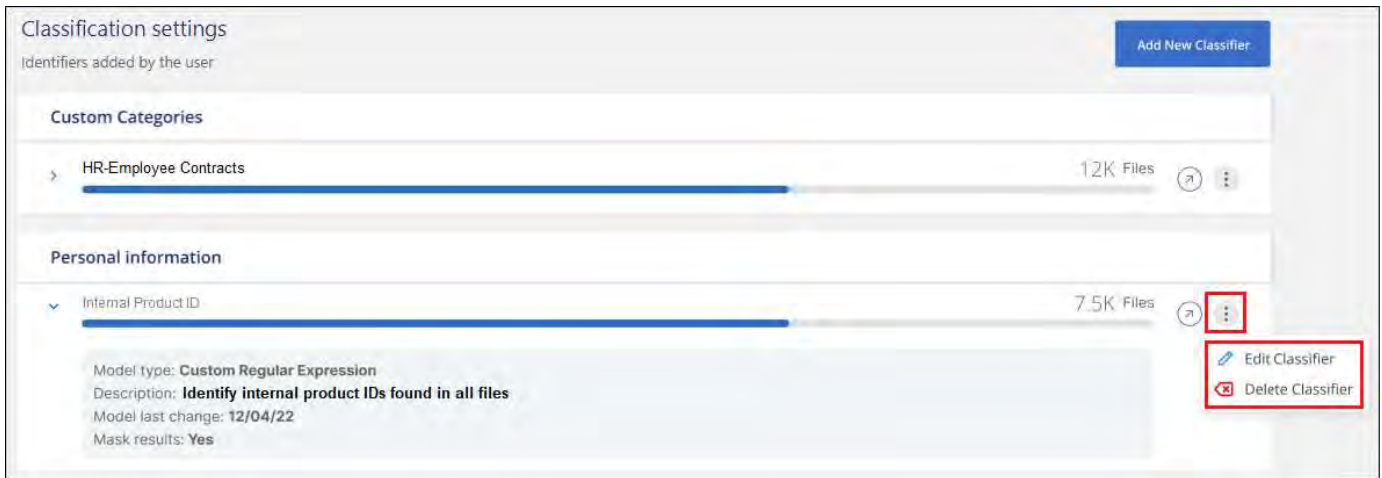


### Manage custom classifiers

You can change any of the custom classifiers that you have created by using the **Edit Classifier** button.

 You can't edit Data Fusion classifiers at this time.

And if you decide at some later point that you don't need BlueXP classification to identify the custom patterns that you added, you can use the **Delete Classifier** button to remove each item.



## View the status of your compliance actions in BlueXP classification

When you run an asynchronous action from the Investigation Results pane across many files, for example, moving or deleting 100 files, the process can take some time. You can monitor the status of these actions in the *Action Status* pane so you'll know when it has been applied to all files.

This allows you to see the actions that have completed successfully, those currently in progress, and those that have failed so you can diagnose and fix any problems. Note that short operations that complete quickly, such as moving a single file, do not appear in the Actions Status pane.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

The status can be:

- Success - A BlueXP classification action is finished and all items succeeded.
- Partial Success - A BlueXP classification action is finished and some items failed and some succeeded.
- In Progress - The action is still in progress.
- Queued - The action has not started.
- Canceled - The action has been canceled.
- Failed - The action failed.

Note that you can Cancel any actions that have the "Queued" or "In Progress" status.

### Steps

1.

In the bottom-right of the BlueXP classification UI you can see the **Actions Status** button



2. Click this button and the most recent 20 actions are listed.

You can click the name of an action to view details corresponding to that operation.

## Audit the history of BlueXP classification actions

BlueXP classification logs management activities that have been performed on files from all the working environments and data sources that BlueXP classification is scanning. BlueXP classification also logs the activities when deploying the BlueXP classification instance.

You can view the contents of the BlueXP classification audit log files, or download them, to see what file changes have occurred, and when. For example, you can see what request was issued, the time of the request, and details such as source location in case a file was deleted, or source and destination location in case a file was moved.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

### Log file contents

Each line in the audit log contains information in this format:

```
<full date> | <status> | ds_audit_logger | <module> | 0 | 0 | File <full file path> deleted from device <device path> - <result>
```

- Date and time - full timestamp for the event
- Status - INFO, WARNING
- Action type (delete, copy, move, create saved search, update saved search, rescan files, download JSON report, etc.)
- File name (if the action is relevant to a file)
- Details for the action - what was done: depends on the action
  - Saved search name
  - For move - Source and destination
  - For copy - Source and destination
  - For tag - tag name
  - For assign to - user name
  - For email alert - email address / account

For example, the following lines from the log file show a successful copy operation and a failed copy operation.

```
2022-06-06 15:23:08,910 | INFO | ds_audit_logger | es_scanned_file | 237 | 49 | Copy file /CIFS_share/data/dopl/random_positives.tsv from device 10.31.133.183 (type: SMB_SHARE) to device 10.31.130.133:/export_reports (NFS_SHARE) - SUCCESS
2022-06-06 15:23:08,968 | WARNING | ds_audit_logger | es_scanned_file | 239 | 153 | Copy file /CIFS_share/data/compliance-netapp.tar.gz from device 10.31.133.183 (type: SMB_SHARE) to device 10.31.130.133:/export_reports (NFS_SHARE) - FAILURE
```

## Log file locations

The management audit log files are located on the BlueXP classification machine in:  
`/opt/netapp/audit_logs/`

The installation audit log files are written to `/opt/netapp/install_logs/`

Each log file can be a maximum of 10 MB in size. When that limit is reached, a new log file is started. The log files are named "DataSense\_audit.log", "DataSense\_audit.log.1", "DataSense\_audit.log.2", and so on. A maximum of 100 log files are retained on the system - older log files are deleted automatically after the maximum has been reached.

## Access the log files

You'll need to log into the BlueXP classification system to access the log files. See how to [log in to the BlueXP classification system](#) depending on whether you manually installed the software on a Linux machine or if you deployed the instance in the cloud.

## Reducing the BlueXP classification scan speed

Data scans have a negligible impact on your storage systems and on your data. However, if you are concerned with even a very small impact, you can configure BlueXP classification to perform "slow" scans.

When enabled, slow scanning is used on all data sources - you can't configure slow scanning for a single working environment or data source.



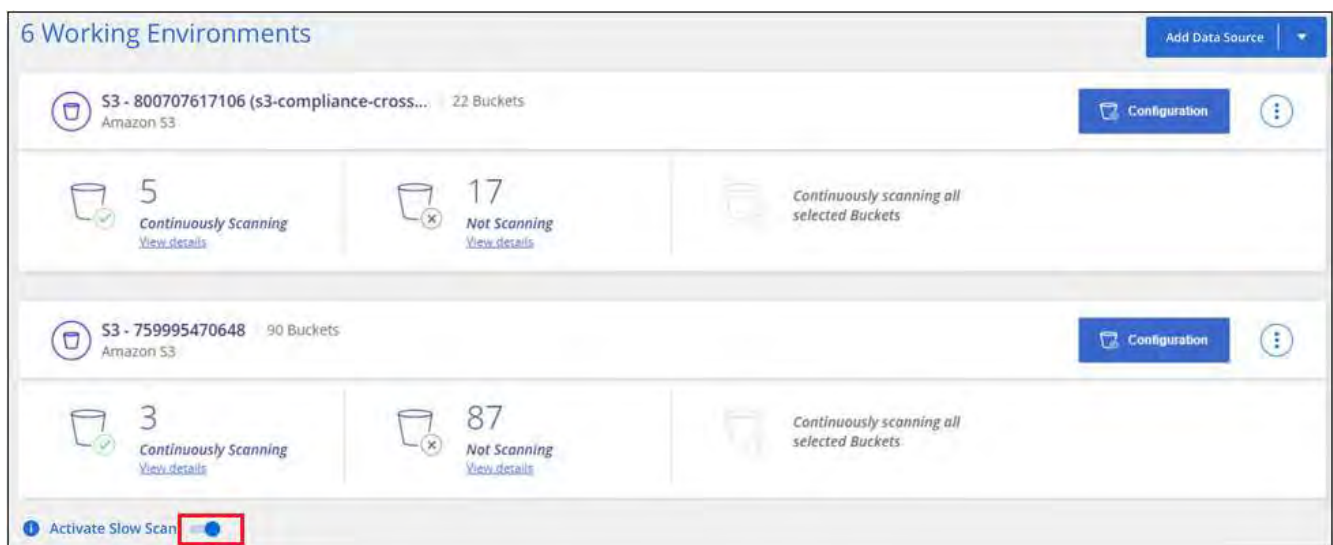
The scan speed can't be reduced when scanning databases.



This information is relevant only for BlueXP classification legacy versions 1.30 and earlier.

## Steps

1. From the bottom of the *Configuration* page, move the slider to the right to activate slow scanning.



The top of the Configuration page indicates that slow scanning is enabled.




2. You can disable slow scanning by clicking **Disable** from this message.

## Remove a OneDrive, SharePoint, or Google Drive account from BlueXP classification

If you no longer want to scan user files from a certain OneDrive account, from a specific SharePoint account, or from a Google Drive account, you can delete the account from the BlueXP classification interface and stop all scans.

### Steps

1. From the *Configuration* page, select the  button in the row for the OneDrive, SharePoint, or Google Drive account, then select **Remove OneDrive Account**, **Remove SharePoint Account**, or **Remove Google Drive account**.



2. Click **Delete Account** from the confirmation dialog.



# Reference

## Supported BlueXP classification instance types

BlueXP classification software must run on a host that meets specific operating system requirements, RAM requirements, software requirements, and so on. When deploying BlueXP classification in the cloud, we recommend that you use a system with the "large" characteristics for full functionality.

You can deploy BlueXP classification on a system with fewer CPUs and less RAM, but there are some limitations when using these less powerful systems. [Learn about these limitations.](#)

In the following tables, if the system marked as "default" is not available in the region where you are installing BlueXP classification, the next system in the table will be deployed.

### AWS instance types

System size	Specs	Instance type
Extra Large	32 CPUs, 128 GB RAM, 1 TiB gp3 SSD	<a href="#">m6i.8xlarge</a> (default)
Large	16 CPUs, 64 GB RAM, 500 GiB SSD	<a href="#">m6i.4xlarge</a> (default) <a href="#">m6a.4xlarge</a> <a href="#">m5a.4xlarge</a> <a href="#">m5.4xlarge</a> <a href="#">m4.4xlarge</a>
Medium	8 CPUs, 32 GB RAM, 200 GiB SSD	<a href="#">m6i.2xlarge</a> (default) <a href="#">m6a.2xlarge</a> <a href="#">m5a.2xlarge</a> <a href="#">m5.2xlarge</a> <a href="#">m4.2xlarge</a>
Small	8 CPUs, 16 GB RAM, 100 GiB SSD	<a href="#">c6a.2xlarge</a> (default) <a href="#">c5a.2xlarge</a> <a href="#">c5.2xlarge</a> <a href="#">c4.2xlarge</a>

### Azure instance types

System size	Specs	Instance type
Extra Large	32 CPUs, 128 GB RAM, OS Disk (2,048 GiB, min 250 MB/s throughput), and Data Disk (1 TiB SSD, min 750 MB/s throughput)	<a href="#">Standard_D32_v3</a> (default)
Large	16 CPUs, 64 GB RAM, 500 GiB SSD	<a href="#">Standard_D16s_v3</a> (default)

## GCP instance types

System size	Specs	Instance type
Large	16 CPUs, 64 GB RAM, 500 GiB SSD	n2-standard-16 (default) n2d-standard-16 n1-standard-16

## Metadata collected from data sources in BlueXP classification

BlueXP classification collects certain metadata when performing classification scans on the data from your data sources and working environments. BlueXP classification can access most of the metadata we need to classify your data, but there are some sources where we are unable to access the data we need.

	Metadata	CIFS	NFS
<b>Time stamps</b>	<i>Creation time</i>	Available	Not available (Unsupported in Linux)
	<i>Last access time</i>	Available	Available
	<i>Last modify time</i>	Available	Available
<b>Permissions</b>	<i>Open permissions</i>	If "EVERYONE" group has access to the file, it is considered "Open to organization"	If "Others" has access to the file, it is considered "Open to organization"
	<i>Users/group access</i>	Users and group information is taken from LDAP	Not available (NFS users are usually managed locally on the server, therefore, the same individual can have a different UID in each server)



- BlueXP classification does not extract the "last accessed time" from the database data sources.
- Older versions of the Windows OS (for example, Windows 7 and Windows 8) disable the collection of the "last accessed time" attribute by default because it can impact system performance. When this attribute is not collected, BlueXP classification analytics that are based on "last accessed time" will be impacted. You can enable the collection of the last access time on these older Windows systems if needed.

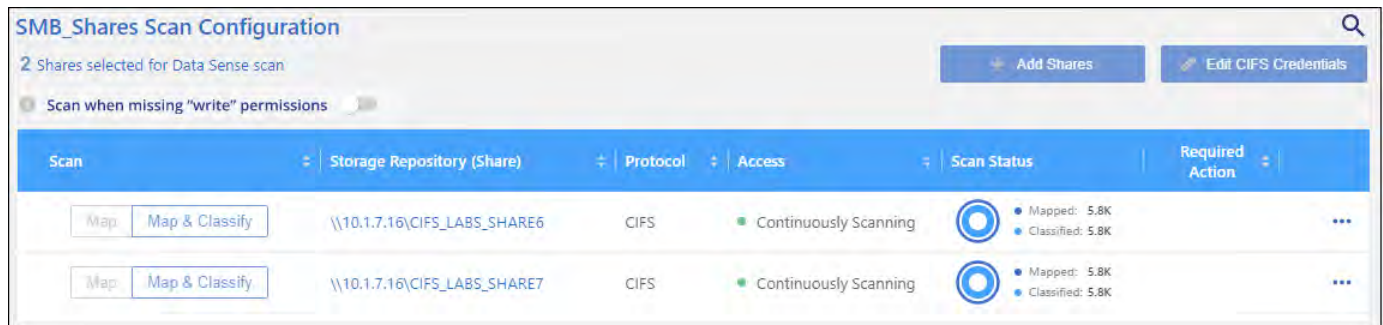
### Last access time timestamp

When BlueXP classification extracts data from file shares, the operating system considers it as accessing the data and it changes the "last access time" accordingly. After scanning, BlueXP classification attempts to revert the last access time to the original timestamp. If BlueXP classification doesn't have write attributes permissions in CIFS, or write permissions in NFS, the system can't revert the last access time to the original timestamp.



ONTAP volumes configured with SnapLock have read-only permissions and also can't revert the last access time to the original timestamp.

By default, if BlueXP classification doesn't have these permissions, the system won't scan those files in your volumes because BlueXP classification can't revert the "last access time" to the original timestamp. However, if you don't care if the last access time is reset to the original time in your files, you can click the **Scan when missing "write attributes" permissions** switch at the bottom of the Configuration page so that BlueXP classification will scan the volumes regardless of permissions.



This functionality is applicable to On-premises ONTAP systems, Cloud Volumes ONTAP, Azure NetApp Files, FSx for ONTAP, and third-party file shares.

Note that there is a filter in the Investigation page called *Scan Analysis Event* that enables you to display either the files that were not classified because BlueXP classification couldn't revert the last accessed time, or the files that were classified even though BlueXP classification couldn't revert the last access time.

The filter selections are:

- "Not classified — Cannot revert last access time" - This shows the files that were not classified due to missing write permissions.
- "Classified and updated last access time" - This shows the files that were classified and BlueXP classification was unable to reset the last access time back to the original date. This filter is relevant only for environments where you turned **Scan when missing "write attributes" permissions** ON.

If needed, you can export these results to a report so you can see which files are, or aren't, being scanned because of permissions. [Learn more about the Data Investigation Report.](#)

## Log in to the BlueXP classification system

At times you may need to log into the BlueXP classification system so you can access log files or edit configuration files.

When BlueXP classification is installed on a Linux machine on your premises, or on a Linux machine you deployed in the cloud, you can access the configuration file and script directly.

When BlueXP classification is deployed in the cloud, you need to SSH to the BlueXP classification instance. You SSH to the system by entering the user and password, or by using the SSH key you provided during the BlueXP Connector installation. The SSH command is:

```
ssh -i <path_to_the_ssh_key> <machine_user>@<datasense_ip>
```

- <path\_to\_the\_ssh\_key> = location of ssh authentication keys
- <machine\_user>:
  - For AWS: use the <ec2-user>
  - For Azure: use the user created for the BlueXP instance
  - For GCP: use the user created for the BlueXP instance
- <datasense\_ip> = IP address of the virtual machine instance

Note that you'll need to modify the security group inbound rules to access the system in the cloud. For details, see:

- [Security group rules in AWS](#)
- [Security group rules in Azure](#)
- [Firewall rules in Google Cloud](#)

## BlueXP classification APIs

The BlueXP classification capabilities that are available through the web UI are also available through the Swagger API.

There are four categories defined within BlueXP classification that correspond to the tabs in the UI:

- Investigation
- Compliance
- Governance
- Configuration

The APIs in the Swagger documentation allow you to search, aggregate data, track your scans, and create actions like copy, move, and more.

### Overview

The API enables you to perform the following functions:

- Export information
  - Everything that is available in the UI can be exported via the API (with the exception of reports)
  - Data is exported in a JSON format (easy to parse and push to 3rd party applications, like Splunk)
- Create queries using "AND" and "OR" statements, include and exclude information, and more.

For example, you can locate files *without* specific Personal Identifiable Information (PII) (functionality not available in the UI). You can also exclude specific fields for the export operation.

- Perform actions
  - Update CIFS credentials

- View and cancel actions
- Re-scan directories
- Export data

The API is secure and it uses the same authentication method as the UI. You can find information on the authentication in: [https://docs.netapp.com/us-en/bluexp-automation/platform/get\\_identifiers.html](https://docs.netapp.com/us-en/bluexp-automation/platform/get_identifiers.html)

## Accessing the Swagger API reference

To get into Swagger you'll need the IP address of the your BlueXP classification instance. In the case of a cloud deployment you'll use the public IP address. Then you'll need to get into this endpoint:

`https://<classification_ip>/documentation`

## Example using the APIs

The following example shows an API call to copy files.

### API Request

You'll initially need to get all the relevant fields and options for a working environment to view all of the filters in the investigation tab.

```
curl -X GET "http://{classification_ip}/api/{classification_version}
/search/options?data_mode=ALL_EXTRACTABLE" -H "accept: application/json"
-H "Authorization: Bearer eyJhbGciOiJSUzI1NiIsInR..... " -H "x-agent-id:
hOXsZNVnA5LsthwMILtjL9xZFYBQxAwMclients"
```

### Response

```
{
  "options": [
    {
      "active_directory_affected": false,
      "data_mode": "ALL_SCANNED",
      "field": "string",
      "is_rulable": true,
      "name": "string",
      "operators": [
        "EQUALS"
      ],
      "optional_values": [
        {}
      ],
      "secondary": {},
      "server_data": false,
      "type": "TEXT"
    }
  ]
}
```

```

]
}
{
  "options": [
    {
      "active_directory_affected": false,
      "data_mode": "ALL_EXTRACTABLE",
      "field": "POLICIES",
      "name": "Policies",
      "operators": [
        "IN",
        "NOT_IN"
      ],
      "server_data": true,
      "type": "SELECT"
    },
    {
      "active_directory_affected": false,
      "data_mode": "ALL_EXTRACTABLE",
      "field": "EXTRACTION_STATUS_RANGE",
      "name": "Scan Analysis Status",
      "operators": [
        "IN"
      ],
      "server_data": true,
      "type": "SELECT"
    },
    {
      "active_directory_affected": false,
      "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
      "field": "SCAN_ANALYSIS_ERROR",
      "name": "Scan Analysis Event",
      "operators": [
        "IN"
      ],
      "server_data": true,
      "type": "SELECT"
    },
    {
      "active_directory_affected": false,
      "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
      "field": "PUBLIC_ACCESS",
      "name": "Open Permissions",
      "operators": [
        "IN",
        "NOT_IN"
      ]
    }
  ]
}

```

```

    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": true,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "USERS_PERMISSIONS_COUNT_RANGE",
    "name": "Number of Users with Access",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": true,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "USER_GROUP_PERMISSIONS",
    "name": "User / Group Permissions",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "FILE_OWNER",
    "name": "File Owner",
    "operators": [
      "EQUALS",
      "CONTAINS"
    ],
    "server_data": true,
    "type": "TEXT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "ENVIRONMENT_TYPE",
    "name": "Working Environment Type",
    "operators": [
      "IN",

```

```

    "NOT_IN"
  ],
  "server_data": true,
  "type": "SELECT"
},
{
  "active_directory_affected": false,
  "data_mode": "ALL_EXTRACTABLE",
  "field": "ENVIRONMENT",
  "name": "Working Environment",
  "operators": [
    "IN",
    "NOT_IN"
  ],
  "server_data": true,
  "type": "SELECT"
},
{
  "active_directory_affected": false,
  "data_mode": "ALL_SCANNED",
  "field": "SCAN_TASK",
  "name": "Storage Repository",
  "operators": [
    "IN",
    "NOT_IN"
  ],
  "server_data": true,
  "type": "SELECT"
},
{
  "active_directory_affected": false,
  "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
  "field": "FILE_PATH",
  "name": "File / Directory Path",
  "operators": [
    "MULTI_CONTAINS",
    "MULTI_EXCLUDE"
  ],
  "server_data": true,
  "type": "MULTI_TEXT"
},
{
  "active_directory_affected": false,
  "data_mode": "ALL_DASHBOARD_EXTRACTABLE",
  "field": "CATEGORY",
  "name": "Category",

```

```

    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "PATTERN_SENSITIVITY_LEVEL",
    "name": "Sensitivity Level",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "NUMBER_OF_IDENTIFIERS",
    "name": "Number of identifiers",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "PATTERN_PERSONAL",
    "name": "Personal Data",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "PATTERN_SENSITIVE",

```

```

    "name": "Sensitive Personal Data",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "DATA_SUBJECT",
    "name": "Data Subject",
    "operators": [
      "EQUALS",
      "CONTAINS"
    ],
    "server_data": true,
    "type": "TEXT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "DIRECTORIES",
    "field": "DIRECTORY_TYPE",
    "name": "Directory Type",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "FILE_TYPE",
    "name": "File Type",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,

```



```

    "data_mode": "ALL_EXTRACTABLE",
    "field": "FILE_SIZE_RANGE",
    "name": "File Size",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "FILE_CREATION_RANGE_RETENTION",
    "name": "Created Time",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "DISCOVERED_TIME_RANGE",
    "name": "Discovered Time",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",
    "field": "FILE_LAST_MODIFICATION_RETENTION",
    "name": "Last Modified",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_FILESYSTEM_EXTRACTABLE",

```

```

    "field": "FILE_LAST_ACCESS_RANGE_RETENTION",
    "name": "Last Accessed",
    "operators": [
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "FILES",
    "field": "IS_DUPLICATE",
    "name": "Duplicates",
    "operators": [
      "EQUALS",
      "IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "FILES",
    "field": "FILE_HASH",
    "name": "File Hash",
    "operators": [
      "EQUALS",
      "IN"
    ],
    "server_data": true,
    "type": "TEXT"
  },
  {
    "active_directory_affected": false,
    "data_mode": "ALL_EXTRACTABLE",
    "field": "USER_DEFINED_STATUS",
    "name": "Tags",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  },
  {
    "active_directory_affected": false,

```

```

    "data_mode": "ALL_EXTRACTABLE",
    "field": "ASSIGNED_TO",
    "name": "Assigned to",
    "operators": [
      "IN",
      "NOT_IN"
    ],
    "server_data": true,
    "type": "SELECT"
  }
]
}

```

We will use that response in our request parameters to filter the desired files we want to copy.

You can apply an action on multiple items. Supported action types include: move, delete, copy, assign to, FlexClone, export data, rescan, and label.

We will create the copy action:

#### API Request

This next API is that action API and it allows you to create multiple actions.

```

curl -X POST "http://
{classification_ip}/api/{classification_version}/actions" -H "accept:
application/json" -H "Authorization: Bearer eyJhbGciOiJSUzI1NiIsInR..... "
-H "x-agent-id: hOXsZNvnA5LsthwMILtjL9xZFYBQxAwMclients " -H "Content-
Type: application/json" -d "{ \"action_type\": \"COPY\", \"data_mode\":
\"FILES\", \"policy_id\": 0, \"request_params\": { destination_nfs_path:
\"{ontap_ip}/{share_name} \" },
\"requested_query\":{\"condition\":\"AND\",\"rules\":[{\"field\":\"ENVIRONMENT_TYPE
\",\"operator\":\"IN\",\"value\":[\"ONPREM\"]},{\"field\":\"CATEGORY\",\"operator\":\"IN\",
\"value\":[\"21\"]}]}}}"

```

#### Response

The response will return the action object, so you can use the get and delete APIs to get status about the action, or to cancel it.

```
{
  "action_type": "COPY",
  "creation_time": "2023-08-08T12:37:21.705Z",
  "data_mode": "FILES",
  "end_time": "2023-08-08T12:37:21.705Z",
  "estimated_time_to_complete": 0,
  "id": 0,
  "policy_id": 0,
  "policy_name": "string",
  "priority": 0,
  "request_params": {},
  "requested_query": {},
  "result": {
    "error_message": "string",
    "failed": 0,
    "in_progress": 0,
    "succeeded": 0,
    "total": 0
  },
  "start_time": "2023-08-08T12:37:21.705Z",
  "status": "QUEUED",
  "title": "string",
  "user_id": "string"
}
```

# Knowledge and support

## Register for BlueXP 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](#)
- [Google Cloud NetApp Volumes](#)

## Support registration overview

There are two forms of registration to activate support entitlement:

- Registering your BlueXP account serial number (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 BlueXP for NetApp support

To register for support and activate support entitlement, one user in your BlueXP organization (or 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 BlueXP organization 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 organization is not registered for support. As long as one user in the organization has followed these steps, then your organization 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.



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 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 organization 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 organization is different than the NSS account that is associated with a BlueXP user login.

These NSS credentials are associated with your specific BlueXP organization ID. Users who belong to the BlueXP organization 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 for BlueXP classification

NetApp provides support for BlueXP and its cloud services in a variety of ways. Extensive free self-support options are available 24/7, 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](#)
- [Google Cloud NetApp Volumes](#)

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 organization (or account), project (or 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 Working Environment


Select 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)  Upload 

No files selected  

### 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

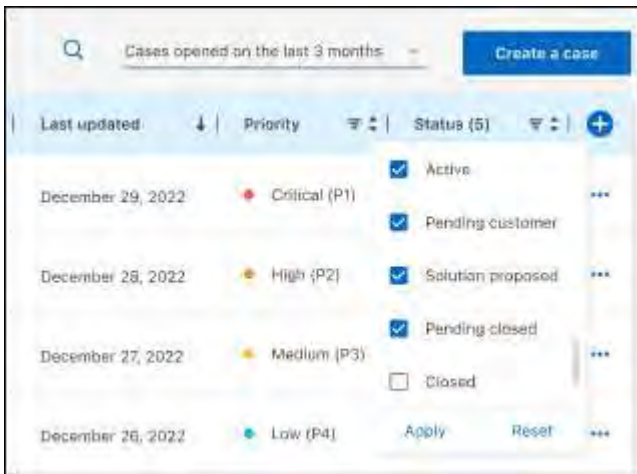
1. In BlueXP, select **Help > Support**.
2. Select **Case Management** and if you're prompted, add your NSS account to BlueXP.


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.



- Filter the contents of the columns.



- Change the columns that appear in the table by selecting  and then choosing the columns that you'd like to display.

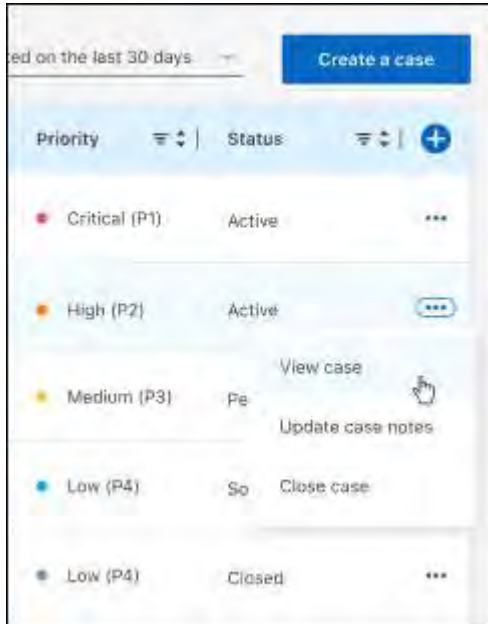


4. Manage an existing case by selecting **...** and selecting one of the available options:

- **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**.



# Frequently asked questions about BlueXP classification

This FAQ can help if you're just looking for a quick answer to a question.

## BlueXP classification service

The following questions provide a general understanding of BlueXP classification.

### How does BlueXP classification work?

BlueXP classification deploys another layer of AI alongside your BlueXP system and storage systems. It then scans the data on volumes, buckets, databases, and other storage accounts and indexes the data insights that are found. BlueXP classification leverages both artificial intelligence and natural language processing, as opposed to alternative solutions that are commonly built around regular expressions and pattern matching.

BlueXP classification uses AI to provide contextual understanding of data for accurate detection and classification. It is driven by AI because it is designed for modern data types and scale. It also understands data context in order to provide strong, accurate, discovery and classification.

[Learn more about how BlueXP classification works.](#)

### Does BlueXP classification have a REST API, and does it work with third-party tools?

Yes, BlueXP classification has a REST API for the supported features in the BlueXP classification version that is part of the BlueXP core platform. See [API documentation](#).

### Is BlueXP classification available through the cloud marketplaces?

BlueXP classification is part of the BlueXP core features, so you do not need to use the marketplaces for this service .

## BlueXP classification scanning and analytics

The following questions relate to BlueXP classification scanning performance and the analytics.

### How often does BlueXP classification scan my data?

While the initial scan of your data might take a little bit of time, subsequent scans only inspect the incremental changes, which reduces system scan times. BlueXP classification scans your data continuously in a round-robin fashion, six repositories at a time, so that all changed data is classified very quickly.

[Learn how scans work.](#)

BlueXP classification scans databases only once per day; databases are not continuously scanned like other data sources.

Data scans have a negligible impact on your storage systems and on your data.

## Does scan performance vary?

Scan performance can vary based on the network bandwidth and the average file size in your environment. It can also depend on the size characteristics of the host system (either in the cloud or on-premises). See [The BlueXP classification instance](#) and [Deploying BlueXP classification](#) for more information.

When initially adding new data sources, you can also choose to perform only a "mapping" (Mapping only) scan instead of a full "classification" (Map & Classify) scan. Mapping can be done on your data sources very quickly because it does not access files to see the data inside. [See the difference between a mapping and classification scan.](#)

## Can I search my data using BlueXP classification?

BlueXP classification offers extensive search capabilities that make it easy to search for a specific file or piece of data across all connected sources. BlueXP classification empowers users to search deeper than just what the metadata reflects. It is a language-agnostic service that can also read the files and analyze a multitude of sensitive data types, such as names and IDs. For example, users can search across both structured and unstructured data stores to find data that may have leaked from databases to user files, in violation of corporate policy. Searches can be saved for later, and policies can be created to search and take action on the results at a set frequency.

Once the files of interest are found, characteristics can be listed, including tags, working environment account, bucket, file path, category (from classification), file size, last modified, permission status, duplicates, sensitivity level, personal data, sensitive data types within the file, owner, file type, file size, created time, file hash, whether the data was assigned to someone seeking their attention, and more. Filters can be applied to screen out characteristics that are not pertinent.

BlueXP classification also has role-based access control (RBAC) to allow files to be moved or deleted, if the right permissions are present. If the right permissions are not present, the tasks can be assigned to someone in the organization who does have the right permissions.

## BlueXP classification management and privacy

The following questions provide information on how to manage BlueXP classification and privacy settings.

### How do I enable or disable BlueXP classification?

First you need to deploy an instance of BlueXP classification in BlueXP, or on an on-premises system. Once the instance is running, you can enable the service on existing working environments, databases, and other data sources from the **Configuration** tab or by selecting a specific working environment. [Learn how to get started.](#)



Activating BlueXP classification on a data source results in an immediate initial scan. Scan results display shortly after.

You can disable BlueXP classification from scanning an individual working environment, database, or file share group from the BlueXP classification Configuration page. See [Remove data sources from BlueXP classification.](#)

To completely remove the BlueXP classification instance, you can manually remove the BlueXP classification instance from your cloud provider's portal or on-prem location.



## Can the service exclude scanning data in certain directories?

Yes. If you want BlueXP classification to exclude scanning data that resides in certain data source directories, you can provide that list to the classification engine. After you apply that change, BlueXP classification will exclude scanning data in the specified directories. [Learn more](#).

## Are snapshots that reside on ONTAP volumes scanned?

No. BlueXP classification does not scan snapshots because the content is identical to the content in the volume.

## What happens if data tiering is enabled on your ONTAP volumes?

When BlueXP classification scans volumes that have cold data tiered to object storage using the Mapping only scans, it scans all of the data—data that's on local disks and cold data tiered to object storage. This is also true for non-NetApp products that implement tiering.

The Mapping only scan doesn't heat up the cold data—it stays cold and remains in object storage. On the other hand, if you perform the Map & Classify scan, some configurations might heat up the cold data.

## Types of source systems and data types

The following questions relate to the types of storage that can be scanned, and the types of data that is scanned.

### Are there any restrictions when deployed in a Government region?

BlueXP classification is supported when the Connector is deployed in a Government region (AWS GovCloud, Azure Gov, or Azure DoD) - also known as "Restricted mode".

### What data sources can I scan if I install BlueXP classification in a site without internet access?

BlueXP classification can only scan data from data sources that are local to the on-premises site. At this time, BlueXP classification can scan the following local data sources in "Private mode" - also known as a "dark" site:

- On-premises ONTAP systems
- Database schemas
- Object Storage that uses the Simple Storage Service (S3) protocol

See [Supported working environments and data sources](#).

### Which file types are supported?

BlueXP classification scans all files for category and metadata insights, and displays all file types in the file types section of the dashboard.

When BlueXP classification detects Personal Identifiable Information (PII), or when it performs a DSAR search, only the following file formats are supported:

.CSV, .DCM, .DOC, .DOCX, .JSON, .PDF, .PPTX, .RTF, .TXT, .XLS, .XLSX, Docs, Sheets, and Slides

## What kinds of data and metadata does BlueXP classification capture?

BlueXP classification enables you to run a general "mapping" scan or a full "classification" scan on your data sources. Mapping provides only a high-level overview of your data, whereas Classification provides deep-level scanning of your data. Mapping can be done on your data sources very quickly because it does not access files to see the data inside.

- **Data mapping scan (Mapping only scan):** BlueXP classification scans the metadata only. This is useful for overall data management and governance, quick project scoping, very large estates, and prioritization. Data mapping is based on metadata and is considered a **fast** scan.

After a fast scan, you can generate a Data Mapping Report. This report is an overview of the data stored in your corporate data sources to assist you with decisions about resource utilization, migration, backup, security, and compliance processes.

- **Data classification deep scan (Map & Classify scan):** BlueXP classification scans using standard protocols and read-only permission throughout your environments. Select files are opened and scanned for sensitive business-related data, private information, and issues related to ransomware.

After a full scan there are many additional BlueXP classification features you can apply to your data, such as view and refine data in the Data Investigation page, search for names within files, copy, move, and delete source files, and more.

BlueXP classification captures metadata such as: file name, permissions, creation time, last access, and last modification. This includes all of the metadata that appears in the Data Investigation Details page and in Data Investigation Reports.

BlueXP classification can identify many types of private data such as personal information (PII) and sensitive personal information (SPII). For details about private data, refer to [Categories of private data that BlueXP classification scans](#).

## Can I limit BlueXP classification information to specific users?

Yes, BlueXP classification is fully integrated with BlueXP. BlueXP users can only see information for the working environments they are eligible to view according to their permissions.

Additionally, if you want to allow certain users to just view BlueXP classification scan results without having the ability to manage BlueXP classification settings, you can assign those users the **Classification viewer** role (when using BlueXP in standard mode) or the **Compliance Viewer** role (when using BlueXP in restricted mode). [Learn more](#).

## Can anyone access the private data sent between my browser and BlueXP classification?

No. The private data sent between your browser and the BlueXP classification instance are secured with end-to-end encryption using TLS 1.2, which means NetApp and non-NetApp parties can't read it. BlueXP classification won't share any data or results with NetApp unless you request and approve access.

The data that is scanned stays within your environment.

## How is sensitive data handled?

NetApp does not have access to sensitive data and does not display it in the UI. Sensitive data is masked, for example, the last four numbers are displayed for credit card information.

## Where is the data stored?

Scan results are stored in Elasticsearch within your BlueXP classification instance.

## How is the data accessed?

BlueXP classification accesses data stored in Elasticsearch through API calls, which require authentication and are encrypted using AES-128. Accessing Elasticsearch directly requires root access.

## Licenses and costs

The following question relates to licensing and costs to use BlueXP classification.

### How much does BlueXP classification cost?

BlueXP classification is a BlueXP core capability and is not charged.

## Connector deployment

The following questions relate to the BlueXP Connector.

### What is the Connector?

The Connector is software running on a compute instance either within your cloud account, or on-premises, that enables BlueXP to securely manage cloud resources. You must deploy a Connector to use BlueXP classification.

### Where does the Connector need to be installed?

When scanning data, the BlueXP Connector needs to be installed in the following locations:

- For Cloud Volumes ONTAP in AWS or Amazon FSx for ONTAP: Connector is in AWS.
- For Cloud Volumes ONTAP in Azure or in Azure NetApp Files: Connector is in Azure.
- For Cloud Volumes ONTAP in GCP: Connector is in GCP.
- For on-premises ONTAP systems: Connector is on-premises.

If you have data in these locations, you may need to use [multiple Connectors](#).

### Does BlueXP classification require access to credentials?

BlueXP classification itself doesn't retrieve storage credentials. Instead, they are stored within the BlueXP Connector.

BlueXP classification uses data plane credentials, for example, CIFS credentials to mount shares before scanning.

### Does communication between the service and the Connector use HTTP?

Yes, BlueXP classification communicates with the BlueXP Connector using HTTP.

# BlueXP classification deployment

The following questions relate to the separate BlueXP classification instance.

## What deployment models does BlueXP classification support?

BlueXP allows the user to scan and report on systems virtually anywhere, including on-premises, cloud, and hybrid environments. BlueXP classification is normally deployed using a SaaS model, in which the service is enabled via the BlueXP interface and requires no hardware or software installation. Even in this click-and-run deployment mode, data management can be done regardless of whether the data stores are on premises or in the public cloud.

## What type of instance or VM is required for BlueXP classification?

When [deployed in the cloud](#):

- In AWS, BlueXP classification runs on an m6i.4xlarge instance with a 500 GiB GP2 disk. You can select a smaller instance type during deployment.
- In Azure, BlueXP classification runs on a Standard\_D16s\_v3 VM with a 500 GiB disk.
- In GCP, BlueXP classification runs on an n2-standard-16 VM with a 500 GiB Standard persistent disk.

[Learn more about how BlueXP classification works.](#)

## Can I deploy the BlueXP classification on my own host?

Yes. You can install BlueXP classification software on a Linux host that has internet access in your network or in the cloud. Everything works the same and you continue to manage your scan configuration and results through BlueXP. See [Deploying BlueXP classification on premises](#) for system requirements and installation details.

## What about secure sites without internet access?

Yes, that's also supported. You can [deploy BlueXP classification in an on-premises site that doesn't have internet access](#) for completely secure sites.

# 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 BlueXP classification](#)

## Copyright information

Copyright © 2025 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.



# **BlueXP digital wallet documentation**

## **BlueXP digital wallet**

NetApp  
July 16, 2025

# Table of Contents

- BlueXP digital wallet documentation . . . . . 1
- Release notes . . . . . 2
  - What's new . . . . . 2
    - 10 March 2025 . . . . . 2
    - 10 February 2025 . . . . . 2
    - 5 March 2024 . . . . . 2
    - 30 July 2023 . . . . . 3
    - 7 May 2023 . . . . . 3
    - 3 April 2023 . . . . . 3
    - 6 November 2022 . . . . . 4
    - 18 September 2022 . . . . . 4
    - 31 July 2022 . . . . . 4
    - 3 July 2022 . . . . . 4
    - 27 February 2022 . . . . . 5
    - 2 January 2022 . . . . . 5
- Get started . . . . . 6
  - Learn about the BlueXP digital wallet . . . . . 6
    - How licenses and subscriptions are displayed in digital wallet . . . . . 6
    - Supported services . . . . . 6
- Use the BlueXP digital wallet . . . . . 7
  - Use the dashboard overview . . . . . 7
  - Manage licenses . . . . . 9
    - Manage licenses for BlueXP data services . . . . . 9
    - Manage licenses for Cloud Volumes ONTAP . . . . . 12
  - Manage PAYGO subscriptions and contracts . . . . . 21
    - View your subscriptions . . . . . 22
    - Rename a subscription . . . . . 23
    - Configure a subscription with a provider credential . . . . . 23
    - Associate a subscription with a BlueXP organization . . . . . 24
    - View credentials associated with a subscription . . . . . 24
    - Add a new marketplace subscription . . . . . 24
  - Manage Keystone subscriptions . . . . . 26
    - Authorize your account . . . . . 26
    - Link a subscription . . . . . 27
    - Request more or less committed capacity . . . . . 27
    - Monitor usage . . . . . 28
    - Unlink a subscription . . . . . 28
  - Manage licenses for on-premises ONTAP . . . . . 29
    - View cluster information and contract details . . . . . 29
    - Discover clusters . . . . . 30
- Knowledge and support . . . . . 31
  - Register for support . . . . . 31
  - Get help . . . . . 35



Legal notices .....	40
Copyright .....	40
Trademarks .....	40
Patents .....	40
Privacy policy .....	41
Open source .....	41

# BlueXP digital wallet documentation

# Release notes

## What's new

Learn what's new with the BlueXP digital wallet.

### 10 March 2025

#### Ability to remove subscriptions

You can now remove subscriptions from the digital wallet if you have unsubscribed from them.

#### View consumed capacity for Marketplace subscriptions

When viewing PAYGO subscriptions, you can now view the consumed capacity of the subscription.

### 10 February 2025

The BlueXP digital wallet has been redesigned for ease of use and now provides additional subscription and license management.

#### New Overview dashboard

The digital wallet homepage has an updated dashboard of your NetApp licenses and Marketplace subscriptions, with the ability to drill-down into specific services, license types and required actions.

#### Configuring subscriptions to credentials

The BlueXP digital wallet now allows you to configure your subscriptions to provider credentials. Typically you do this when you first subscribe to a Marketplace subscription or annual contract. Previously changing the subscription's credentials could only be done on the Credentials page.

#### Associating subscriptions with organizations

You can now update the organization to which a subscription is associated directly from digital wallet.

#### Managing Cloud Volume ONTAP licenses

You now manage Cloud Volumes ONTAP licenses through the home page or the **Direct licenses** tab. Use the **Marketplace subscriptions** tab to view your subscription information.

### 5 March 2024

#### BlueXP disaster recovery

The BlueXP digital wallet now enables you to manage licenses for BlueXP disaster recovery. You can add licenses, update licenses, and view details about licensed capacity.

[Learn how to manage licenses for BlueXP data services](#)

## 30 July 2023

### Usage reports enhancements

Several improvements to the Cloud Volumes ONTAP usage reports are now available:

- The TiB unit is now included in the name of columns.
- A new *node(s)* field for serial numbers is now included.
- A new *Workload Type* column is now included under the Storage VMs usage report.
- Working environment names are now included in the Storage VMs and Volume usage reports.
- The volume type *file* is now labeled *Primary (Read/Write)*.
- The volume type *secondary* is now labeled *Secondary (DP)*.

For more information about the usage reports, refer to [Download usage reports](#).

## 7 May 2023

### Google Cloud private offers

The BlueXP digital wallet now identifies Google Cloud Marketplace subscriptions that are associated with a private offer and shows the end date and term of the subscription. This enhancement enables you to verify that you've successfully accepted the private offer and to validate its terms.

### Charging usage breakdown

Now you can find out what you're being charged for when you're subscribed to capacity-based licenses. The following types of usage reports are available for download from the BlueXP digital wallet. The usage reports provide capacity details of your subscriptions and tell you how you're being charged for the resources in your Cloud Volumes ONTAP subscriptions. The downloadable reports can be easily shared with others.

- Cloud Volumes ONTAP package usage
- High-level usage
- Storage VMs usage
- Volumes usage

For more information about the usage reports, refer to [Download usage reports](#).

## 3 April 2023

### Email notifications

Email notifications are now supported with the BlueXP digital wallet.

If you configure your notification settings, you can receive email notifications when your BYOL licenses are about to expire (a "Warning" notification) or if they have already expired (an "Error" notification).

[Learn how to set up email notifications](#)

## **Licensed capacity for marketplace subscriptions**

When viewing capacity-based licensing for Cloud Volumes ONTAP, the BlueXP digital wallet now shows the licensed capacity that you purchased with marketplace private offers.

[Learn how to view the consumed capacity in your account.](#)

## **6 November 2022**

### **Subscriptions and annual contracts**

Your PAYGO subscriptions and annual contracts for BlueXP are now available to view and manage from the BlueXP digital wallet.

[Learn how to manage your subscriptions.](#)

## **18 September 2022**

### **Optimized I/O and WORM capacity**

The BlueXP digital wallet now shows a summary of the Optimized I/O licensing package and the provisioned WORM capacity for Cloud Volumes ONTAP systems across your account.

These details can help you better understand how you're being charged and whether you need to purchase additional capacity.

[Learn how to view the consumed capacity in your account.](#)

## **31 July 2022**

### **Change charging method**

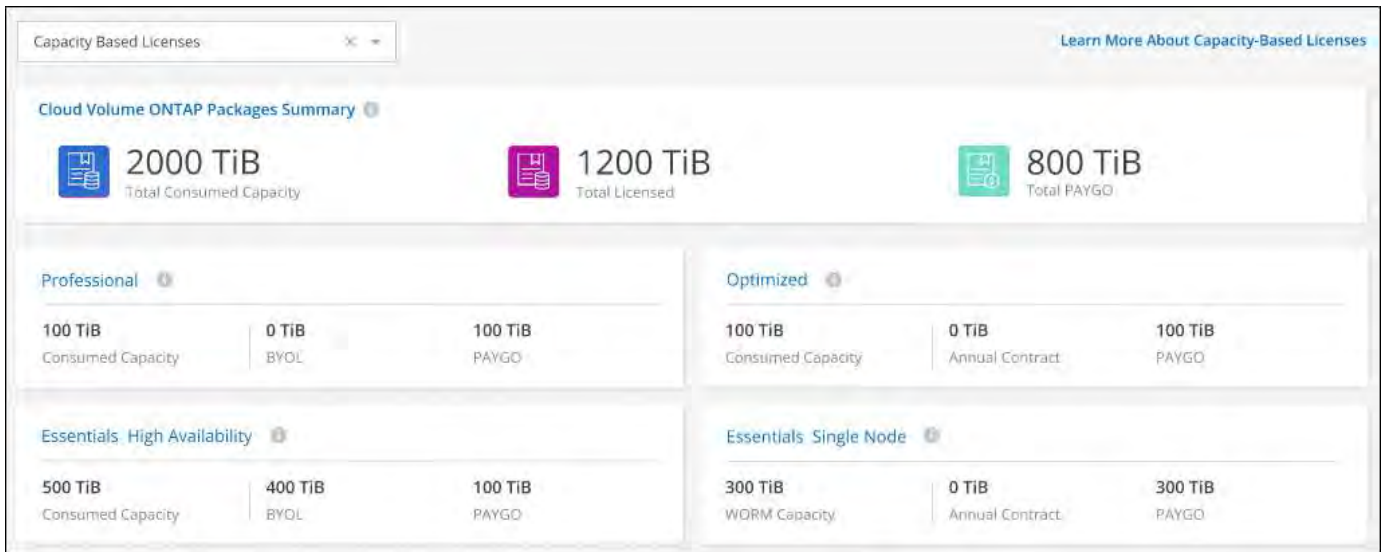
You can now change the charging method for a Cloud Volumes ONTAP system that uses capacity-based licensing. For example, if you deployed a Cloud Volumes ONTAP system with the Essentials package, you can change it to the Professional package if your business needs changed.

[Learn how to change charging methods.](#)

## **3 July 2022**

### **Consumed capacity**

The now shows you the total consumed capacity in your account and the consumed capacity by licensing package. This can help you understand how you're being charged and whether you need to purchase additional capacity.



## 27 February 2022

### Licenses for on-premises ONTAP clusters

You can now view an inventory of your on-prem ONTAP clusters along with their hardware and service contracts expiration dates. Additional details about the clusters are also available.

[Learn how to manage licenses for on-prem ONTAP clusters.](#)

## 2 January 2022

### Licensing terms update automatically

If you change the capacity or term for any of your licenses, the license terms now automatically update in the . You don't need to manually update the license yourself.

The automatic license update works with all types of Cloud Volumes ONTAP licenses and all licenses for data services.

# Get started

## Learn about the BlueXP digital wallet

The BlueXP digital wallet enables you to manage and monitor BlueXP licenses purchased from NetApp (BYOL), BlueXP data services marketplace subscriptions (including NetApp Cloud Volumes ONTAP), and NetApp Keystone.

### How licenses and subscriptions are displayed in digital wallet

Both licenses and subscriptions automatically display in digital wallet when the BlueXP account used to subscribe is also an NetApp Support Site account. If you used a BlueXP account that is not associated with your NetApp Support Site (NSS) account, you'll need to manually add and update licenses and you won't see usage, capacity, and other details change.



You need to have deployed a Connector in order to view subscription information in digital wallet. A Connector is also needed to view node licenses for Cloud Volumes ONTAP.

In the meantime, BlueXP automatically queries NetApp's licensing service to obtain details about the licenses associated with your NetApp Support Site account. If there are no errors, BlueXP automatically adds the licenses to the digital wallet.

If BlueXP can't add the license, you'll need to manually add it to the digital wallet yourself. For example, if your BlueXP user account is not associated with your NetApp Support Site account, you'll need to add the licenses yourself.

After you purchase a license from your NetApp sales representative, NetApp sends you an email with the serial number and additional licensing details. You'll need that serial number to add or update the respective license in digital wallet.

[Learn how to add your NetApp Support Site account to BlueXP user credentials.](#)

### Supported services

The BlueXP digital wallet enables you to manage licenses and subscriptions for the following services:

- [Backup and recovery](#)
- [Cloud Volumes ONTAP](#)
- [Disaster recovery](#) (BYOL only)
- [Ransomware protection](#)
- [On-prem ONTAP clusters](#)
- [Tiering](#)

# Use the BlueXP digital wallet

## Use the dashboard overview

Use the **Overview** dashboard in digital wallet to monitor the health of your licenses and subscriptions to ensure that you can manage costs and maintain service as subscriptions reach expiration dates or capacity limits.

You can view specific license and subscription information about each of your data services (including Cloud Volumes ONTAP) and drill-down into details for each.

You'll also be alerted to upcoming capacity limits or expirations and be prompted to take action.

## View details about a specific data service

The digital wallet **Overview** dashboard provides the ability to view details on the following data services:

- Cloud Volumes ONTAP
- Disaster recovery
- Ransomware protection
- Backup and recovery
- Tiering

For example, administrators tasked with managing Cloud Volumes ONTAP resources can view the current capacity of each license or subscription which gives them a focused view resources within their purview.

Note: Although you can update and remove licenses from the dashboard, you can't add a new license or subscription.

[Learn more about managing licenses.](#)

[Learn more about managing subscriptions.](#)

## Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Overview**.
3. Select **View** for the respective data service (including Cloud Volumes ONTAP) to view details and manage licenses and subscriptions.

Each tile provides a graphical overview of capacity. You can view the details page to see and manage payment methods for that capacity.

4. Select either the **Licenses** or **Subscriptions** tab.
5. Filter the table by selecting a column filter and a value to filter by. For example, on the **Subscriptions** tab, you can filter the Type column by Annual Contract or Subscription.



The screenshot shows a table with the following columns: Provider, Name, Provider credentials, Type, Start date, End date, Status, and Configuration. The first row is selected, and a dropdown menu is open over the 'Configuration' column, showing 'Configured' (checked) and 'Unconfigured' (unchecked) options, with 'Apply' and 'Clear' buttons at the bottom.

Provider	Name	Provider credentials	Type	Start date	End date	Status	Configuration
aws	Annual_small_1TB_all_servi...	1 View	Annual Contract	Mar 31, 2024	Mar 31, 2025	Subscribed	Configured
gordan-demo-annual-azure		N/A	Annual Contract	Nov 14, 2024	Nov 12, 2025	Subscribed	Unconfigured
Gordon_BXP_DR_RPS_Ann...		N/A	Annual Contract	Dec 13, 2024	Dec 11, 2025	Subscribed	Unconfigure

6. Customize the columns shown in the table. You can add and remove columns from the table by selecting the **Column** icon indicating which columns you want to show or hide. The following screenshot indicates that only the Type, End Date, Status, and Configuration columns.

The screenshot shows the same table as above, but with a column selection dropdown menu open. The menu lists columns with checkboxes: Provider credentials (unchecked), Type (checked), Start date (unchecked), End date (checked), Status (checked), Configuration (checked), and Service (unchecked). 'Apply' and 'Restore Defaults' buttons are at the bottom.

Provider	Name	Provider credentials	Type	Start date	End date	Status	Configuration
aws	Annual_small_1TB_all_servi...	1 View	Annual Contract	Mar 31, 2024	Mar 31, 2025	Subscribed	Configured
gordan-demo-annual-azure		N/A	Annual Contract	Nov 14, 2024	Nov 12, 2025	Subscribed	Unconfigured
Gordon_BXP_DR_RPS_Ann...		N/A	Annual Contract	Dec 13, 2024	Dec 11, 2025	Subscribed	Unconfigure

- Expand any row to view details of what the subscription or license includes, capacity purchased, and terms.
- Use the actions menu to manage a specific license or subscription, such as update a license or associate a subscription with a different organization.

You cannot add new licenses or subscriptions from the **Overview** page for a specific service, you can only add licenses or subscriptions from the **Direct licenses** tab or the **Marketplace subscriptions** tab, respectively.

[Learn more about managing licenses.](#)  
[Learn more about managing subscriptions.](#)

## Resolve license or subscription issues

You can view license and subscription issues that need to be resolved. Issues include licenses and subscriptions that are expiring or reaching capacity.

- From the BlueXP navigation menu, select **Governance > Digital wallet**.
- Select **Overview**.
- Select **Resolve** for the **Requires action** tile to view issues that need to be adjusted. If **Resolve** does not display there are no issues that require action at this time.
- From the **Require action** page, select either the **Licenses** tab or the **Subscriptions** tab.

5. Use the action menu to resolve the issue.

[Learn more about managing licenses.](#)

[Learn more about managing subscriptions.](#)

## Manage licenses

### Manage licenses for BlueXP data services

The BlueXP digital wallet enables you to manage licenses that you purchased directly from NetApp (BYOL) for use with BlueXP data services, including Cloud Volumes for ONTAP. You can view used license capacity, how much free capacity you have left, and you'll see notifications if you reached the capacity limit or the expiration date.



The **Direct licenses** page lists all licenses. If you want license details for a specific data service, use the data service tiles on the **Overview** dashboard. [Learn more about the Overview dashboard.](#)

The instructions on this page provide information that applies to each service. For more specific information about the licensing for these services, refer to the following pages:

- [Set up licensing for BlueXP backup and recovery](#)
- [Set up licensing for BlueXP disaster recovery](#)
- [Set up licensing for BlueXP ransomware protection](#)
- [Set up licensing for BlueXP tiering](#)
- [Set up licensing for Cloud Volumes ONTAP](#)


### Obtain a license file

You should obtain a NetApp license file to upload if BlueXP does not have internet access (private mode installations).

After you purchase a license from your NetApp sales representative, NetApp sends you an email with the serial number and additional licensing details. In the case where you do not see your licenses automatically, you'll need that serial number to obtain the respective license file.

### Steps

1. Find your BlueXP account ID:

- a. In the upper right of the BlueXP console, select  > **Identity & Access Management**.
- b. On the Organization page, look for your account ID and copy it.

If there is no account ID listed and you just have an organization ID, then you'll need to copy the first eight characters of the organization ID and append it to *account-*

For example, let's say this is your organization ID:

ea10e1c6-80cc-4219-8e99-3c3e6b161ba5

Your account ID would be as follows:

account-ea10e1c6

2. Sign in to the [NetApp Support Site](#) and select **Systems > Software Licenses**.
3. Enter the serial number for your license.

Software Licenses

Serial Number: 4810

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

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

**Get License**

SERIAL NUMBER: 4810

LICENSE: CLOUD\_BKP\_SERVICE

SALES ORDER: 3005

TENANT ID: Enter Tenant ID  
Example: account-xxxxxxx

Cancel Submit

## Add a license

License information automatically displays in digital wallet when the BlueXP account associated with the license is also a NetApp Support Site account and BlueXP has access to the internet. If you used a BlueXP account that is not associated with your NSS account, you'll need to manually add a license.

You can add the license in BlueXP either by entering the serial number and the associated NSS account, or by uploading the NetApp license file (NLF). You should obtain a NetApp license file to upload if BlueXP does not have internet access (private mode installations).

After you purchase a license from your NetApp sales representative, NetApp sends you an email with the serial number and additional licensing details. You'll need that serial number to add or update the respective license in digital wallet.



If you want to enter the serial number, you first need to [add your NetApp Support Site account to BlueXP](#). This is the NetApp Support Site account that's authorized to access the serial number.

## Steps

1. From the BlueXP menu, select **Governance > Digital wallet** and then select the **Direct licenses** tab.
2. Select **Add license**.
3. In the *Add license* dialog, enter the license information and select **Add license**:

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

If you entered a serial number, you also need to select the NetApp Support Site account that's authorized to access the serial number.

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

- If you have the license file (required when using BlueXP in private mode), select the **Upload license file** option and follow the prompts to attach the file.

## Update a license

If your licensed term is nearing expiration, or if your licensed capacity is reaching the limit, you'll be notified in the BlueXP digital wallet. You can update your license before it expires so that there is no interruption in your ability to use a BlueXP data service or Cloud Volumes ONTAP.

After you purchase additional capacity from NetApp or extend the term of your license, BlueXP automatically updates the license in the digital wallet when the BlueXP account associated with the license is also a NetApp Support Site account and BlueXP has access to the internet.

If your BlueXP account is not associated with your NSS account, you'll need to manually update a license.

You can manually update your license in BlueXP either by entering the serial number and the associated NSS account, or by uploading the NetApp license file (or *files* if you have a Cloud Volumes ONTAP HA pair). You should obtain a NetApp license file to upload if BlueXP does not have internet access (private mode installations).

After you purchase a license from your NetApp sales representative, NetApp sends you an email with the serial number and additional licensing details. You'll need that serial number to add or update the respective license in digital wallet.



If you want to enter the serial number, you first need to [add your NetApp Support Site account to BlueXP](#). This is the NetApp Support Site account that's authorized to access the serial number.

## Steps

1. Contact your NetApp representative to buy a new license.

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 **Direct licenses** page will reflect the change in 5 to 10 minutes.

2. If BlueXP can't automatically update the license (for example, when using BlueXP in private mode), then you'll need to obtain a NetApp license file from support and manually upload the license file. [Learn how to obtain a license file](#).
3. On the **Direct licenses** tab, select **...** for the serial number you are updating, and select **Update license**.

4. In the **Update license** page, upload the license file and select **Update license**.

## View license status

To manage licenses, you can group licenses based on the service name. This allows you to see all licenses related to a specific service. You can expand a row to view detailed information about each license related to the service. The root row for each service displays the service name and the used capacity for that service. The licenses are automatically grouped by service name. The root row for each service shows the service name and the used capacity for that service.

### Steps

1. From the BlueXP menu, select **Governance > Digital wallet**, and then select the **Direct licenses** tab.
2. Click a service name row to expand it. This displays all licenses related to that service. Each expanded row displays detailed information about the licenses, including license ID, serial number, capacity, and expiration date.

## Manage licenses for Cloud Volumes ONTAP

### Manage capacity-based Cloud Volume ONTAP licenses

Manage your capacity-based licenses from the BlueXP digital wallet to ensure that your NetApp account has enough capacity for your Cloud Volumes ONTAP systems.

*Capacity-based licenses* enable you to pay for Cloud Volumes ONTAP per TiB of capacity.

The *BlueXP digital wallet* enables you to manage licenses for Cloud Volumes ONTAP from a single location. You can add new licenses and update existing licenses.



While the actual usage and metering for the products and services managed in BlueXP are always calculated in GiB and TiB, the terms GB/GiB and TB/TiB are used interchangeably. This is reflected in the Cloud Marketplace listings, price quotes, listing descriptions, and in other supporting documentation

[Learn more about Cloud Volumes ONTAP licenses.](#)

### How licenses are added to the BlueXP digital wallet

After you purchase a license from your NetApp sales representative, NetApp will send you an email with the serial number and additional licensing details.

In the meantime, BlueXP automatically queries NetApp's licensing service to obtain details about the licenses associated with your NetApp Support Site account. If there are no errors, BlueXP automatically adds the licenses to the digital wallet.

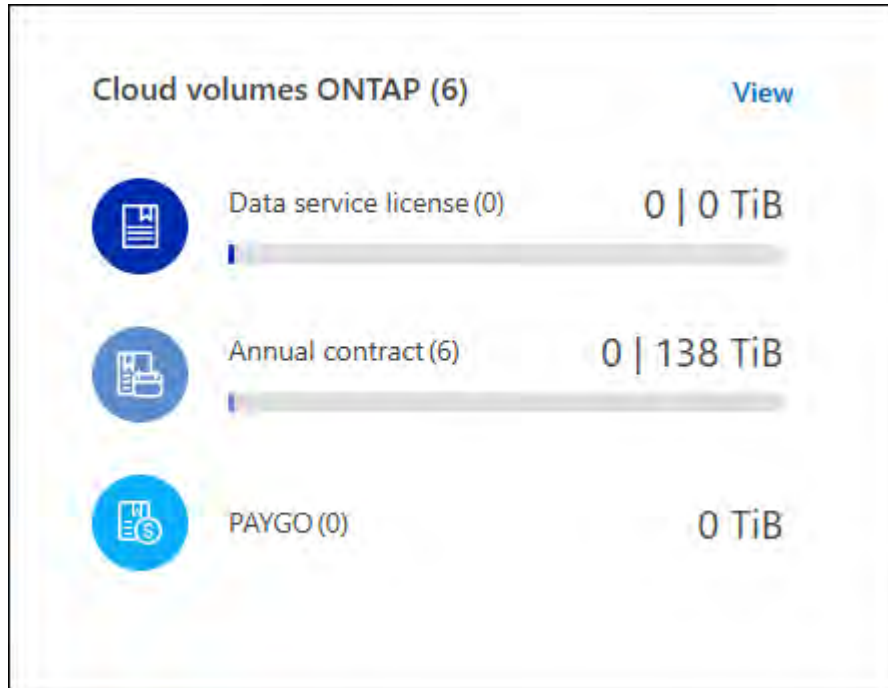
If BlueXP can't add the license, you'll need to manually add them to the digital wallet yourself. For example, if the Connector is installed in a location that doesn't have internet access, you'll need to add the licenses yourself. [Learn how to add purchased licenses to your account.](#)

### View the consumed capacity in your account

The BlueXP digital wallet shows you the total consumed capacity in your account and the consumed capacity by licensing package. This can help you understand how you're being charged and whether you need to purchase additional capacity.

## Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. On the **Overview** tab, the Cloud Volumes ONTAP tile displays the current capacity provisioned for your account.



- *Direct license* is the total provisioned capacity of all Cloud Volumes ONTAP systems in your NetApp account. The charging is based on each volume's provisioned size, regardless of local, used, stored, or effective space within the volume.
  - *Annual contract* is the total licensed capacity (bring your own license (BYOL) or Marketplace Contract) that you purchased from NetApp.
  - *PAYGO* is the total provisioned capacity using cloud marketplace subscriptions. Charging via PAYGO is used only if the consumed capacity is higher than the licensed capacity or if there is no BYOL license available in the BlueXP digital wallet.
3. Select **View** to see the consumed capacity for each of your licensing packages.
  4. Select the **Licenses** tab to see details for each package license that you have purchased.

To better understand the capacities that display for the Essentials package, you should be familiar with how charging works. [Learn about charging for the Essentials package.](#)

5. Select the **Subscriptions** tab to see the consumed capacity by license consumption model. This tab includes both PAYGO and annual contract licenses.

You'll only see the subscriptions that are associated with the organization that you are that you're currently viewing.

6. As you view the information about your subscriptions, you can interact with the details in the table as follows:
  - Expand a row to view more details.



Provider	Name	Type	Service	Start Date	Status	
aws	AWS subscription	PAYGO	NetApp BlueXP	Apr 04, 2024	Subscribed	⋮
NetApp BlueXP			N/A	N/A		
Product Title			Term	Auto Renew		

- Select to choose which columns appear in the table. Note that the Term and Auto Renew columns don't appear by default. The Auto Renew column displays renewal information for Azure contracts only.

### Viewing package details

You can view details about the capacity used per package by switching to legacy mode on the Cloud Volumes ONTAP page.

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. On the **Overview** tab, the Cloud Volumes ONTAP tile displays the current capacity provisioned for your account.
3. Select **View** to see the provisioned capacity for each of your licensing packages.
4. Select **Switch to advanced view**.

Digital Wallet Overview

Cloud Volumes ONTAP

View: Capacity Based Licenses

Usage report | **Switch to legacy view**

Data service... (1) ⚠ 0   10 TiB	Annual contract(5) 6   132 ...	PAYGO (1) 9 TiB
----------------------------------	--------------------------------	-----------------

5. View the details of the package you want to see.

[A screenshot of switch to standard view button]

### Change charging methods

Capacity-based licensing is available in the form of a *package*. When you create a Cloud Volumes ONTAP working environment, you can choose from several licensing packages based on your business needs. If your needs change after you create the working environment, you can change the package at any time. For example, you might change from the Essentials package to the Professional package.

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

## About this task

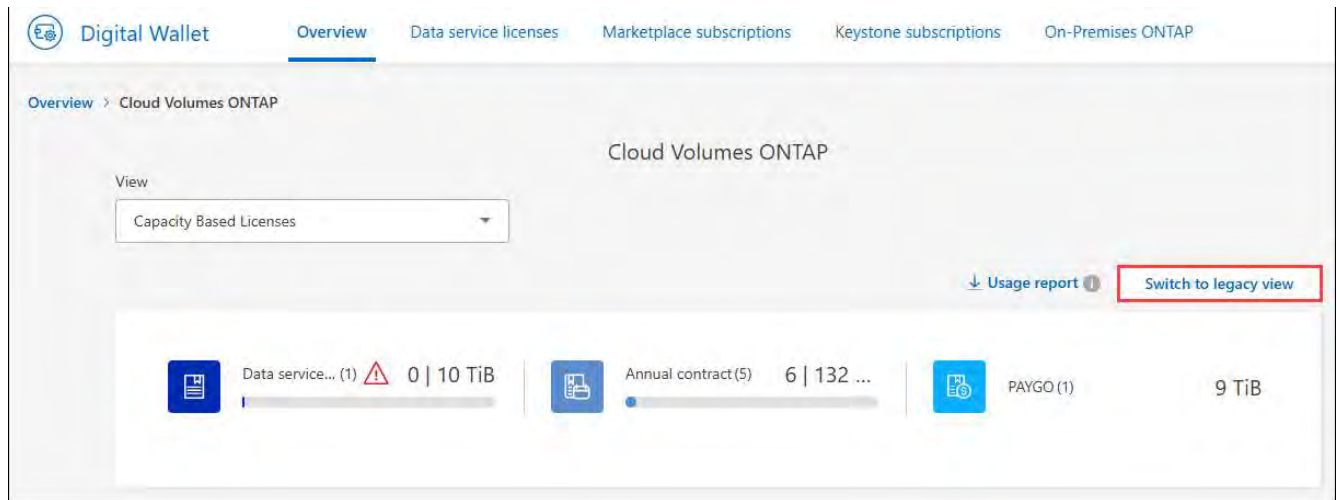
- Changing the charging method doesn't affect whether you're charged through a license purchased from NetApp (BYOL) or from your cloud provider's marketplace pay-as-you-go (PAYGO) subscription.

BlueXP always attempts to charge against a license first. If a license isn't available, it charges against a marketplace subscription. No "conversion" is required for BYOL to marketplace subscription or vice versa.

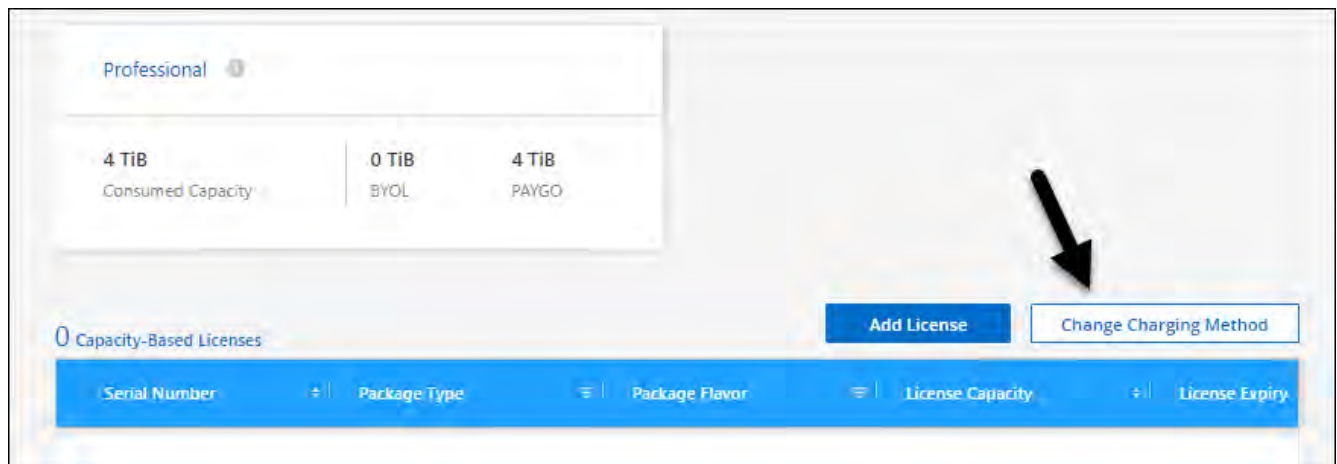
- If you have a private offer or contract from your cloud provider's marketplace, changing to a charging method that's not included in your contract will result in charging against BYOL (if you purchased a license from NetApp) or PAYGO.

## Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Switch to advanced view**.



5. Scroll down to the **Capacity-based license** table and select **Change charging method**.



6. On the **Change charging method** pop-up, select a working environment, choose the new charging method, and then confirm your understanding that changing the package type will affect service charges.
7. Select **Change charging method**.



## Download usage reports

You can download four usage reports from the BlueXP digital wallet. These usage reports provide capacity details of your subscriptions and tell you how you're being charged for the resources in your Cloud Volumes ONTAP subscriptions. The downloadable reports capture data at a point in time and can be easily shared with others.



The following reports are available for download. Capacity values shown are in TiB.

- **High-level usage:** This report includes the following information:
  - Total consumed capacity
  - Total precommitted capacity
  - Total BYOL capacity
  - Total Marketplace contracts capacity
  - Total PAYGO capacity
- **Cloud Volumes ONTAP package usage:** This report includes the following information for each package except the Optimized I/O package:
  - Total consumed capacity
  - Total precommitted capacity
  - Total BYOL capacity
  - Total Marketplace contracts capacity
  - Total PAYGO capacity
- **Storage VMs usage:** This report shows how charged capacity is broken down across Cloud Volumes ONTAP systems and storage virtual machines (SVMs). This information is only available in the report. It contains the following information:
  - Working environment ID and name (appears as the UUID)
  - Cloud
  - NetApp account ID
  - Working environment configuration
  - SVM name
  - Provisioned capacity
  - Charged capacity roundup
  - Marketplace billing term
  - Cloud Volumes ONTAP package or feature
  - Charging SaaS Marketplace subscription name
  - Charging SaaS Marketplace subscription ID

- Workload type
- **Volumes usage:** This report shows how charged capacity is broken down by volumes in a working environment. This information is not available on any screen in the digital wallet. It includes the following information:
  - Working environment ID and name (appears as the UUID)
  - SVN name
  - Volume ID
  - Volume type
  - Volume provisioned capacity



FlexClone volumes aren't included in this report because these types of volumes don't incur charges.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. On the **Overview** tab, select **View** from the Cloud Volumes ONTAP tile.
3. Select **Usage report**.

The usage report downloads.

4. Open the downloaded file to access the reports.

### Manage node-based licenses

Manage node-based licenses in the BlueXP digital wallet to ensure that each Cloud Volumes ONTAP system has a valid license with the required capacity.

*Node-based licenses* are the previous generation licensing model (and not available for new customers):

- Bring your own license (BYOL) licenses purchased from NetApp
- Hourly pay-as-you-go (PAYGO) subscriptions from your cloud provider's marketplace

The *BlueXP digital wallet* enables you to manage licenses for Cloud Volumes ONTAP from a single location. You can add new licenses and update existing licenses.

[Learn more about Cloud Volumes ONTAP licenses.](#)

### Manage PAYGO licenses

The BlueXP digital wallet page enables you to view details about each of your PAYGO Cloud Volumes ONTAP systems, including the serial number and PAYGO license type.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.

5. Click **PAYGO**.
6. View details in the table about each of your PAYGO licenses.

[A screenshot that shows a table in the BlueXP digital wallet page with three paygo licenses. Each row shows the name, type of system, serial number, package, and a link to manage the license.]

7. If needed, click **Manage PAYGO License** to change the PAYGO license or to change the instance type.

### Manage BYOL licenses

Manage licenses that you purchased directly from NetApp by adding and removing system licenses and extra capacity licenses.

### Add unassigned licenses

Add a node-based license to the BlueXP digital wallet so that you can select the license when you create a new Cloud Volumes ONTAP system. The digital wallet identifies these licenses as *unassigned*.

#### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. Click **Unassigned**.
6. Click **Add Unassigned Licenses**.
7. Enter the serial number of the license or upload the license file.

If you don't have the license file yet, refer to the section below.

8. Click **Add License**.

#### Result

BlueXP adds the license to the digital wallet. The license will be identified as unassigned until you associate it with a new Cloud Volumes ONTAP system. After that happens, the license moves to the **BYOL** tab in the digital wallet.

### Exchange unassigned node-based licenses

If you have an unassigned node-based license for Cloud Volumes ONTAP that you haven't used, you can exchange the license by converting it to a BlueXP backup and recovery license, a BlueXP classification license, or a BlueXP tiering license.

Exchanging the license revokes the Cloud Volumes ONTAP license and creates a dollar-equivalent license for the service:

- Licensing for a Cloud Volumes ONTAP HA pair is converted to a 51 TiB direct license
- Licensing for a Cloud Volumes ONTAP single node is converted to a 32 TiB direct license

The converted license has the same expiration date as the Cloud Volumes ONTAP license.

[View walkthrough of how to exchange node-based licenses.](#)

## Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. Click **Unassigned**.
6. Click **Exchange License**.

[A screenshot of the Exchange License option that appears on the Unassigned license page.]

7. Select the service that you'd like to exchange the license with.
8. If you're prompted, select an additional license for the HA pair.
9. Read the legal consent and click **Agree**.

## Result

BlueXP converts the unassigned license to the service that you selected. You can view the new license in the **Data Services Licenses** tab.

## Obtain a system license file

In most cases, BlueXP can automatically obtain your license file using your NetApp Support Site account. But if it can't, then you'll need to manually upload the license file. If you don't have the license file, you can obtain it from netapp.com.

## Steps

1. Go to the [NetApp License File Generator](#) and log in using your NetApp Support Site credentials.
2. Enter your password, choose your product, enter the serial number, confirm that you have read and accepted the privacy policy, and then click **Submit**.

## Example

[Screen shot: Shows an example of the NetApp License Generator web page with the available product lines.]

3. Choose whether you want to receive the serialnumber.NLF JSON file through email or direct download.

## Update a system license

When you renew a BYOL subscription by contacting a NetApp representative, BlueXP automatically obtains the new license from NetApp and installs it on the Cloud Volumes ONTAP system.

If BlueXP can't access the license file over the secure internet connection, you can obtain the file yourself and then manually upload the file to BlueXP.

## Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.

5. In the **BYOL** tab, expand the details for a Cloud Volumes ONTAP system.
6. Click the action menu next to the system license and select **Update License**.
7. Upload the license file (or files if you have an HA pair).
8. Click **Update License**.

### Result

BlueXP updates the license on the Cloud Volumes ONTAP system.

### Manage extra capacity licenses

You can purchase extra capacity licenses for a Cloud Volumes ONTAP BYOL system to allocate more than the 368 TiB of capacity that's provided with a BYOL system license. For example, you might purchase one extra license capacity to allocate up to 736 TiB of capacity to Cloud Volumes ONTAP. Or you could purchase three extra capacity licenses to get up to 1.4 PiB.

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

### Add capacity licenses

Purchase an extra capacity license by contacting us through the chat icon in the lower-right of BlueXP. After you purchase the license, you can apply it to a Cloud Volumes ONTAP system.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. In the **BYOL** tab, expand the details for a Cloud Volumes ONTAP system.
6. Click **Add Capacity License**.
7. Enter the serial number or upload the license file (or files if you have an HA pair).
8. Click **Add Capacity License**.

### Update capacity licenses

If you extended the term of an extra capacity license, you'll need to update the license in BlueXP.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. In the **BYOL** tab, expand the details for a Cloud Volumes ONTAP system.
6. Click the action menu next to the capacity license and select **Update License**.
7. Upload the license file (or files if you have an HA pair).
8. Click **Update License**.

## Remove capacity licenses

If an extra capacity license expired and is no longer in use, then you can remove it at any time.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select the **Overview** tab.
3. On the Cloud Volumes ONTAP tile, select **View**.
4. Select **Node Based Licenses** from the drop-down.
5. In the **BYOL** tab, expand the details for a Cloud Volumes ONTAP system.
6. Click the action menu next to the capacity license and select **Remove License**.
7. Click **Remove**.

### Change between PAYGO and BYOL

Converting a system from PAYGO by-node licensing to BYOL by-node licensing (and vice versa) isn't supported. If you want to switch between a pay-as-you-go subscription and a BYOL subscription, then you need to deploy a new system and replicate data from the existing system to the new system.

### Steps

1. Create a new Cloud Volumes ONTAP working environment.
2. Set up a one-time data replication between the systems for each volume that you need to replicate.

[Learn how to replicate data between systems](#)

3. Terminate the Cloud Volumes ONTAP system that you no longer need by deleting the original working environment.

[Learn how to delete a Cloud Volumes ONTAP working environment.](#)

### Related links

link: [End of availability of node-based licenses](#)  
[Convert node-based licenses to capacity based](#)

## Manage PAYGO subscriptions and contracts

When you subscribe to BlueXP data services (including Cloud Volumes ONTAP) from a cloud provider's marketplace, you're redirected to the BlueXP website where you need to save your subscription and associate it with your BlueXP organization. After you've subscribed, each subscription is available to manage from the BlueXP digital wallet.

- [Learn how to subscribe to BlueXP data services \(standard mode\)](#)
- [Learn how to subscribe to BlueXP data services \(restricted mode\)](#)



The **Marketplace subscriptions** page lists all licenses. If you want license details for a specific data service, use the data service tiles on the **Overview** dashboard. [Learn more about the Overview dashboard.](#)

## View your subscriptions

The BlueXP digital wallet provides details about each PAYGO subscription and annual contract associated with your BlueXP organization or account.

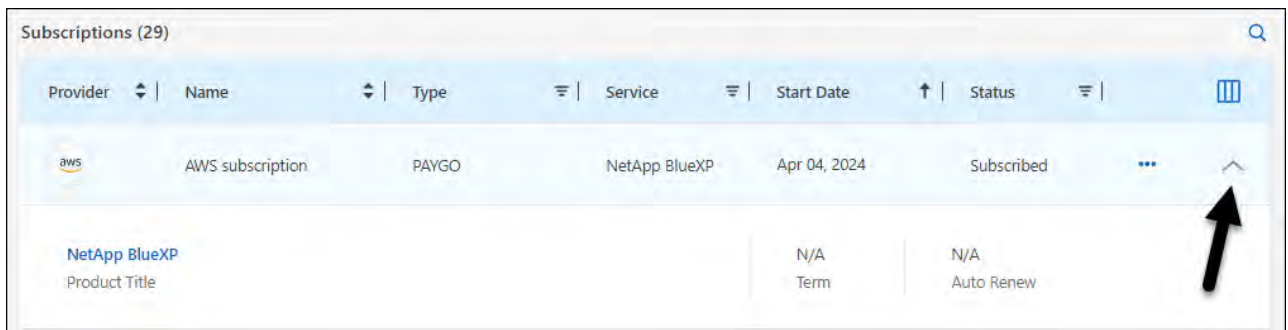
For Google Cloud, you can also identify marketplace subscriptions that are associated with a private offer, which enables you to verify that you've successfully accepted the offer.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Marketplace Subscriptions**.

You'll only see the subscriptions that are associated with the organization that you are that you're currently viewing.

3. As you view the information about your subscriptions, you can interact with the details in the table as follows:
  - Use Advanced Search and Filtering to determine which subscriptions are shown in the table. You can search by a specific name of a subscription and filter by a variety of subscription parameters such as type and configuration.
  - Expand a row to view more details.



Provider	Name	Type	Service	Start Date	Status	
AWS	AWS subscription	PAYGO	NetApp BlueXP	Apr 04, 2024	Subscribed	⋮
NetApp BlueXP	Product Title			N/A	N/A	
				Term	Auto Renew	

- Select  to choose which columns appear in the table.

Note that the Term and Auto Renew columns don't appear by default. The Auto Renew column displays renewal information for Azure contracts only.

Note the following about what you see in the table:

### Start date

The start date is when you successfully associated the subscription with your account and charging started.

### N/A

If you see N/A in the table, the information isn't available from the cloud provider's API at this time.

### Term

If your Google Cloud subscription is associated with a private offer and that private offer was modified after it was created and accepted, then the term shows N/A. In this scenario, the API response that we receive from the Google Cloud Marketplace doesn't include term-related information.

## Contracts

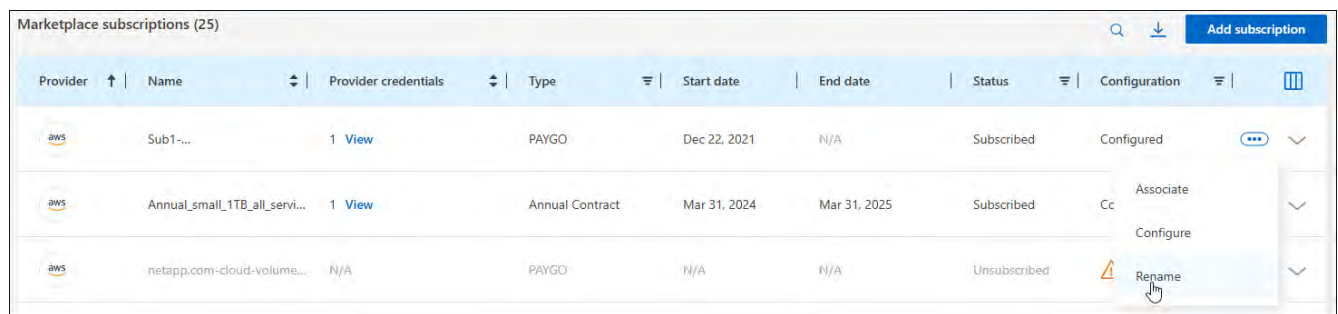
- If you expand the details for a contract, the BlueXP digital wallet shows what's available for your current plan: the contract options and units (capacity or number of nodes).
- The BlueXP digital wallet identifies the end date and whether the contract will renew soon, end soon, or whether it has already ended.
- If you have an AWS contract and you changed any of the contract's options after the start date, be sure to validate your contract options from the AWS Marketplace.
- If you have a Google Cloud private offer, contract options aren't available.

## Rename a subscription

You can rename a subscription to better identify how it is used in your organization.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Marketplace Subscriptions**.
3. Select the action menu in the row that corresponds to the subscription that you want to manage and choose **Rename**.



Provider	Name	Provider credentials	Type	Start date	End date	Status	Configuration
aws	Sub1-...	1 View	PAYGO	Dec 22, 2021	N/A	Subscribed	Configured
aws	Annual_small_1TB_all_servi...	1 View	Annual Contract	Mar 31, 2024	Mar 31, 2025	Subscribed	Cc
aws	netapp.com-cloud-volume...	N/A	PAYGO	N/A	N/A	Unsubscribed	Configure

## Configure a subscription with a provider credential

Subscriptions are typically configured with the provider credential that you created when you subscribed. In some cases, you may need to reconfigure a subscription to use a different credential if you want to change the way it is charged. The credential you associate with a subscription must be a credential that is also associated with a connector.

The format of the credential depends on the marketplace you are using. For example, Azure marketplace subscriptions are associated with the name of the Azure subscription, while AWS marketplace subscriptions use the AWS account ID. You can see a list of available credentials from the Credentials page.

The Configure option is grayed out if you have unsubscribed to a subscription.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Marketplace Subscriptions**.
3. Select the action menu in the row that corresponds to the subscription that you want to configure and choose **Configure**.
4. In the *Configure* dialog, choose a credential to which to configure the subscription. You can only choose from credentials that are associated with the currently selected connector. If you don't see the credential



that you want to use, try switching to a different connector view.

## Associate a subscription with a BlueXP organization

Associating a subscription with an organization ensures members of that organization can use that subscription for charging.

You can limit use of a subscription to a specific organization or share the subscription between multiple organizations.

You must have the organization admin role in order to associate a subscription with an organization.



BlueXP supports Identity and Access Management (IAM) in standard mode which uses organizations to manage users and resources. If you're using BlueXP in private or restricted mode, then you use a BlueXP *account* to manage users and resources, including subscriptions.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Marketplace Subscriptions**.
3. In the row for the subscription you want to associate, open the action menu and select **Associate**.
4. In the **Associate the subscription** dialog, select one or more organizations to associate with the subscription.
5. Select **Associate**.

## View credentials associated with a subscription

You can view the credentials for a specific subscription from the **Marketplace Subscriptions** page in the digital wallet. This allows you to verify how the subscription is being billed. Because credentials are also tied to the connector you are using, you must select the connector associated with the subscription you want to see.



Use the Connector drop-down in the top navigation bar to switch connectors if you need.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Marketplace Subscriptions**.
3. On the row that contains the subscription whose credentials you want to view, select **View**. If there are multiple credentials associated with a subscription, no credentials may show and you are directed to select a different connector.

## Add a new marketplace subscription

You can subscribe to a marketplace subscription directly from digital wallet.

## AWS

The following video shows the steps to subscribe to BlueXP from the AWS Marketplace:

[Subscribe to BlueXP from the AWS Marketplace](#)

## Azure

The following video shows the steps to subscribe from the Azure Marketplace:

[Subscribe to BlueXP from the Azure Marketplace](#)

## Google Cloud

The following video shows the steps to subscribe from the Google Cloud Marketplace:

[Subscribe to BlueXP from the Google Cloud Marketplace](#)

## Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Marketplace Subscriptions**.
3. Above the **Subscriptions** table, select **Add Subscription**.
4. In the *Add Subscription* dialog, select a cloud provider.
  - a. If choosing an AWS subscription, choose whether you want an annual contract or PAYGO subscription.
5. Select **Add subscription** to navigate to the provider's marketplace and complete the steps provided.
6. When finished at the cloud provider marketplace, return to BlueXP to complete the process.

## Unconfigure a subscription

Before you can remove a subscription, you must unconfigure it. This clears all associated data and settings.

## Steps

1. In the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Marketplace Subscriptions**.
3. In the row for the subscription you want to unconfigure, click the action menu and select **Unconfigure**.
4. Follow the prompts to remove or reset any associated settings or data.
5. Wait for the status to update to **Unconfigured**.

## Remove a subscription

When you unsubscribe from a BlueXP subscription in your cloud provider (AWS, Google Cloud, or Azure), the digital wallet shows the subscription status as **Unsubscribed**.

You can remove **Unsubscribed** subscriptions from the digital wallet so they no longer appear.



You can only remove a subscription if it is both **Unsubscribed** and **Unconfigured**. This means all related settings, data, and configuration must be cleared or reset before removal.

If the subscription is still configured, the **Remove** option is not displayed. To make the option available, unconfigure the subscription by clearing any associated settings, services, or data.

### Steps

1. In the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Marketplace Subscriptions**.
3. In the row for the subscription you want to remove, open the action menu and select **Remove**.

You can only remove subscriptions with a status of **Unsubscribed** and **Unconfigured**.

4. In the **Remove subscription** dialog, confirm that you want to remove the subscription.

## Manage Keystone subscriptions

Manage your Keystone subscriptions from the BlueXP digital wallet by enabling subscriptions for use with Cloud Volumes ONTAP and by requesting changes to the committed capacity for your subscription's service levels. Requesting additional capacity for a service level provides more storage for on-premises ONTAP clusters or for Cloud Volumes ONTAP systems.

NetApp Keystone is a flexible pay-as-you-grow subscription-based service that delivers a hybrid cloud experience for customers who prefer OpEx to CapEx or leasing.

[Learn more about Keystone](#)

### Authorize your account

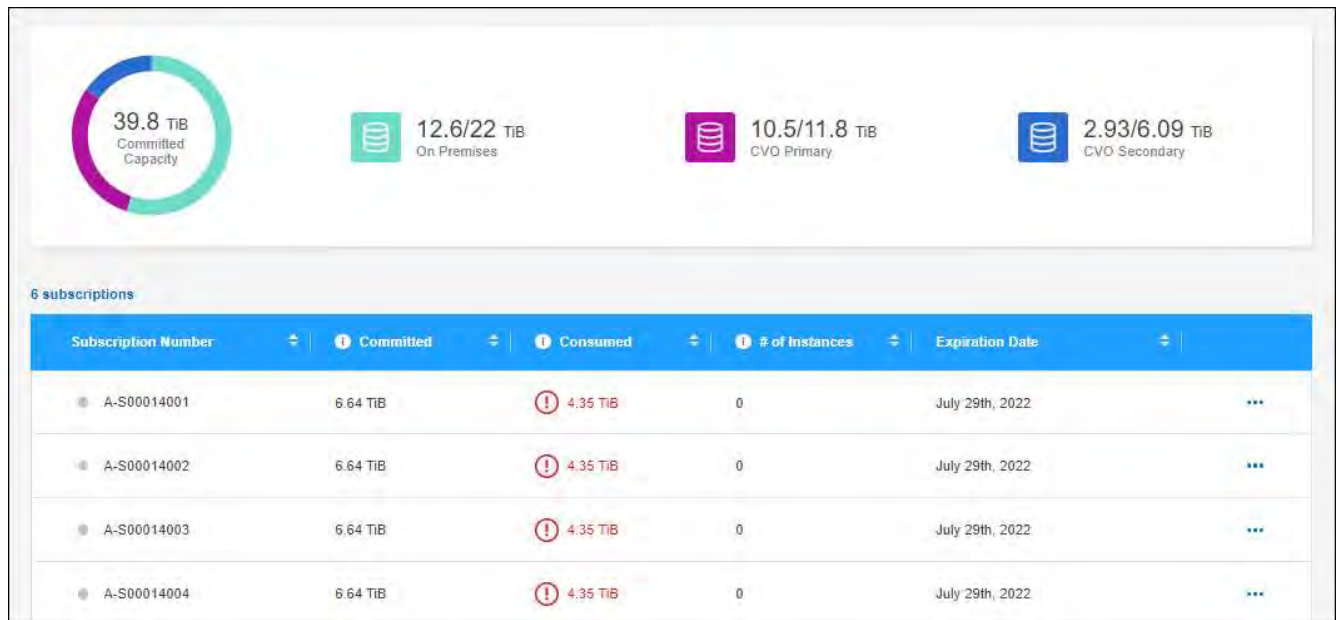
Before you can use and manage Keystone subscriptions in BlueXP, you need to contact NetApp to authorize your BlueXP user account with your Keystone subscriptions.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Keystone Subscriptions**.
3. If you see the **Welcome to NetApp Keystone** page, send an email to the address listed on the page.

A NetApp representative will process your request by authorizing your user account to access the subscriptions.

4. Come back to the **Keystone Subscriptions** tab to view your subscriptions.

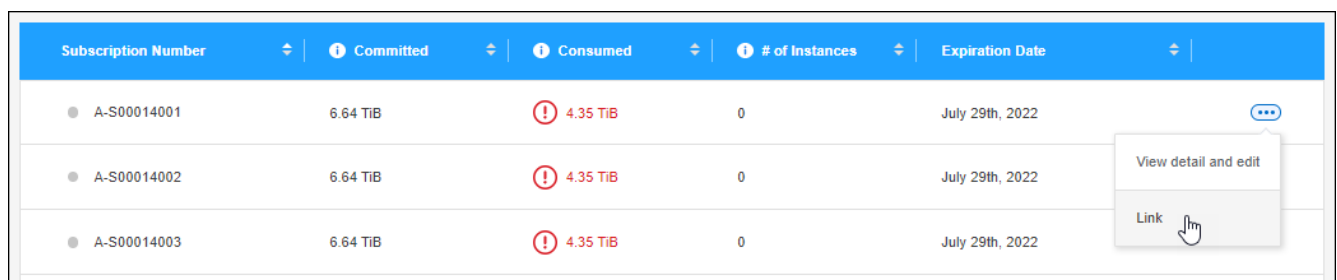


## Link a subscription

After NetApp authorizes your account, you can link Keystone subscriptions for use with Cloud Volumes ONTAP. This action enables users to select the subscription as the charging method for new Cloud Volumes ONTAP systems.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **Keystone Subscriptions**.
3. For the subscription that you want to link, click **...** and select **Link**.



### Result

The subscription is now linked to your BlueXP organization or account and available to select when creating a Cloud Volumes ONTAP working environment.

## Request more or less committed capacity

If you want to change the committed capacity for your subscription's service levels, you can send a request to NetApp directly from BlueXP. Requesting additional capacity for a service level provides more storage for on-premises clusters or for Cloud Volumes ONTAP systems.

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.

2. Select **Keystone Subscriptions**.
3. For the subscription that you want adjust the capacity, click **...** and select **View detail and edit**.
4. Enter the requested committed capacity for one or more subscriptions.

**Subscription Modification for A-S00014001**

Service Level	Current Committed Capacity	Current Consumed Capacity	Requested Committed Capacity
Extreme	0.977 TiB	0.293 TiB	<input type="text" value="Enter amount"/> TiB
Premium	0.977 TiB	0.488 TiB	<input type="text" value="Enter amount"/> TiB
Performance	0 TiB	0 TiB	<input type="text" value="Enter amount"/> TiB
Standard	0.732 TiB	0.439 TiB	<input type="text" value="Enter amount"/> TiB
Value	0.977 TiB	⚠ 0.879 TiB	<input type="text" value="Enter amount"/> TiB
Data Tiering	0 TiB	0 TiB	<input type="text" value="Enter amount"/> TiB
CVO Primary	1.96 TiB	⚠ 1.76 TiB	<input type="text" value="3"/> TiB
CVO Secondary	1.02 TiB	0.488 TiB	<input type="text" value="Enter amount"/> TiB

**Additional Information**

Is there anything else we should know about your request?  
Please be as descriptive as possible.

5. Scroll down, enter any additional details for the request, and then click **Submit**.

## Result

Your request creates a ticket in NetApp's system for processing.

## Monitor usage

The BlueXP digital advisor dashboard enables you to monitor Keystone subscription usage and to generate reports.

[Learn more about monitoring subscription usage](#)

## Unlink a subscription

If you no longer want to use a Keystone subscription with BlueXP, you can unlink the subscription. Note that you can only unlink a subscription that isn't attached to an existing Cloud Volumes ONTAP subscription.

## Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.

2. Select **Keystone**.
3. For the subscription that you want to unlink, click **...** and select **Unlink**.

## Result

The subscription is unlinked from your BlueXP organization or account and no longer available to select when creating a Cloud Volumes ONTAP working environment.

# Manage licenses for on-premises ONTAP

The BlueXP digital wallet enables you to view contract details for each of your on-prem ONTAP clusters. If you haven't discovered a cluster in BlueXP yet, you can also discover them from the digital wallet.

## Before you begin

The BlueXP digital wallet displays details about the on-premises ONTAP clusters that you discovered as a working environment or that are associated with a NetApp Support Site account that you added to BlueXP.

- [Learn how to discover an on-premises ONTAP cluster](#)
- [Learn how to manage NSS credentials associated with your BlueXP organization or account](#)

## View cluster information and contract details

View the status of the hardware and software contracts for your on-premises ONTAP clusters so that you can renew them before they expire.

## Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **On-Premises ONTAP**.

The Software Contract and Hardware Contract expiration dates appear on the line for each cluster.

If you are prompted to enter your NetApp Support Site (NSS) account credentials first, select **Add NSS Account**. After you add the account, the clusters associated with that NSS account appear.

3. If the contract is close to the expiration date, or has expired, you can select the chat icon in the lower-right of the BlueXP console to request an extension to the contract.
4. To view more details, select **▼** to expand the cluster information.

Host name	Status	Capacity	Software Contract	HW Contract	
OnPremisesHostName#1	Discovered	10.25 TB Used   50.25 TB Allocated	January 1, 2025	January 1, 2025	▼

Cluster Name	OnPremises_Cluster_#1	Support Offering	Standard
Cluster Management IP Address	196.10.10.196		
UUID	OnPremises_UUID_#1		

## Discover clusters

If you haven't discovered an on-premises ONTAP cluster as a working environment, you can do that from the BlueXP digital wallet. Once discovered, a cluster is available as a working environment in BlueXP so that you can manage it.

### Before you begin

You should understand your discovery and management options (discovery using a Connector or direct discovery without a Connector) as well as discovery requirements.

[Learn about discovery options and requirements.](#)

### Steps

1. From the BlueXP navigation menu, select **Governance > Digital wallet**.
2. Select **On-Premises ONTAP**.

[A screenshot of the Digital Wallet page for on-premises ONTAP clusters.]

Your ONTAP clusters display with a status of whether they have been discovered in BlueXP.

If you are prompted to enter your NetApp Support Site (NSS) account credentials first, select **Add NSS Account**. After you add the account, the clusters associated with that NSS account appear.

3. Select **Discover** for the cluster that you want to manage through BlueXP.
4. On the *Discover ONTAP Cluster* page, enter the password for the admin user account and select **Discover**.

Note that the cluster management IP address is populated based on information from your NetApp Support Site account.

### Result

BlueXP discovers the cluster and adds it as a working environment in the Canvas. The status for the cluster turns to **Discovered** in the *On-Premises ONTAP* page. Note that the working environment name is the name of the cluster.



You can now start managing the cluster.

- [Learn how to manage clusters discovered with a Connector](#)
- [Learn how to manage clusters discovered directly](#)

# 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 serial number (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 BlueXP for NetApp support

To register for support and activate support entitlement, one user in your BlueXP organization (or 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 BlueXP organization 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 organization is not registered for support. As long as one user in the organization has followed these steps, then your organization 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.



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 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 organization 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 organization is different than the NSS account that is associated with a BlueXP user login.

These NSS credentials are associated with your specific BlueXP organization ID. Users who belong to the BlueXP organization 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 organization (or account), project (or 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 Working Enviroment


Select 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) Upload 

No files selected  

### 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.

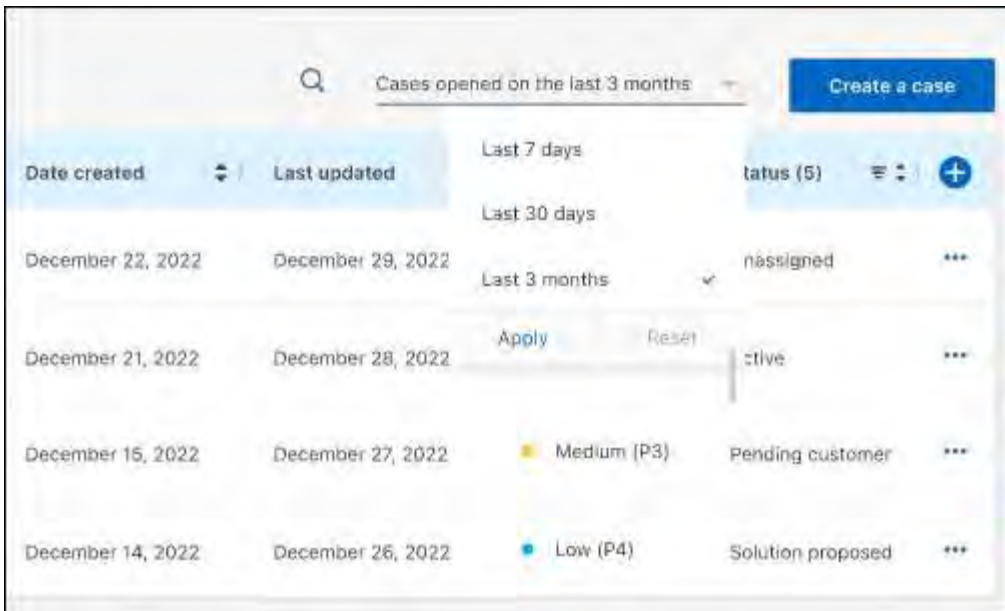
- At a per-case level, we offer the ability to update case notes or close a case that is not already in Closed or Pending Closed status.

### 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.



- Filter the contents of the columns.



- Change the columns that appear in the table by selecting  and then choosing the columns that you'd like to display.



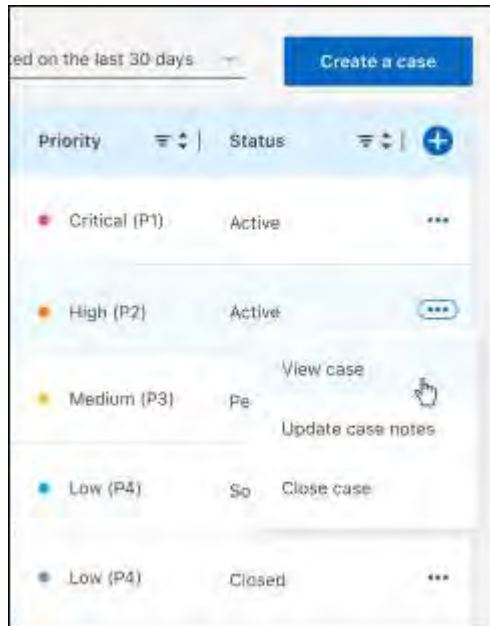


4. Manage an existing case by selecting **...** and selecting one of the available options:

- **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](#)

## Copyright information

Copyright © 2025 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.



# On-premises ONTAP cluster management using BlueXP

On-premises ONTAP clusters

NetApp  
July 24, 2025

# Table of Contents

- On-premises ONTAP cluster management using BlueXP ..... 1
- Release notes ..... 2
  - What's new with on-premises ONTAP clusters ..... 2
    - 12 May 2025 ..... 2
    - 26 November 2024 ..... 2
    - 7 October 2024 ..... 2
    - 22 April 2024 ..... 2
    - 30 July 2023 ..... 2
    - 2 July 2023 ..... 3
    - 4 May 2023 ..... 3
    - 3 April 2023 ..... 3
    - 1 January 2023 ..... 4
    - 4 December 2022 ..... 4
    - 18 September 2022 ..... 5
    - 7 June 2022 ..... 6
    - 27 February 2022 ..... 7
    - 11 January 2022 ..... 7
  - Known limitations ..... 7
    - Limitations related to ASA r2 systems ..... 7
    - Unsupported clusters ..... 7
    - System Manager limitations ..... 8
- Get started ..... 9
  - Learn about on-premises ONTAP cluster management in BlueXP ..... 9
    - Features ..... 9
    - Cost ..... 9
  - Discover on-premises ONTAP clusters ..... 9
    - Step 1: Review discovery and management options ..... 9
    - Step 2: Set up your environment ..... 10
    - Step 3: Discover a cluster ..... 11
- Manage on-prem ONTAP clusters ..... 14
  - Manage clusters that were discovered directly ..... 14
  - Manage clusters that were discovered with a Connector ..... 15
    - Create FlexVol volumes from BlueXP ..... 16
    - Create FlexGroup volumes with the BlueXP API ..... 17
    - Access ONTAP System Manager from BlueXP ..... 17
    - Enable BlueXP services ..... 19
  - View cluster information and contract details ..... 20
  - Optimize clusters using BlueXP digital advisor ..... 20
    - Features ..... 20
    - Supported ONTAP systems ..... 20
    - More information ..... 21
  - Remove an on-prem ONTAP working environment ..... 21
- Knowledge and support ..... 22

Register for support .....	22
Support registration overview .....	22
Register BlueXP for NetApp support .....	22
Associate NSS credentials for Cloud Volumes ONTAP support .....	24
Get help .....	26
Get support for a cloud provider file service .....	26
Use self-support options .....	26
Create a case with NetApp support .....	26
Manage your support cases (Preview) .....	29
Legal notices .....	32
Copyright .....	32
Trademarks .....	32
Patents .....	32
Privacy policy .....	32
Open source .....	32

# On-premises ONTAP cluster management using BlueXP

# Release notes

## What's new with on-premises ONTAP clusters

Learn what's new with on-premises ONTAP cluster management in BlueXP.

### 12 May 2025

#### BlueXP access role needed

You now need one of the following access roles to view, discover or manage on-prem ONTAP clusters: Organization admin, Folder or project admin, Storage admin, or System health specialist. [Learn about BlueXP access roles.](#)

### 26 November 2024

#### Support for ASA r2 systems with private mode

You can now discover NetApp ASA r2 systems when using BlueXP in private mode. This support is available starting with the 3.9.46 private mode release of BlueXP.

- [Learn more about ASA r2 systems](#)
- [Learn about BlueXP deployment modes](#)

### 7 October 2024

#### Support for ASA r2 systems

You can now discover NetApp ASA r2 systems in BlueXP when using BlueXP in standard mode or restricted mode. After you discover a NetApp ASA r2 system and open the working environment, you're brought directly to System Manager.

No other management options are available with ASA r2 systems. You can't use the Standard view and you can't enable BlueXP services.

Discovery of ASA r2 systems is not supported when using BlueXP in private mode.

- [Learn more about ASA r2 systems](#)
- [Learn about BlueXP deployment modes](#)

### 22 April 2024

#### Volume templates no longer supported

You can no longer create a volume from a template. This action was associated with the BlueXP remediation service, which is no longer available.

### 30 July 2023



## Create FlexGroup volumes

If you're managing a cluster with a Connector, you can now create FlexGroup volumes using the BlueXP API.

- [Learn how to create a FlexGroup volume](#)
- [Learn what a FlexGroup volume is](#)

## 2 July 2023

### Cluster discovery from My estate

You can now discover on-premises ONTAP clusters from **Canvas > My estate** by selecting a cluster that BlueXP pre-discovered based on the ONTAP clusters that are associated with the email address for your BlueXP login.

[Learn how to discover clusters from the My estate page.](#)

## 4 May 2023

### Enable BlueXP backup and recovery

Beginning with ONTAP 9.13.1, you can use System Manager (advanced view) to enable BlueXP backup and recovery if you discovered the cluster using a Connector. [Learn more about enabling BlueXP backup and recovery](#)

### Upgrade ONTAP version image and hardware firmware

Beginning with ONTAP 9.10.1, you can use System Manager (advanced view) to upgrade the ONTAP version image and hardware firmware. You can choose to receive automatic upgrades to stay up to date, or you can make manual updates from your local machine or a server that can be accessed using BlueXP. [Learn more about upgrading ONTAP and firmware](#)

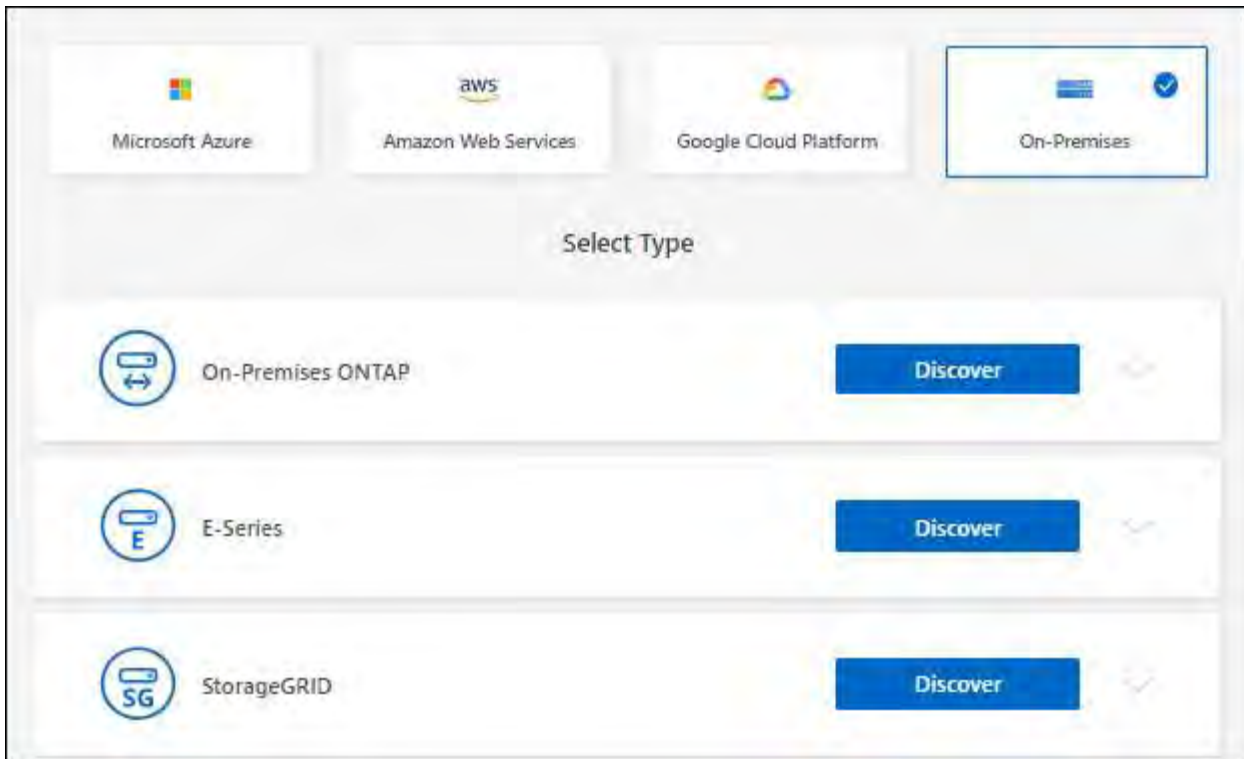


If you do *not* have a Connector, you cannot make updates from your local machine, only from a server that can be accessed using BlueXP.

## 3 April 2023

### Single discovery option from the BlueXP console

When you discover an on-prem ONTAP cluster from the BlueXP console, you'll now see a single option:



Previously, there were separate flows for direct discovery and for discovery with a Connector. Both of those options are still available, but merged into a single flow.

When you start the discovery process, BlueXP discovers the cluster as follows:

- If you have an active Connector that has a connection to your ONTAP cluster, BlueXP will use that Connector to discover and manage the cluster.
- If you don't have a Connector or if your Connector doesn't have a connection to the ONTAP cluster, then BlueXP will automatically use the direct discovery and management option.

[Learn more about the discovery and management options.](#)

## 1 January 2023

### Save ONTAP credentials

When you open an on-premises ONTAP working environment that was discovered directly without using a Connector, you now have the option to save your ONTAP cluster credentials so that you don't need to enter them each time that you open the working environment.

[Learn more about this option.](#)

## 4 December 2022

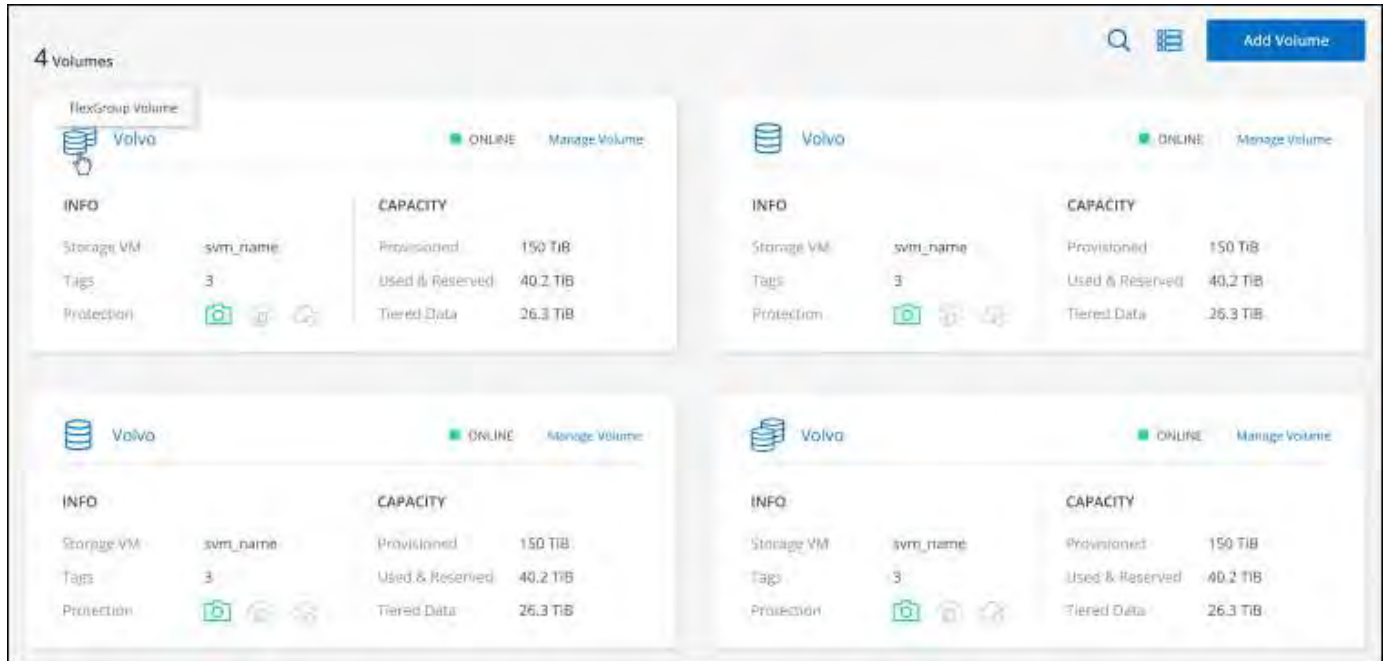
### New way to discover on-premises ONTAP clusters

You can now directly discover your on-premises ONTAP clusters without using a Connector. This option enables cluster management through System Manager only. You can't enable any BlueXP data services on this type of working environment.

[Learn more about this discovery and management option.](#)

## FlexGroup volumes

For on-premises ONTAP clusters that are discovered through a Connector, the Standard view in BlueXP now shows the FlexGroup volumes that were created through System Manager or the ONTAP CLI. You can also manage these volumes by cloning them, editing their settings, deleting them, and more.



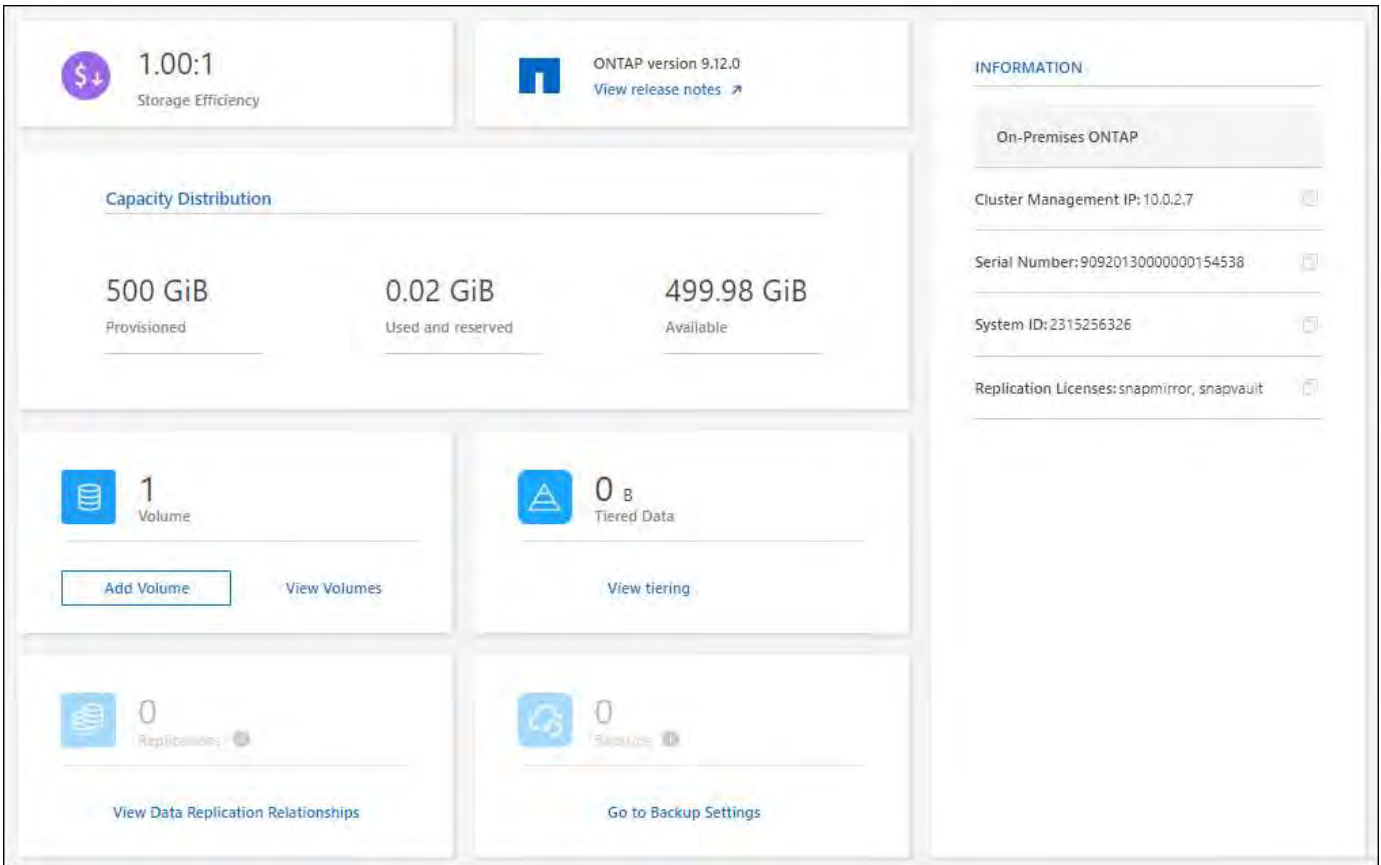
BlueXP does not support creating FlexGroup volumes. You'll need to continue using System Manager or the CLI to create FlexGroup volumes.

## 18 September 2022

### New Overview page

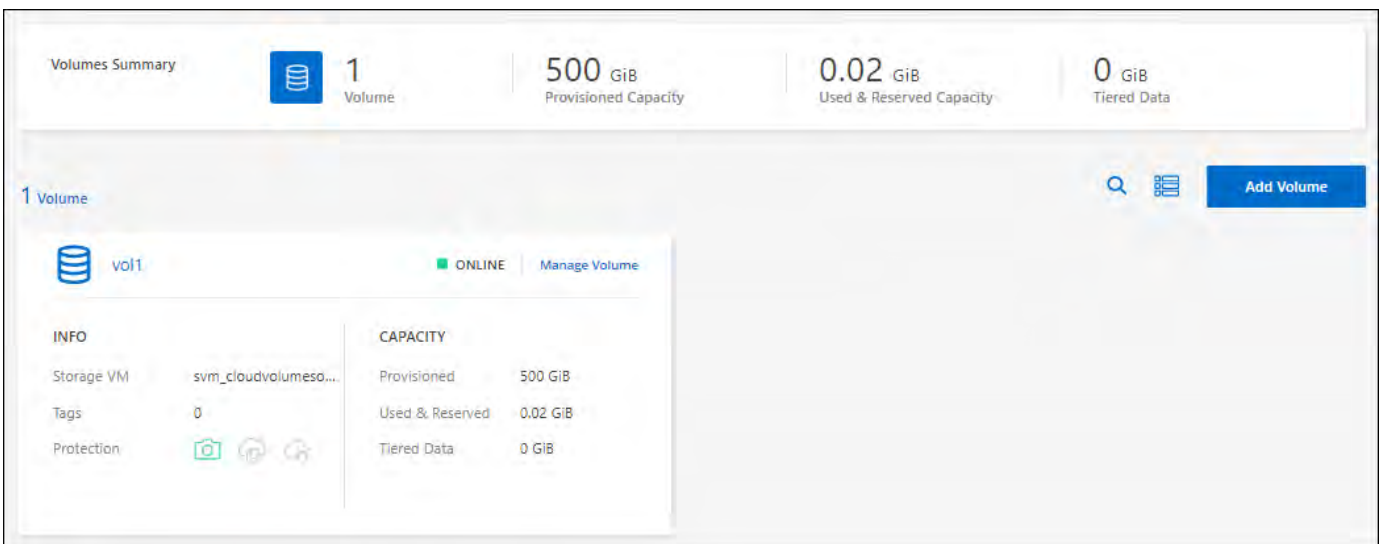
We've introduced a new Overview page to provide key details about an on-premises ONTAP cluster. For example, you can now view details like storage efficiency, capacity distribution, and system information.

You can also view details about integration with other BlueXP services that enable data tiering, data replication, and backups.



## Redesigned Volumes page

We redesigned the Volumes page to provide a summary of the volumes on a cluster. The summary shows you the total number of volumes, the amount of provisioned capacity, used and reserved capacity, and the amount of tiered data.



7 June 2022

## New Advanced View

If you need to perform advanced management of an ONTAP on-premises cluster, you can do so using ONTAP

System Manager, which is a management interface that's provided with an ONTAP system. We have included the System Manager interface directly inside Cloud Manager so that you don't need to leave Cloud Manager for advanced management.

This Advanced View is available as a Preview with on-premises ONTAP clusters running 9.10.0 or later. We plan to refine this experience and add enhancements in upcoming releases. Please send us feedback by using the in-product chat.

- [Learn how to manage clusters discovered directly](#)
- [Learn how to manage clusters discovered with a Connector](#)

## 27 February 2022

### An "On-Premises ONTAP" tab is available in the Digital Wallet

Now you can view an inventory of your on-premises ONTAP clusters along with their hardware and service contracts expiration dates. Additional details about the clusters are also available.

[Learn how to view this important on-prem cluster information.](#) You'll need to have a NetApp Support Site account (NSS) for the clusters, and the NSS credentials will need to be attached to your Cloud Manager account.

## 11 January 2022

### Tags that you add to volumes on on-premises ONTAP clusters can be use with the Tagging service

Tags that you add to a volume are now associated with the tagging feature of the Application Templates service, which can help you organize and simplify the management of your resources.

## Known limitations

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

### Limitations related to ASA r2 systems

After you discover a NetApp ASA r2 system and open the working environment, you're brought directly to System Manager.

No other management options are available with ASA r2 systems. You can't use the Standard view and you can't enable BlueXP services.

[Learn more about ASA r2 systems](#)

### Unsupported clusters

On-premises ONTAP clusters that are configured with SAML authentication are not supported with BlueXP.

## System Manager limitations

The following System Manager features are not supported from BlueXP:

- Cluster setup

After you set the management IP address and configure the admin password on an on-premises ONTAP cluster, you can discover the cluster in BlueXP.

- Role-based access control (Connector only)

Role-based access control from System Manager is not supported when discovering and managing an on-prem ONTAP cluster using a Connector. You are prompted to enter your admin credentials during the discovery process. Those credentials are used for all actions taken from System Manager.

With the direct discovery option, you're prompted to log in with your ONTAP credentials each time that you open the working environment.

- BlueXP backup and recovery (Cloud Backup) activation

The cluster version must be 9.13.1 to enable BlueXP backup and recovery from System Manager.

If you did *not* discover a cluster using a Connector, then you can't use System Manager (advanced view) to enable backup and recovery. However, you can enable backup and recovery on an on-premises cluster directly from BlueXP. [Learn how to get started](#)

- On-demand upgrades

On-demand upgrades of firmware and software are not available if the cluster version is ONTAP 9.9.1 or earlier.

If you do *not* have a Connector, you cannot make updates from your local machine, only from a server that can be accessed using BlueXP.

- Global search
- User interface settings

# Get started

## Learn about on-premises ONTAP cluster management in BlueXP

BlueXP can discover the ONTAP clusters running on AFF/FAS controllers and ONTAP Select. Adding on-premises ONTAP systems to BlueXP enables you to manage all of your storage and data assets from a single interface.

### Features

- Manage NFS and CIFS volumes
- Access ONTAP System Manager for any managed cluster through BlueXP
- Get health and performance observability with BlueXP analysis and control
- Use BlueXP services to replicate, back up, scan, classify, and tier data
- View hardware and software contract status information in the BlueXP digital wallet

### Cost

A cost might be associated, but it depends on the following:

- Whether you deploy a Connector to discover and manage your clusters.

You can install the Connector in the cloud or on your premises. Costs are incurred when you install a Connector in the cloud.

- Whether you use BlueXP services such as backup and recovery, tiering, and classification.

## Discover on-premises ONTAP clusters

Discover on-premises ONTAP clusters from BlueXP so that you can start managing volumes and performing advanced management using ONTAP System Manager, which is available from BlueXP.

### Required BlueXP role:

Organization admin, Folder or project admin, Storage admin, or System health specialist. [Learn about BlueXP access roles.](#)

### Step 1: Review discovery and management options

BlueXP provides two discovery and management options for on-prem ONTAP clusters:

#### Discovery and management using a Connector

This option enables you to manage clusters running ONTAP 8.3 and later by using the following features:

- Provides basic volume operations natively through the BlueXP interface
- ONTAP System Manager (supported with ONTAP 9.10.0 and later), access System Manager for each

respective cluster directly from BlueXP

- Integration with BlueXP services that provide data replication, back up and recovery, data classification, and data tiering
- You must have the Organization admin role to install a Connector. If you don't know if your organization has a Connector, or if you need one created, contact your BlueXP administrator. [Contact your Organization admin.](#)

### Direct discovery and management

This option enables you to manage clusters running ONTAP 9.12.1 and later by using System Manager. No other management options are available. You can't use the Standard view and you can't enable BlueXP services.

This option doesn't require a Connector.

When you access System Manager on an on-premises ONTAP cluster running 9.12.1 or later with connectivity to the BlueXP service, you'll be prompted to manage the cluster directly from BlueXP. If you follow this prompt, it discovers the cluster in BlueXP using the direct discovery option.

Once discovered, your clusters are available as a working environment on the BlueXP Canvas.

If you decide to add a Connector to your BlueXP deployment, you'll need to re-discover your on-prem cluster as a separate working environment on the Canvas. This will enable native BlueXP management and access to BlueXP data services. You then have the option to remove the other working environment.

## Step 2: Set up your environment

Before you discover your on-prem ONTAP clusters, ensure that you've met the following requirements.

### General requirements

- You'll need to have one of the following BlueXP access roles: Organization admin, Folder or project admin, or Storage admin.
- You need the cluster management IP address and the password for the admin user account.
- BlueXP discovers ONTAP clusters using HTTPS. If you use custom firewall policies, the ONTAP cluster must allow inbound HTTPS access through port 443.

The default "mgmt" firewall policy allows inbound HTTPS access from all IP addresses. If you modified this default policy, or if you created your own firewall policy, you must associate the HTTPS protocol with that policy and enable access from the Connector host.

### Requirements for Connector discovery

- The on-premises cluster must be running ONTAP 8.3 or later.
- A Connector must be installed in a cloud provider or on your premises.

If you want to tier cold data to the cloud, then you should review requirements for the Connector based on where you plan to tier cold data.

- [Learn about Connectors](#)
- [Learn how to switch between multiple Connectors](#)
- [Learn about BlueXP tiering](#)
- The Connector host must allow outbound connections through port 443 (HTTPS) and the ONTAP



cluster must allow inbound HTTP access through port 443 to the cluster management LIF.

If the Connector is in the cloud, all outbound communication is allowed by the predefined security group.

### Requirements for direct discovery

- The on-premises cluster must be running ONTAP 9.12.1 or later.
- The cluster must have inbound and outbound connectivity to the BlueXP service:

<https://cloudmanager.cloud.netapp.com/ontap-service/check-service-connection>

- The computer that you're using to access the BlueXP console must have a network connection to the on-prem ONTAP cluster, similar to how you would provide connections to other resources in your private network.

### Step 3: Discover a cluster

Discover your on-prem ONTAP clusters from the Canvas in one of two ways:

- From **Canvas > My Working Environments** by manually adding details about the on-premises ONTAP cluster.
- From **Canvas > My estate** by selecting a cluster that BlueXP pre-discovered based on the ONTAP clusters that are associated with the email address for your BlueXP login.

When you start the discovery process, BlueXP discovers a cluster as follows:

- If you have an active Connector that has a connection to an ONTAP cluster, then BlueXP will use that Connector to discover and manage the cluster.
- If you don't have a Connector or if your Connector doesn't have a connection to the ONTAP cluster, then BlueXP will automatically use the direct discovery and management option.

## Discover a cluster manually

Discover an on-premises ONTAP cluster in BlueXP by entering the cluster management IP address and the password for the admin user account.

### Steps

1. From the navigation menu, select **Storage > Canvas**.
2. On the Canvas page, select **Add Working Environment > On-Premises**.
3. Next to On-Premises ONTAP, select **Discover**.
4. On the *Discover* page, enter the cluster management IP address, and the password for the admin user account.
5. If you're discovering the cluster directly (without a Connector), you can select **Save the credentials**.

If you select this option, you won't need to re-enter the credentials each time that you open the working environment. These credentials are only associated with your BlueXP user login. They aren't saved for use by anyone else in the BlueXP organization.

6. Select **Discover**.

If you don't have a Connector and the IP address isn't reachable from BlueXP, then you'll be prompted to create a Connector.

### Result

BlueXP discovers the cluster and adds it as a working environment on the Canvas. You can now start managing the cluster.

- [Learn how to manage clusters discovered directly](#)
- [Learn how to manage clusters discovered with a Connector](#)

## Add a pre-discovered cluster

BlueXP automatically discovers information about the ONTAP clusters that are associated with the email address for your BlueXP login and displays them on the **My estate** page as undiscovered clusters. You can view the list of undiscovered clusters and add them one at a time.


### About this task

Note the following about the on-premises ONTAP clusters that appear on the My estate page:

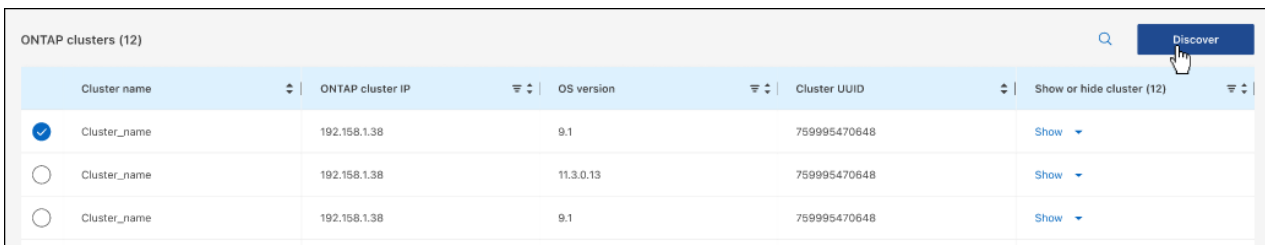
- The email address that you use to log in to BlueXP must be associated with a registered, full-level NetApp Support Site (NSS) account.
  - If you log in to BlueXP with your NSS account and navigate to the My estate page, BlueXP uses that NSS account to find the clusters that are associated with the account.
  - If you log in to BlueXP with a cloud account or a federated connection and you navigate to the My estate page, BlueXP prompts you to verify your email. If that email address is associated with an NSS account, BlueXP uses that information to find the clusters that are associated with the account.
- BlueXP only shows the ONTAP clusters that have successfully sent AutoSupport messages to NetApp.
- To refresh the inventory list, exit the My estate page, wait 5 minutes, and then go back to it.

## Steps

1. From the navigation menu, select **Storage > Canvas**.
2. Select **My estate**.
3. On the My estate page, select **Discover** for on-premises ONTAP.



4. Select a cluster and then select **Discover**.



ONTAP clusters (12)						Discover
Cluster name	ONTAP cluster IP	OS version	Cluster UUID	Show or hide cluster (12)		
<input checked="" type="radio"/> Cluster_name	192.158.1.38	9.1	759995470648	Show		
<input type="radio"/> Cluster_name	192.158.1.38	11.3.0.13	759995470648	Show		
<input type="radio"/> Cluster_name	192.158.1.38	9.1	759995470648	Show		

5. Enter the password for the admin user account.
6. Select **Discover**.

If you don't have a Connector and the IP address isn't reachable from BlueXP, then you'll be prompted to create a Connector.

## Result

BlueXP discovers the cluster and adds it as a working environment on the Canvas. You can now start managing the cluster.

- [Learn how to manage clusters discovered directly](#)
- [Learn how to manage clusters discovered with a Connector](#)

# Manage on-prem ONTAP clusters

## Manage clusters that were discovered directly

If you discovered your on-prem ONTAP cluster directly without using a Connector, you can open the working environment to manage the cluster by using ONTAP System Manager.

### Required BlueXP role:

Organization admin, Folder or project admin, Storage admin, or System health specialist. [Learn about BlueXP access roles.](#)

### Before you begin

The computer that you're using to access the BlueXP console must have a network connection to the on-prem ONTAP cluster, similar to how you would provide connections to other resources in your private network.

### Limitations

A few System Manager features are not supported from BlueXP.

[Review the list of limitations.](#)

### Steps

1. On the Canvas page, select the on-premises ONTAP working environment.

The working environment icon identifies clusters that were discovered directly:



2. If prompted, enter your ONTAP credentials.

You're prompted to log in with your ONTAP credentials each time that you open the working environment, if you don't save the credentials. You have the option to save the credentials so that you don't need to enter them each time. You can manage these credentials on the User Credentials page. In some cases, your BlueXP administrator (Organization admin role) may have disabled this option and require you to enter your credentials each time.

### ONTAP Cluster Credentials

Enter credentials for ONTAP Cluster

ONTAP Cluster IP: 192.168. 1.1

User name

Password

Save the credentials ⓘ

### 3. Use System Manager to manage ONTAP.

If you need help using System Manager with ONTAP, you can refer to [ONTAP documentation](#) for step-by-step instructions. Here are a few links that might help:

- [Volume and LUN management](#)
- [Network management](#)
- [Data protection](#)

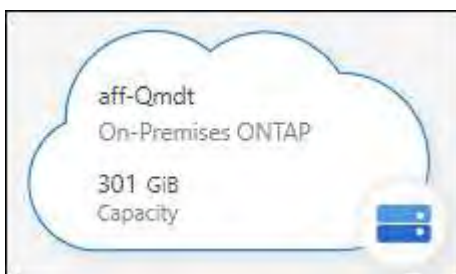
## Manage clusters that were discovered with a Connector

If you discovered an on-premises ONTAP cluster using a Connector, you can create volumes directly from the BlueXP interface, access ONTAP System Manager (from BlueXP) for advanced management, and enable BlueXP data services.

### Required BlueXP role:

Organization admin, Folder or project admin, Storage admin, or System health specialist. [Learn about BlueXP access roles.](#)

On the Canvas, the working environment icon for a cluster that you discovered with a Connector looks similar to the following:



If a working environment was discovered directly, the working environment icon includes the word "Direct."

## Create FlexVol volumes from BlueXP

After you discover your on-prem ONTAP cluster from BlueXP using a Connector, you can provision and manage FlexVol volumes directly from BlueXP.

BlueXP enables you to create NFS or CIFS volumes on existing aggregates. You can't create new aggregates on an on-prem ONTAP cluster from the native BlueXP interface; however, you can access the respective ONTAP System Manager from BlueXP to create aggregates.

### Steps

1. From the navigation menu, select **Storage > Canvas**.
2. On the Canvas page, select the on-prem ONTAP cluster on which you want to provision volumes.
3. Select **Volumes > Add Volume**.
4. Follow the steps in the wizard to create the volume.
  - a. **Details, Protection, & Tags:** Enter details about the volume like its name and size and choose a Snapshot policy.

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

Field	Description
Size	The maximum size that you can enter largely depends on whether you enable thin provisioning, which enables you to create a volume that is bigger than the physical storage currently available to it.
Snapshot Policy	A Snapshot copy policy specifies the frequency and number of automatically created NetApp Snapshot copies. A NetApp Snapshot copy is a point-in-time file system image that has no performance impact and requires minimal storage. You can choose the default policy or none. You might choose none for transient data: for example, tempdb for Microsoft SQL Server.

- b. **Protocol:** Choose the protocol for the volume (NFS, CIFS, or iSCSI) and then set the access control or permissions for the volume.

If you choose CIFS and a server isn't set up yet, then BlueXP prompts you to set up a CIFS server using either Active Directory or a workgroup.

The following list describes fields for which you might need guidance:

Field	Description
Access Control	An NFS export policy defines the clients in the subnet that can access the volume. By default, BlueXP enters a value that provides access to all instances in the subnet.
Permissions and Users/Groups	These fields enable you to control the level of access to an SMB share for users and groups (also called access control lists or ACLs). You can specify local or domain Windows users or groups, or UNIX users or groups. If you specify a domain Windows user name, you must include the user's domain using the format domain\username.

- c. **Usage Profile:** Choose whether to enable or disable storage efficiency features on the volume in order to reduce the total amount of storage that you need.
- d. **Review:** Review details about the volume and then select **Add**.

## Create FlexGroup volumes with the BlueXP API

You can use the BlueXP API to create FlexGroup volumes. A FlexGroup volume is a scale-out volume that provides high performance along with automatic load distribution.

- [Learn how to create a FlexGroup volume using the API](#)
- [Learn what a FlexGroup volume is](#)

## Access ONTAP System Manager from BlueXP

If you need to perform advanced management of an on-premises ONTAP cluster, you can do so using ONTAP System Manager. You can access the ONTAP System Manager interface directly inside BlueXP so that you don't need to leave BlueXP for advanced management.

### Features

When you access ONTAP System Manager from BlueXP you have access to additional management features:

- **Advanced storage management**  
Manage consistency groups, shares, qtrees, quotas, and Storage VMs.
- **Networking management**  
Manage IPspaces, network interfaces, portsets, and ethernet ports.
- **Events and jobs**  
View event logs, system alerts, jobs, and audit logs.
- **Advanced data protection**  
Protect storage VMs, LUNs, and consistency groups.
- **Host management**  
Set up SAN initiator groups and NFS clients.

### Supported configurations

Advanced management through System Manager is supported with on-premises ONTAP clusters running 9.10.0 or later.

System Manager integration is not supported in GovCloud regions or in regions that have no outbound internet access.

### Limitations

A few System Manager features are not supported with on-premises ONTAP clusters when accessing ONTAP

System Manager through BlueXP.

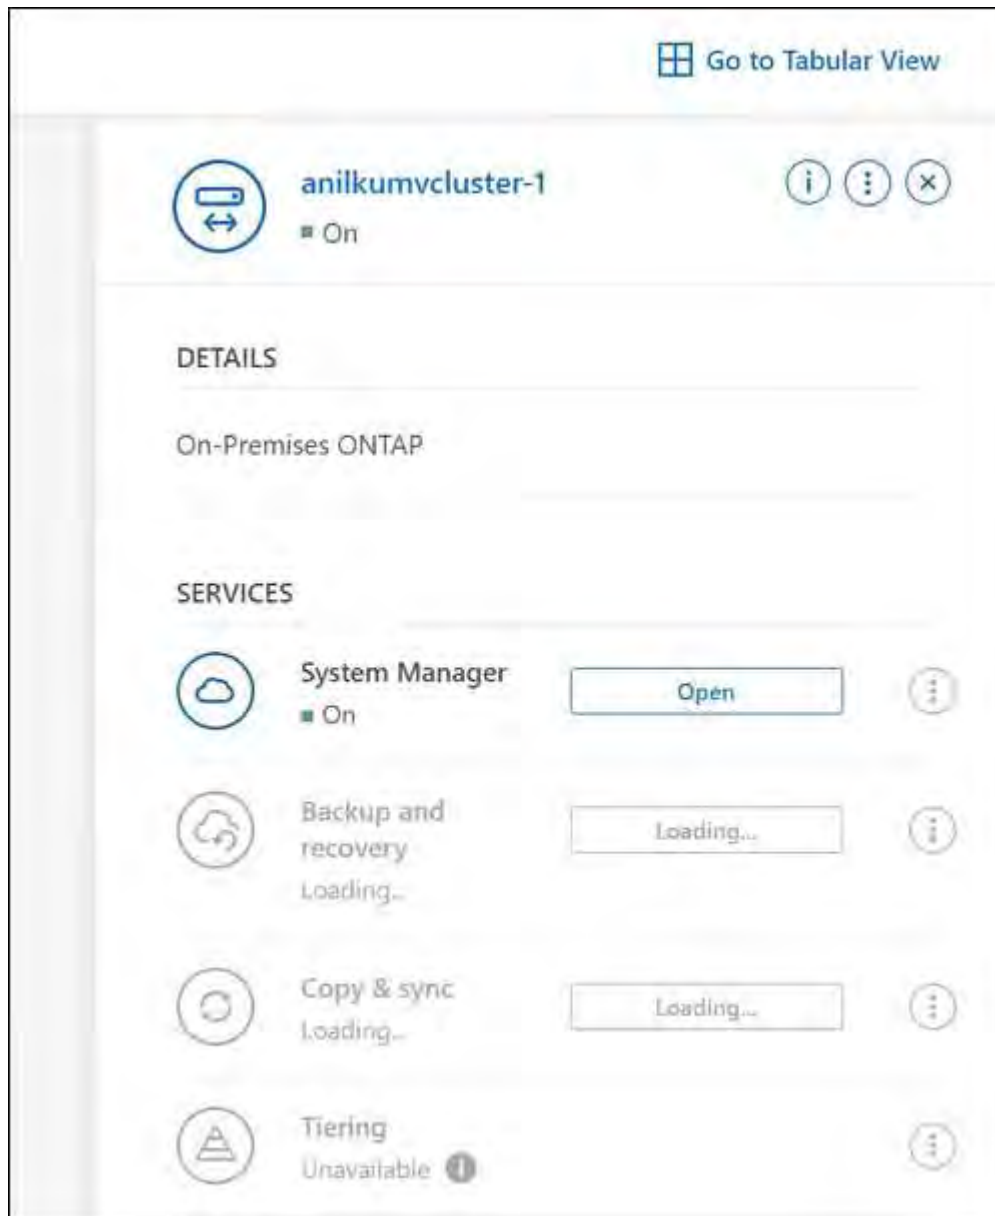
[Review the list of limitations.](#)

### Access ONTAP System Manager from BlueXP

Open an on-premises ONTAP working environment and open the System Manager for the environment.

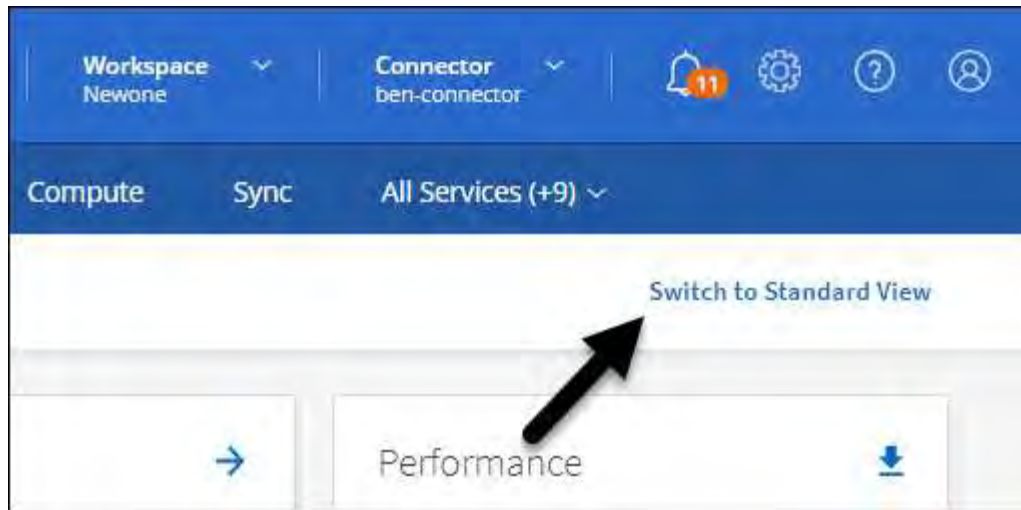
#### Steps

1. On the Canvas page, select the on-prem ONTAP cluster on which you want to provision volumes.
2. From the right panel, under **Services**, find **System Manager** and select **Open**.



3. If the confirmation message appears, read through it and select **Close**.
4. Use System Manager to manage ONTAP.
5. If needed, select **Switch to Standard View** to return to standard management through BlueXP.





## Get help with System Manager

If you need help using System Manager with ONTAP, you can refer to [ONTAP documentation](#) for step-by-step instructions. Here are a few links that might help:

- [Volume and LUN management](#)
- [Network management](#)
- [Data protection](#)

## Enable BlueXP services

Enable BlueXP data services on your working environments to replicate data, back up data, tier data, and more.

### Replicate data

Replicate data between Cloud Volumes ONTAP systems, Amazon FSx for ONTAP file systems, and ONTAP clusters. Choose a one-time data replication, which can help you move data to and from the cloud, or a recurring schedule, which can help with disaster recovery or long-term data retention.

[Replication documentation](#)

### Back up data

Back up data from your on-premises ONTAP system to low-cost object storage in the cloud.

[Backup and recovery documentation](#)

### Scan, map, and classify your data

Scan your corporate on-premises clusters to map and classify data, and to identify private information. This can help reduce your security and compliance risk, decrease storage costs, and assist with your data migration projects.

[Classification documentation](#)

### Tier data to the cloud

Extend your data center to the cloud by automatically tiering inactive data from ONTAP clusters to object storage.

[Tiering documentation](#)

### **Maintain health, uptime, and performance**

Implement suggested remediations to ONTAP clusters before an outage or failure occurs.

[Operational resiliency documentation](#)

### **Identify clusters with low capacity**

Identify clusters that are showing low capacity, review clusters for current and forecasted capacity, and more.

[Economic efficiency documentation](#)

## **View cluster information and contract details**

The BlueXP digital wallet enables you to view contract details for each of your on-prem ONTAP clusters. If you haven't discovered a cluster in BlueXP yet, you can also do that from the digital wallet.

### **Required BlueXP role:**

Organization admin, or Folder or project admin. [Learn about BlueXP access roles.](#)

[Learn more about managing licenses for on-prem ONTAP clusters from the BlueXP digital wallet](#)

## **Optimize clusters using BlueXP digital advisor**

BlueXP digital advisor enables you to optimize the operations, security, and performance of your ONTAP clusters.

### **Features**

You can view the overall status of your storage system, high-level information about the wellness of the system, inventory, planning, upgrades, and valuable insights at a watchlist level using BlueXP digital advisor.

- Analyze and optimize the health of your storage systems
- Gain insights regarding all the risks to your storage systems and the actions to mitigate the risks
- Analyze the performance of your storage devices by viewing the graphical format of performance data
- Get details about systems that have exceeded 90% capacity or are nearing 90% capacity
- Get information about the hardware and software that have expired or are near-expiration within the next 6 months
- Upgrade your storage system software, and update your ONTAP firmware using Ansible

### **Supported ONTAP systems**

Digital advisor provides information for all the on-premises ONTAP systems and Cloud Volumes ONTAP systems associated with your NetApp Support Site (NSS) account.

## More information

[Digital advisor documentation](#)

# Remove an on-prem ONTAP working environment

Remove an on-premises ONTAP working environment if you no longer want to manage it from BlueXP.

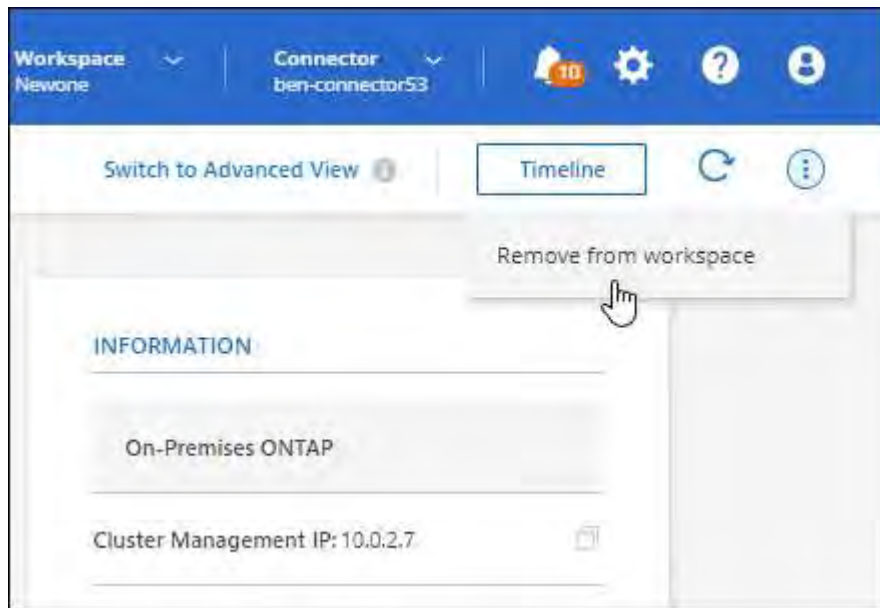
### Required BlueXP role:

Organization admin, Folder or project admin, or Storage admin. [Learn about BlueXP access roles.](#)

Removing the working environment doesn't affect the ONTAP cluster. You can rediscover it from BlueXP at any time.

### Steps

1. On the Canvas page, select the on-premises ONTAP working environment.
2. Select the menu icon and select **Remove from workspace**.



3. Select **Remove** to confirm.

# 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](#)
- [Google Cloud NetApp Volumes](#)

## Support registration overview

There are two forms of registration to activate support entitlement:

- Registering your BlueXP account serial number (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 BlueXP for NetApp support

To register for support and activate support entitlement, one user in your BlueXP organization (or 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 BlueXP organization 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 organization is not registered for support. As long as one user in the organization has followed these steps, then your organization 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.



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 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 organization 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 organization is different than the NSS account that is associated with a BlueXP user login.

These NSS credentials are associated with your specific BlueXP organization ID. Users who belong to the BlueXP organization 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](#)
- [Google Cloud NetApp Volumes](#)

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 organization (or account), project (or 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 Working Environment


Select 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)  Upload 

No files selected  

### 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.



- Filter the contents of the columns.



- Change the columns that appear in the table by selecting  and then choosing the columns that you'd like to display.

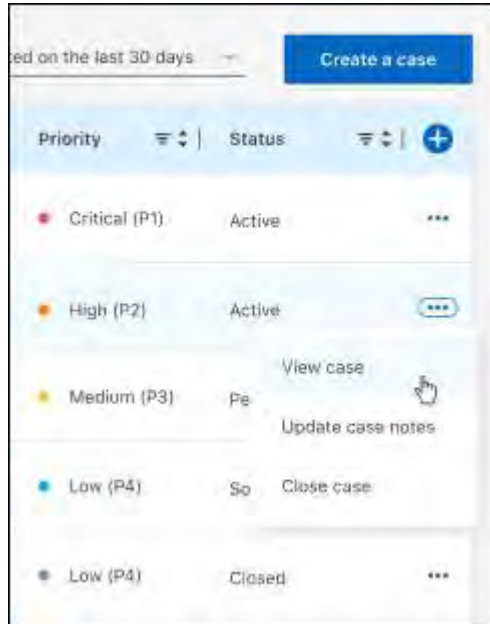


4. Manage an existing case by selecting **...** and selecting one of the available options:

- **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](#)

## Copyright information

Copyright © 2025 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.