



Activate scanning on your data sources

BlueXP classification

NetApp
September 12, 2024

Table of Contents

- Activate scanning on your data sources. 1
 - Scan Azure NetApp Files volumes with BlueXP classification 1
 - Scan Amazon FSx for ONTAP volumes with BlueXP classification 5
 - Scan Cloud Volumes ONTAP and on-premises ONTAP volumes with BlueXP classification 11
 - Scan database schemas with BlueXP classification 17
 - Scan file shares with BlueXP classification 20
 - Scan StorageGRID data with BlueXP classification. 24

Activate scanning on your data sources

Scan Azure NetApp Files volumes with BlueXP classification

Complete a few steps to get started with BlueXP classification for Azure NetApp Files.

Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

Discover the Azure NetApp Files systems you want to scan

Before you can scan Azure NetApp Files volumes, [BlueXP must be set up to discover the configuration](#).

2

Deploy the BlueXP classification instance

[Deploy BlueXP classification in BlueXP](#) if there isn't already an instance deployed.

3

Enable BlueXP classification and select the volumes to scan

Click **Compliance**, select the **Configuration** tab, and activate compliance scans for volumes in specific working environments.

4

Ensure access to volumes

Now that BlueXP classification is enabled, ensure that it can access all volumes.

- The BlueXP classification instance needs a network connection to each Azure NetApp Files subnet.
- Make sure these ports are open to the BlueXP classification instance:
 - For NFS – ports 111 and 2049.
 - For CIFS – ports 139 and 445.
- NFS volume export policies must allow access from the BlueXP classification instance.
- BlueXP classification needs Active Directory credentials to scan CIFS volumes.

Click **Compliance** > **Configuration** > **Edit CIFS Credentials** and provide the credentials.

5

Manage the volumes you want to scan

Select or deselect the volumes that you want to scan and BlueXP classification will start or stop scanning them.

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.

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 on your Azure NetApp Files volumes.

1. From the BlueXP left navigation menu, click **Governance > Classification** and then select the **Configuration** tab.



2. 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.

See [Enable and disable compliance scans on volumes](#) for details.

3. 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.



- 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.

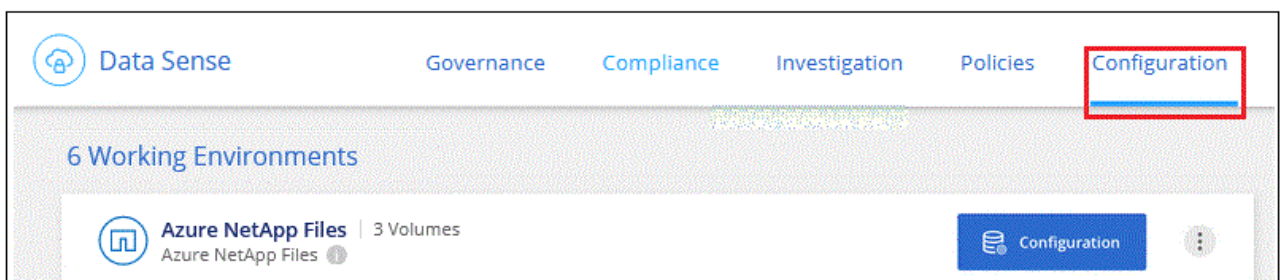
Steps

1. Make sure that there's a network connection between the BlueXP classification instance and each network that includes volumes for Azure NetApp Files.



For Azure NetApp Files, BlueXP classification can only scan volumes that are in the same region as BlueXP.

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, click **Governance > Classification** and then select the **Configuration** tab.



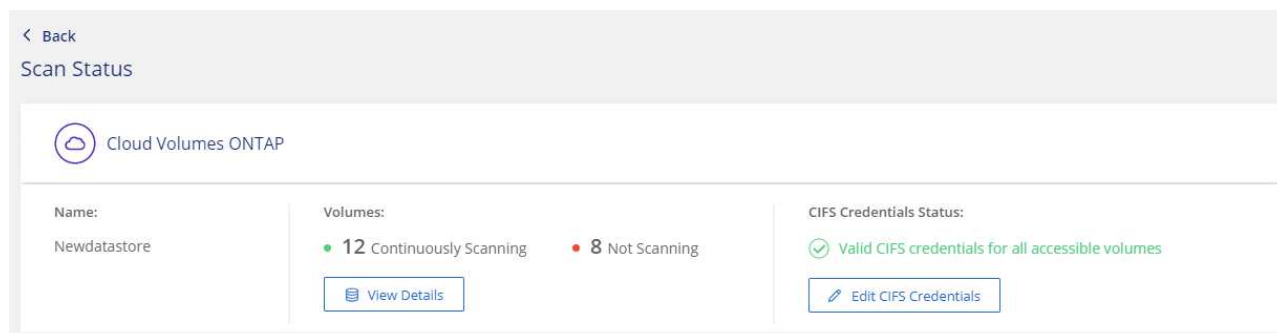
- b. For each working environment, click **Edit CIFS Credentials** and enter the user name and password that BlueXP classification needs to access CIFS volumes on the system.

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

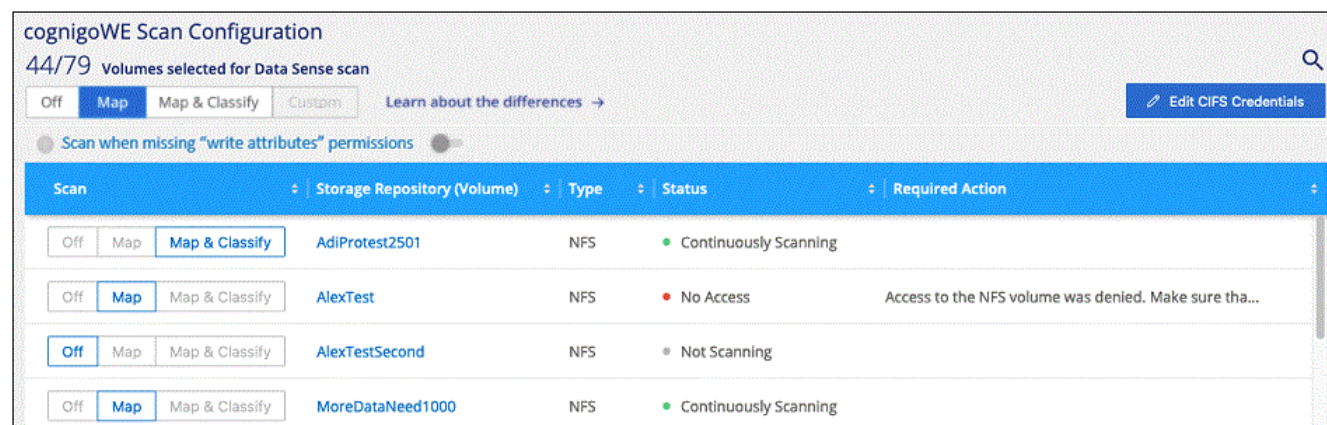
scans, we recommend that the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, we recommend making the Active Directory configured user part of a parent group in the organization which has permissions to all files.

After you enter the credentials, you should see a message that all CIFS volumes were authenticated successfully.



5. On the *Configuration* page, click **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.

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](#).

cognitoWE Scan Configuration
44/79 Volumes selected for Data Sense scan

Off
Map
Map & Classify
Custom
Learn about the differences →
Edit CIFS Credentials

☐ Scan when missing "write attributes" permissions

Scan	Storage Repository (Volume)	Type	Status	Required Action
Off Map Map & Classify	AdiNFSVol_copy	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
Off Map Map & Classify	AdiProtest2501	NFS	Continuously Scanning	
Off Map Map & Classify	AlexTest	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
Off Map Map & Classify	AlexTestSecond	NFS	Not Scanning	

To:	Do this:
Enable mapping-only scans on a volume	In the volume area, click Map
Enable full scanning on a volume	In the volume area, click Map & Classify
Disable scanning on a volume	In the volume area, click Off
Enable mapping-only scans on all volumes	In the heading area, click Map
Enable full scanning on all volumes	In the heading area, click Map & Classify
Disable scanning on all volumes	In the heading area, click 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 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 \(ENI\)](#)

Quick start

Get started quickly by following these steps or scroll down for full details.

1

Discover the FSx for ONTAP file systems you want to scan

Before you can scan FSx for ONTAP volumes, [you must have an FSx working environment with volumes configured](#).

2

Deploy the BlueXP classification instance

[Deploy BlueXP classification in BlueXP](#) if there isn't already an instance deployed.

3

Enable BlueXP classification and select the volumes to scan

Select the **Configuration** tab and activate compliance scans for volumes in specific working environments.

4

Ensure access to volumes

Now that BlueXP classification is enabled, ensure that it can access all volumes.

- The BlueXP classification instance needs a network connection to each FSx for ONTAP subnet.
- Make sure the following ports are open to the BlueXP classification instance:
 - For NFS – ports 111 and 2049.
 - For CIFS – ports 139 and 445.
- NFS volume export policies must allow access from the BlueXP classification instance.
- BlueXP classification needs Active Directory credentials to scan CIFS volumes.

Click **Compliance > Configuration > Edit CIFS Credentials** and provide the credentials.

5

Manage the volumes you want to scan

Select or deselect the volumes you want to scan and BlueXP classification will start or stop scanning them.

Discover the FSx for ONTAP file system that you want to scan

If the FSx for ONTAP file 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 or create the FSx for ONTAP file system in BlueXP](#).

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, click **Governance > Classification** and then select the **Configuration** tab.



2. 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.

See [Enabling and disabling compliance scans on volumes](#) for details.

3. 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.



- 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. On the *Configuration* page, click **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
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map & Classify"/>	jrmclone	NFS	● No Access	Check network connectivity between the Data Sense ...

2. 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.

3. Ensure the following ports are open to the BlueXP classification instance.
 - For NFS – ports 111 and 2049.
 - For CIFS – ports 139 and 445.
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 left navigation menu, click **Governance > Classification** and then select the **Configuration** tab.
 - b. For each working environment, click **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, we recommend that the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, we recommend making the Active Directory configured user 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](#).

cognigoWE Scan Configuration

44/79 Volumes selected for Data Sense scan

OffMapMap & ClassifyCustom

Learn about the differences →

Edit CIFS Credentials

Scan when missing "write attributes" permissions

Scan	Storage Repository (Volume)	Type	Status	Required Action
<div>OffMapMap & Classify</div>	AdiNFSVol_copy	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
<div>OffMapMap & Classify</div>	AdiProtest2501	NFS	Continuously Scanning	
<div>OffMapMap & Classify</div>	AlexTest	NFS	No Access	Access to the NFS volume was denied. Make sure tha...
<div>OffMapMap & Classify</div>	AlexTestSecond	NFS	Not Scanning	

To:	Do this:
Enable mapping-only scans on a volume	In the volume area, click Map
Enable full scanning on a volume	In the volume area, click Map & Classify
Disable scanning on a volume	In the volume area, click Off
Enable mapping-only scans on all volumes	In the heading area, click Map
Enable full scanning on all volumes	In the heading area, click Map & Classify
Disable scanning on all volumes	In the heading area, click 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*.

'Working Environment Name' Configuration 22/28 Volumes selected for compliance scan

Enable Access to DP Volumes
Edit CIFS Credentials

[Learn about the differences →](#)

☒ Scan when missing "write attributes" permissions

Scan	Storage Repository (Volume)	Type	Status	Required Action
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map & Classify"/>	VolumeName1	DP	Not Scanning	Enable access to DP Volumes ⓘ
<input type="button" value="Off"/> <input checked="" type="button" value="Map"/> <input type="button" value="Map & Classify"/>	VolumeName2	NFS	Continuously Scanning	
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map & Classify"/>	VolumeName3	CIFS	Not Scanning	

Steps

If you want to scan these data protection volumes:

- Click **Enable Access to DP volumes** at the top of the page.
- Review the confirmation message and click **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.

Provide Active Directory Credentials

☒ Use existing CIFS Scanning Credentials (user1@domain2)
 ☐ Use Custom Credentials

Active Directory Domain ⓘ
 DNS IP Address ⓘ

DP Volumes, created from a SnapMirror relationship, do not allow external access by default. Continuing will create NFS shares from DP Volumes which have been activated for **Data Sense**. The shares' export policies will allow access only from the Cloud **Data Sense** instance. [Learn More](#)

Provide Active Directory Credentials

☐ Use existing CIFS Scanning Credentials (user1@domain2)
 ☒ Use Custom Credentials

Username ⓘ
 Password

Active Directory Domain ⓘ
 DNS IP Address ⓘ

DP Volumes, created from a SnapMirror relationship, do not allow external access by default. Continuing will create NFS shares from DP Volumes which have been activated for **Data Sense**. The shares' export policies will allow access only from the Cloud **Data Sense** instance. [Learn More](#)

- Activate each DP volume that you want to scan [the same way you enabled other volumes](#).

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.

Note: 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. Click this button and add CIFS credentials to enable access to these CIFS DP volumes.



Active Directory credentials are only registered 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.

Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

Discover the data sources that you want to scan

Before you can scan volumes, you must add the systems as working environments in BlueXP:

- For Cloud Volumes ONTAP systems, these working environments should already be available in BlueXP
- For on-premises ONTAP systems, [BlueXP must discover the ONTAP clusters](#)

2

Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

3

Enable BlueXP classification and select the volumes to scan

Select the **Configuration** tab and activate compliance scans for volumes in specific working environments.

4

Ensure access to volumes

Now that BlueXP classification is enabled, ensure that it can access all volumes.

- The BlueXP classification instance 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.
- Make sure these ports are open to the BlueXP classification instance:
 - For NFS - ports 111 and 2049.
 - For CIFS - ports 139 and 445.
- NFS volume export policies must allow access from the BlueXP classification instance.
- BlueXP classification needs Active Directory credentials to scan CIFS volumes.

Click **Compliance > Configuration > Edit CIFS Credentials** and provide the credentials.

5

Manage the volumes you want to scan

Select or deselect the volumes that you want to scan and BlueXP classification will start or stop scanning them.

Discover the data sources that you want to scan

If the data sources you want to scan are not already in your BlueXP environment, you can add them to the canvas at this time.

Your Cloud Volumes ONTAP systems should already be available in the Canvas in BlueXP. For on-premises ONTAP systems, you'll need to have [BlueXP discover these clusters](#).

Deploy the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

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](#).

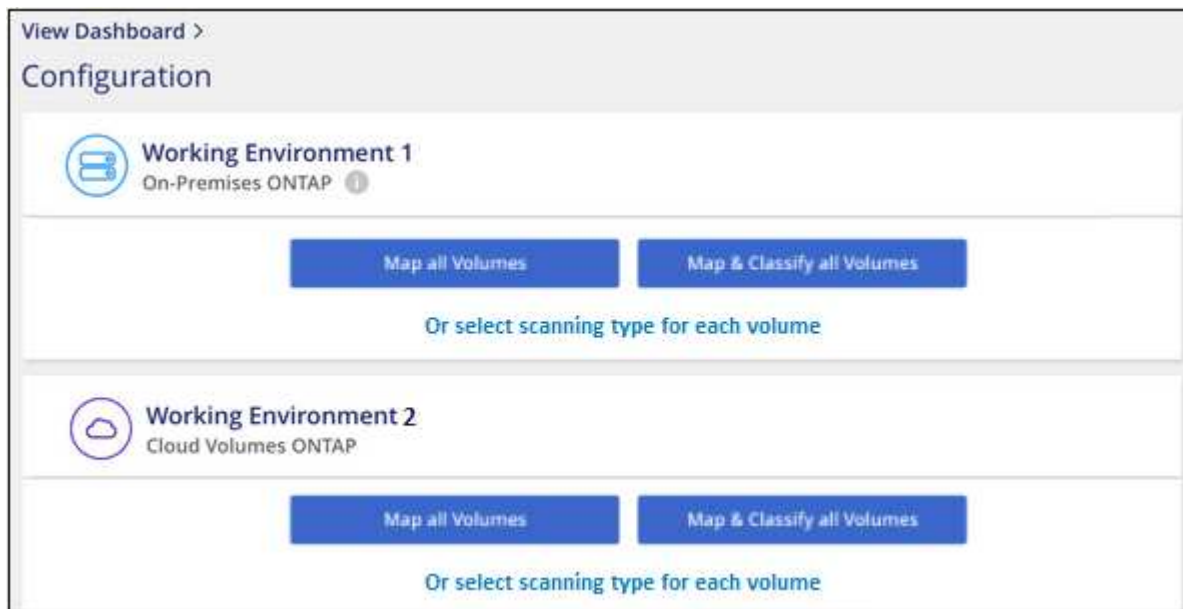
If you are scanning on-premises ONTAP systems that have been installed in a dark site that has no internet access, you need to [deploy BlueXP classification in the same on-premises location that has no internet access](#). This also requires that the BlueXP Connector is deployed in that same on-premises location.

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 on Cloud Volumes ONTAP systems in any supported cloud provider, and on on-premises ONTAP clusters.

1. From the BlueXP left navigation menu, click **Governance > Classification** and then select the **Configuration** tab.



2. 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.

See [Enable and disable compliance scans on volumes](#) for details.

3. 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.



- 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.

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 the following ports are open to the BlueXP classification instance:
 - For NFS - ports 111 and 2049.
 - For CIFS - ports 139 and 445.
4. Ensure that 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 left navigation menu, click **Governance > Classification** and then select the **Configuration** tab.




- b. For each working environment, click **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, we recommend that the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, we recommend making the Active Directory configured user 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, click **View Details** to review the status for each CIFS and NFS volume and correct any errors.

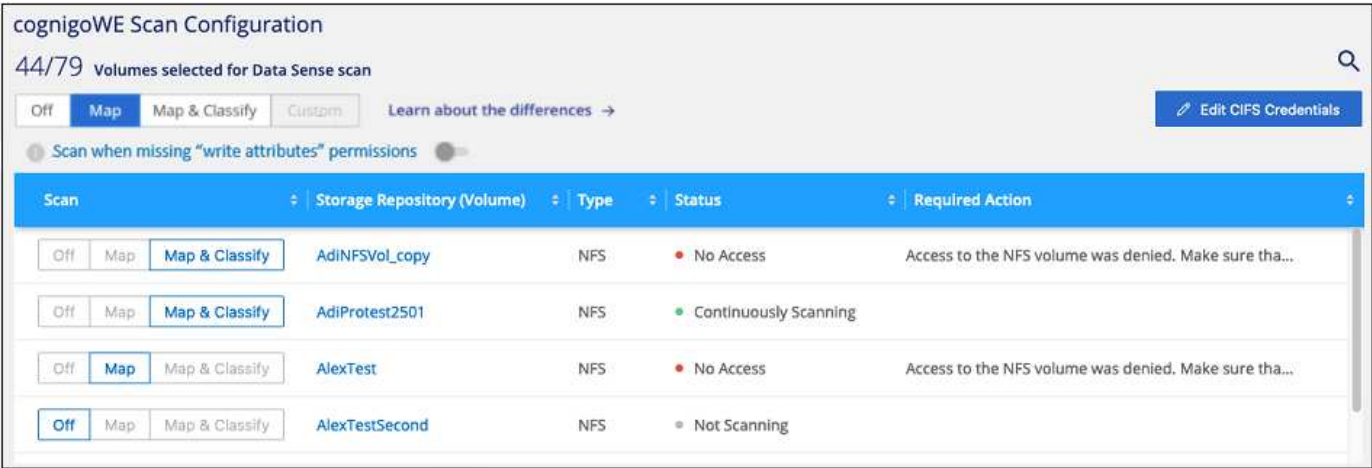
For example, the following image shows four volumes; one of which BlueXP classification can't scan due to network connectivity issues between the BlueXP classification instance and the volume.



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.

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.](#)



To:	Do this:
Enable mapping-only scans on a volume	In the volume area, click Map
Enable full scanning on a volume	In the volume area, click Map & Classify
Disable scanning on a volume	In the volume area, click Off
Enable mapping-only scans on all volumes	In the heading area, click Map
Enable full scanning on all volumes	In the heading area, click Map & Classify
Disable scanning on all volumes	In the heading area, click 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 on-premises ONTAP system or from a Cloud Volumes ONTAP 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*.

'Working Environment Name' Configuration 22/28 Volumes selected for compliance scan

Enable Access to DP Volumes
Edit CIFS Credentials

[Learn about the differences →](#)

☒ Scan when missing "write attributes" permissions

Scan	Storage Repository (Volume)	Type	Status	Required Action
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map & Classify"/>	VolumeName1	DP	Not Scanning	Enable access to DP Volumes ⓘ
<input type="button" value="Off"/> <input checked="" type="button" value="Map"/> <input type="button" value="Map & Classify"/>	VolumeName2	NFS	Continuously Scanning	
<input type="button" value="Off"/> <input type="button" value="Map"/> <input type="button" value="Map & Classify"/>	VolumeName3	CIFS	Not Scanning	

Steps

If you want to scan these data protection volumes:

1. Click **Enable Access to DP volumes** at the top of the page.
2. Review the confirmation message and click **Enable Access to DP volumes** again.
 - Volumes that were initially created as NFS volumes in the source ONTAP system are enabled.
 - Volumes that were initially created as CIFS volumes in the source ONTAP 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.

Provide Active Directory Credentials

☒ Use existing CIFS Scanning Credentials (user1@domain2)
 ☐ Use Custom Credentials

Active Directory Domain ⓘ
 DNS IP Address ⓘ

DP Volumes, created from a SnapMirror relationship, do not allow external access by default. Continuing will create NFS shares from DP Volumes which have been activated for **Data Sense**. The shares' export policies will allow access only from the Cloud **Data Sense** instance. [Learn More](#)

Enable Access to DP Volumes
Cancel

Provide Active Directory Credentials

☐ Use existing CIFS Scanning Credentials (user1@domain2)
 ☒ Use Custom Credentials

Username ⓘ
 Password

Active Directory Domain ⓘ
 DNS IP Address ⓘ

DP Volumes, created from a SnapMirror relationship, do not allow external access by default. Continuing will create NFS shares from DP Volumes which have been activated for **Data Sense**. The shares' export policies will allow access only from the Cloud **Data Sense** instance. [Learn More](#)

Enable Access to DP Volumes
Cancel

3. Activate each DP volume that you want to scan [the same way you enabled other volumes](#).

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.

Note: 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. Click this button and add CIFS credentials to enable access to these CIFS DP volumes.



Active Directory credentials are only registered 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 database schemas with BlueXP classification

Complete a few steps to start scanning your database schemas with BlueXP classification.

Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

Review database prerequisites

Ensure that your database is supported and that you have the information necessary to connect to the database.

2

Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

3

Add the database server

Add the database server that you want to access.

4

Select the schemas

Select the schemas that you want to scan.

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.

Upgrades to BlueXP classification software is automated as long as the instance has internet connectivity.

Add the database server

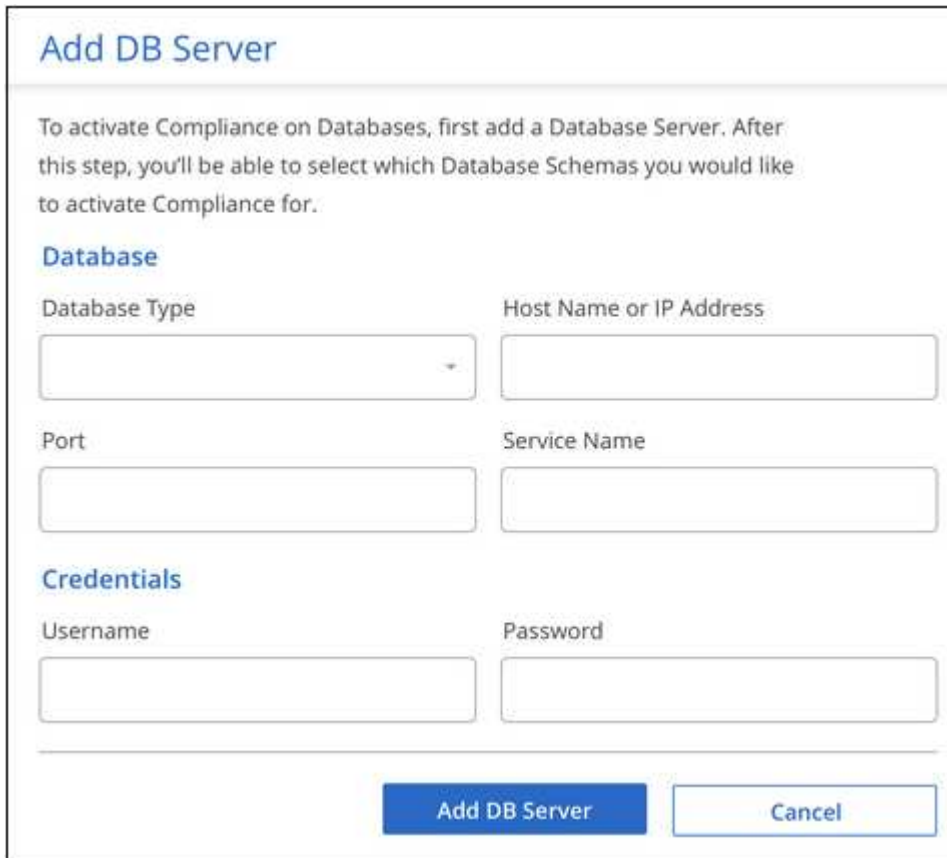
Add the database server where the schemas reside.

1. From the Working Environments Configuration page, click **Add Data Source > Add Database Server**.



2. Enter the required information to identify the database server.
 - a. Select the database type.
 - b. Enter the port and the host name or IP address to connect to the database.
 - c. For Oracle databases, enter the Service name.

- d. Enter the credentials so that BlueXP classification can access the server.
- e. Click **Add DB Server**.



Add DB Server

To activate Compliance on Databases, first add a Database Server. After this step, you'll be able to select which Database Schemas you would like to activate Compliance for.

Database

Database Type: Host Name or IP Address:

Port: Service Name:

Credentials

Username: Password:

Add DB Server **Cancel**

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, click 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.

'Working Environment Name' Configuration			
28/28 Schemas selected for compliance scan		<input type="text"/> Edit Credentials	
Scan	Schema Name	Status	Required Action
<input checked="" type="checkbox"/>	DB1 - SchemaName1	Not Scanning	Add Credentials ⓘ
<input checked="" type="checkbox"/>	DB1 - SchemaName2	Continuously Scanning	
<input checked="" type="checkbox"/>	DB1 - SchemaName3	Continuously Scanning	
<input checked="" type="checkbox"/>	DB1 - SchemaName4	Continuously Scanning	

Result

BlueXP classification starts scanning the database schemas that you enabled. If there are any errors, they'll appear in the Status column, alongside the required action to fix the error.

Note that BlueXP classification scans your databases once per day - databases are not continuously scanned like other data sources.

Scan file shares with BlueXP classification

Complete a few steps to start scanning NFS or CIFS file shares from Google Cloud NetApp Volumes and from older NetApp 7-mode systems. These file shares can reside on-premises or in the cloud.



Scanning data from non-NetApp file shares is not supported in the BlueXP classification core version.

Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

Review file share prerequisites

For CIFS (SMB) shares, ensure that you have credentials to access the shares.

2

Deploy the BlueXP classification instance

[Deploy BlueXP classification](#) if there isn't already an instance deployed.

3

Create a group to hold the file shares

The group is a container for the file shares that you want to scan, and it is used as the working environment name for those file shares.

4

Add the file shares to the group

Add the list of file shares that you want to scan and select the type of scanning. You can add up to 100 file shares at a time.

Review file share requirements

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.

Note that BlueXP classification can't extract permissions or the "last access time" from 7-Mode systems. Additionally, because of a known issue between some Linux versions and CIFS shares on 7-Mode systems, you must configure the share to use only SMB v1 with NTLM authentication enabled.

- There needs to be network connectivity between the BlueXP classification instance and the shares.
- Make sure these ports are open to the BlueXP classification instance:
 - For NFS – ports 111 and 2049.
 - For CIFS – ports 139 and 445.
- You can add a DFS (Distributed File System) share as a regular CIFS share. However, because BlueXP classification is not aware 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, we recommend that the user has Write Attributes permissions in CIFS or write permissions in NFS. If possible, we recommend making the Active Directory configured user part of a parent group in the organization which has permissions to all files.

- You will need the list of shares you want to add in the format `<host_name>:/<share_path>`. You can enter the shares individually, or you can supply a line-separated list of the file shares you want to scan.

Deploy the BlueXP classification instance

Deploy BlueXP classification if there isn't already an instance deployed.

Upgrades to BlueXP classification software are automated as long as the instance has internet connectivity.

Create the group for the file shares

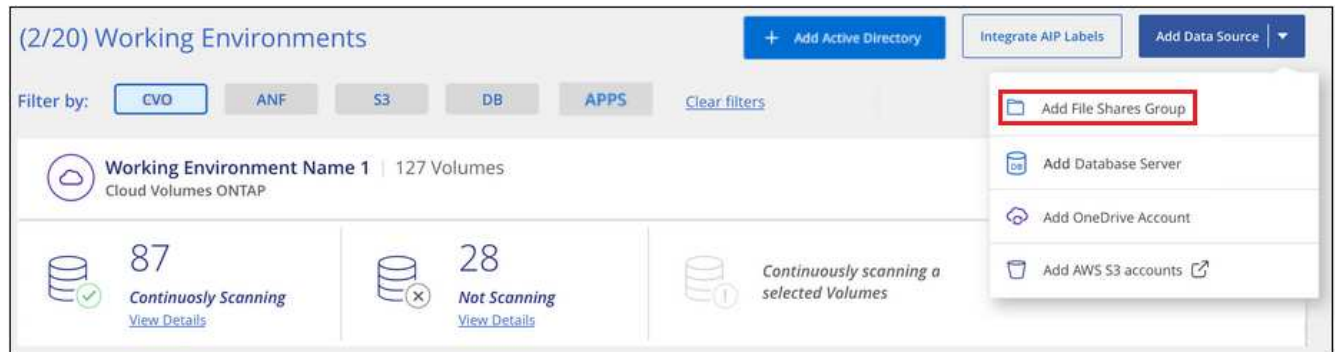
You must add a file shares "group" before you can add your file shares. The group is a container for the file shares that you want to scan, and the group name is used as the working environment name for those file shares.

You can mix NFS and CIFS shares in the same group, however, all CIFS file shares in a group need to be using the same Active Directory credentials. If you plan to add CIFS shares that use different credentials, you

must make a separate group for each unique set of credentials.

Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add File Shares Group**.



2. In the Add Files Shares Group dialog, enter the name for the group of shares and click **Continue**.

The new File Shares Group is added to the list of working environments.

Add file shares to a group

You add file shares to the File Shares Group so that the files in those shares will be scanned by BlueXP classification. You add the shares in the format `<host_name>:/<share_path>`.

You can add individual file shares, or you can supply a line-separated list of the file shares you want to scan. You can add up to 100 shares at a time.


When adding both NFS and CIFS shares in a single group, you'll need to run through the process twice - once adding NFS shares, and then again adding the CIFS shares.

Steps

1. From the *Working Environments* page, click the **Configuration** button for the File Shares Group.



2. If this is the first time adding file shares for this File Shares Group, click **Add your first Shares**.



If you are adding file shares to an existing group, click **Add Shares**.

SMB_Shares Scan Configuration

2 Shares selected for Data Sense scan

Scan when missing "write" permissions ☐

+ Add Shares Edit CIFS Credentials

Scan	Storage Repository (Share)	Protocol	Access	Scan Status	Required Action
Map Map & Classify	.E6	CIFS	Continuously Scanning	Mapped: 5.8K Classified: 5.8K	...
Map Map & Classify	7	CIFS	Continuously Scanning	Mapped: 5.8K Classified: 5.8K	...

3. Select the protocol for the file shares you are adding, add the file shares that you want to scan - one file share per line - and click **Continue**.

When adding CIFS (SMB) shares, you need to enter the Active Directory credentials that provide read access to the shares. Admin credentials are preferred.

Adding Shares

Directly add any NFS or CIFS (SMB) File Shares, located in the cloud or on-premises.

Select Protocol

You'll be able to add additional shares from the other protocol later.

☒ NFS ☐ CIFS (SMB)

Type or paste below the Shares to add

Provide a list of shares, line-separated. You can add up to 100 at a time (you can add more later).

Hostname/SHAREPATH
Hostname/SHAREPATH
Hostname/SHAREPATH

Continue Cancel

☐ NFS ☒ CIFS (SMB)

Provide CIFS Credentials

Username Password

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. In some cases you can re-add the share with a corrected host name or share name.

4. Enable mapping-only scans, or mapping and classification scans, on each file share.

To:	Do this:
Enable mapping-only scans on file shares	Click Map
Enable full scans on file shares	Click Map & Classify
Disable scanning on file shares	Click 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, 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](#).

Result

BlueXP classification starts scanning the files in the file shares you added, and the results are displayed in the Dashboard and in other locations.

Remove a file share from compliance scans

If you no longer need to scan certain file shares, you can remove individual file shares from having their files scanned at any time. Just click **Remove Share** from the Configuration page.



Scan StorageGRID data with BlueXP classification

Complete a few steps to start scanning data within StorageGRID directly with BlueXP classification.

Quick start

Get started quickly by following these steps, or scroll down to the remaining sections for full details.

1

Review StorageGRID prerequisites

You need to have the endpoint URL to connect with the StorageGRID service.

You need to have the Access Key and Secret Key from StorageGRID 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 StorageGRID Service

Add the StorageGRID 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.

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.

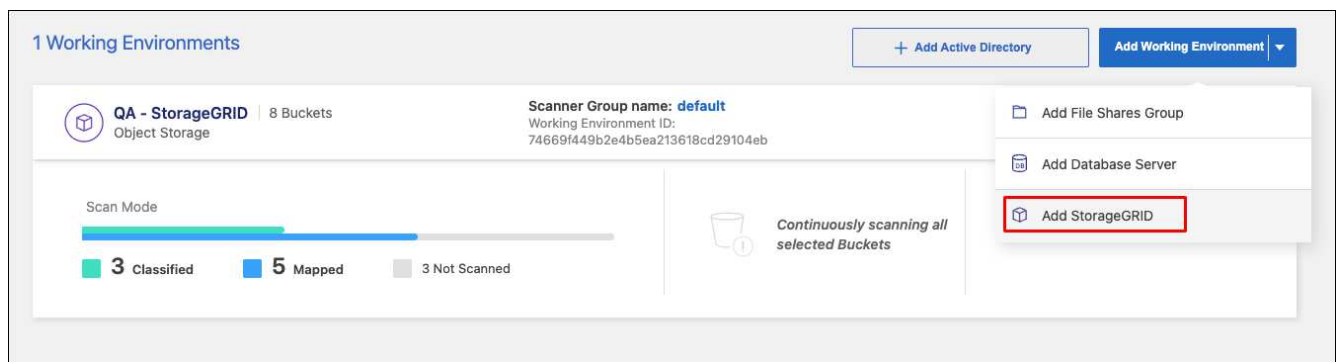
Upgrades to BlueXP classification software are automated as long as the instance has internet connectivity.

Add the StorageGRID service to BlueXP classification

Add the StorageGRID service.

Steps

1. From the Working Environments Configuration page, click **Add Data Source > Add StorageGRID**.



2. 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.

Add StorageGRID

BlueXP Classification can scan data from NetApp StorageGRID, which uses the S3 protocol. [Learn more](#)

To continue, provide the following details. Next, you'll select the buckets you want to scan.

Name the Working Environment

Endpoint URL

Access Key

Secret Key

Continue

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, click **Configuration** from the StorageGRID working environment.

1 Working Environments

[+ Add Active Directory](#)
[Add Working Environment](#)

QA - StorageGRID

8 Buckets

Object Storage

Scanner Group name: default

Working Environment ID: 74669f449b2e4b5ea213618cd29104eb

Configuration

Scan Mode

3 Classified

5 Mapped

3 Not Scanned

Continuously scanning all selected Buckets

2. Enable mapping-only scans, or mapping and classification scans, on your buckets.

Buckets selected for Classification scan (5/8)

Scan	Storage Repository (Bucket)	Mapping status	Classification status	Required Action
Off Map Map & Classify	bucketadipro	Finished 2024-09-05 10:33 Last full cycle: 2024-09-05 10:33	Mapped: 84 Classified: 5	...
Off Map Map & Classify	datasense-0-files	Finished 2024-09-05 08:00 Last full cycle: 2024-09-05 08:00		...
Off Map Map & Classify	datasense-10tb	Running 2024-09-04 07:25	Mapped: 3.7M Classified: 2.1M	...
Off Map Map & Classify	datasense-1tb	Running 2024-09-05 09:05 Last full cycle: 2024-09-05 03:04	Mapped: 1.3M	...
Off Map 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		...

To:	Do this:
Enable mapping-only scans on a bucket	Click Map
Enable full scans on a bucket	Click Map & Classify
Disable scanning on a bucket	Click Off

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.

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

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