



Reference

BlueXP backup and recovery

NetApp
April 18, 2024

Table of Contents

- Reference 1
 - AWS S3 archival storage classes and restore retrieval times 1
 - Azure archival tiers and restore retrieval times 2
 - Google archival storage classes and restore retrieval times 3
 - Configure backup for multi-account access in Azure 4
 - Restore BlueXP backup and recovery data in a dark site 11
 - Restart the BlueXP backup and recovery service 15

Reference

AWS S3 archival storage classes and restore retrieval times

BlueXP backup and recovery supports two S3 archival storage classes and most regions.

Supported S3 archival storage classes for BlueXP backup and recovery

When backup files are initially created they're stored in S3 *Standard* storage. This tier is optimized for storing data that's infrequently accessed; but that also allows you to access it immediately. After 30 days the backups transition to the S3 *Standard-Infrequent Access* storage class to save on costs.

If your source clusters are running ONTAP 9.10.1 or greater, you can choose to tier backups to either S3 *Glacier* or S3 *Glacier Deep Archive* storage after a certain number of days (typically more than 30 days) for further cost optimization. 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 about [restoring data from archival storage](#).

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

Note that when you configure BlueXP backup and recovery with this type of lifecycle rule, you must not configure any lifecycle rules when setting up the bucket in your AWS account.

[Learn about S3 storage classes.](#)

Restoring data from archival storage

While storing older backup files in archival storage is much less expensive than Standard or Standard-IA storage, accessing data from a backup file in archive storage for restore operations will take a longer amount of time and will cost more money.

How much does it cost to restore data from Amazon S3 Glacier and Amazon S3 Glacier Deep Archive?

There are 3 restore priorities you can choose when retrieving data from Amazon S3 Glacier, and 2 restore priorities when retrieving data from Amazon S3 Glacier Deep Archive. S3 Glacier Deep Archive costs less than S3 Glacier:

| Archive Tier | Restore Priority & Cost | | |
|--------------|---------------------------------|------------------------------|--------------------------------|
| | High | Standard | Low |
| S3 Glacier | Fastest retrieval, highest cost | Slower retrieval, lower cost | Slowest retrieval, lowest cost |

| Archive Tier | Restore Priority & Cost | | |
|--------------------------------|-------------------------|-------------------------------|-------------------------------|
| S3 Glacier Deep Archive | | Faster retrieval, higher cost | Slower retrieval, lowest cost |

Each method has a different per-GB retrieval fee and per-request fee. For detailed S3 Glacier pricing by AWS Region, visit the [Amazon S3 pricing page](#).

How long will it take to restore my objects archived in Amazon S3 Glacier?

There are 2 parts that make up the total restore time:

- **Retrieval time:** The time to retrieve the backup file from archive and place it in Standard storage. This is sometimes called the "rehydration" time. The retrieval time is different depending on the restore priority you choose.

| Archive Tier | Restore Priority & Retrieval Time | | |
|--------------------------------|-----------------------------------|-----------|------------|
| | High | Standard | Low |
| S3 Glacier | 3-5 minutes | 3-5 hours | 5-12 hours |
| S3 Glacier Deep Archive | | 12 hours | 48 hours |

- **Restore time:** The time to restore the data from the backup file in Standard storage. This time is no different than the typical restore operation directly from Standard storage - when not using an archival tier.

For more information about Amazon S3 Glacier and S3 Glacier Deep Archive retrieval options, refer to [the Amazon FAQ about these storage classes](#).

Azure archival tiers and restore retrieval times

BlueXP backup and recovery supports one Azure archival access tier and most regions.

Supported Azure Blob access tiers for BlueXP backup and recovery

When backup files are initially created they're stored in the *Cool* access tier. This tier is optimized for storing data that's infrequently accessed; but when needed, can be accessed immediately.

If your source clusters are running ONTAP 9.10.1 or greater, you can choose to tier backups from *Cool* to *Azure Archive* storage after a certain number of days (typically more than 30 days) for further cost optimization. Data in this tier can't be accessed immediately when needed, and will require a higher retrieval cost, so you need to consider how often you may need to restore data from these archived backup files. See the next section about [restoring data from archival storage](#).

Note that when you configure BlueXP backup and recovery with this type of lifecycle rule, you must not configure any lifecycle rules when setting up the container in your Azure account.

[Learn about Azure Blob access tiers](#).

Restoring data from archival storage

While storing older backup files in archival storage is much less expensive than Cool storage, accessing data from a backup file in Azure Archive for restore operations will take a longer amount of time and will cost more

money.

How much does it cost to restore data from Azure Archive?

There are two restore priorities you can choose when retrieving data from Azure Archive:

- **High:** Fastest retrieval, higher cost
- **Standard:** Slower retrieval, lower cost

Each method has a different per-GB retrieval fee and per-request fee. For detailed Azure Archive pricing by Azure Region, visit the [Azure pricing page](#).



The High priority is not supported when restoring data from Azure to StorageGRID systems.

How long will it take to restore my data archived in Azure Archive?

There are 2 parts that make up the restore time:

- **Retrieval time:** The time to retrieve the archived backup file from Azure Archive and place it in Cool storage. This is sometimes called the "rehydration" time. The retrieval time is different depending on the restore priority you choose:
 - **High:** < 1 hour
 - **Standard:** < 15 hours
- **Restore time:** The time to restore the data from the backup file in Cool storage. This time is no different than the typical restore operation directly from Cool storage - when not using an archival tier.

For more information about Azure Archive retrieval options, refer to [this Azure FAQ](#).

Google archival storage classes and restore retrieval times

BlueXP backup and recovery supports one Google archival storage class and most regions.

Supported Google archival storage classes for BlueXP backup and recovery

When backup files are initially created they're stored in *Standard* storage. This tier is optimized for storing data that's infrequently accessed; but that also allows you to access it immediately.

If your on-prem cluster is using ONTAP 9.12.1 or greater, you can choose to tier older backups to *Archive* storage in the BlueXP backup and recovery UI after a certain number of days (typically more than 30 days) for further cost optimization. Data in this tier will require a higher retrieval cost, so you need to consider how often you may need to restore data from these archived backup files. See the section about [restoring data from archival storage](#).

Note that when you configure BlueXP backup and recovery with this type of lifecycle rule, you must not configure any lifecycle rules when setting up the bucket in your Google account.

[Learn about Google storage classes.](#)

Restoring data from archival storage

While storing older backup files in Archive storage is much less expensive than Standard storage, accessing data from a backup file in Archive storage for restore operations will take a slightly longer amount of time and

will cost more money.

How much does it cost to restore data from Google Archive?

For detailed Google Cloud Storage pricing by region, visit the [Google Cloud Storage pricing page](#).

How long will it take to restore my objects archived in Google Archive?

There are 2 parts that make up the total restore time:

- **Retrieval time:** The time to retrieve the backup file from Archive and place it in Standard storage. This is sometimes called the "rehydration" time. Unlike the "coldest" storage solutions provided by other cloud providers, your data is accessible within milliseconds.
- **Restore time:** The time to restore the data from the backup file in Standard storage. This time is no different than the typical restore operation directly from Standard storage - when not using an archival tier.

Configure backup for multi-account access in Azure

BlueXP backup and recovery enables you to create backup files in an Azure account that is different than where your source Cloud Volumes ONTAP volumes reside. Both of those accounts can be different than the account where the BlueXP Connector resides.

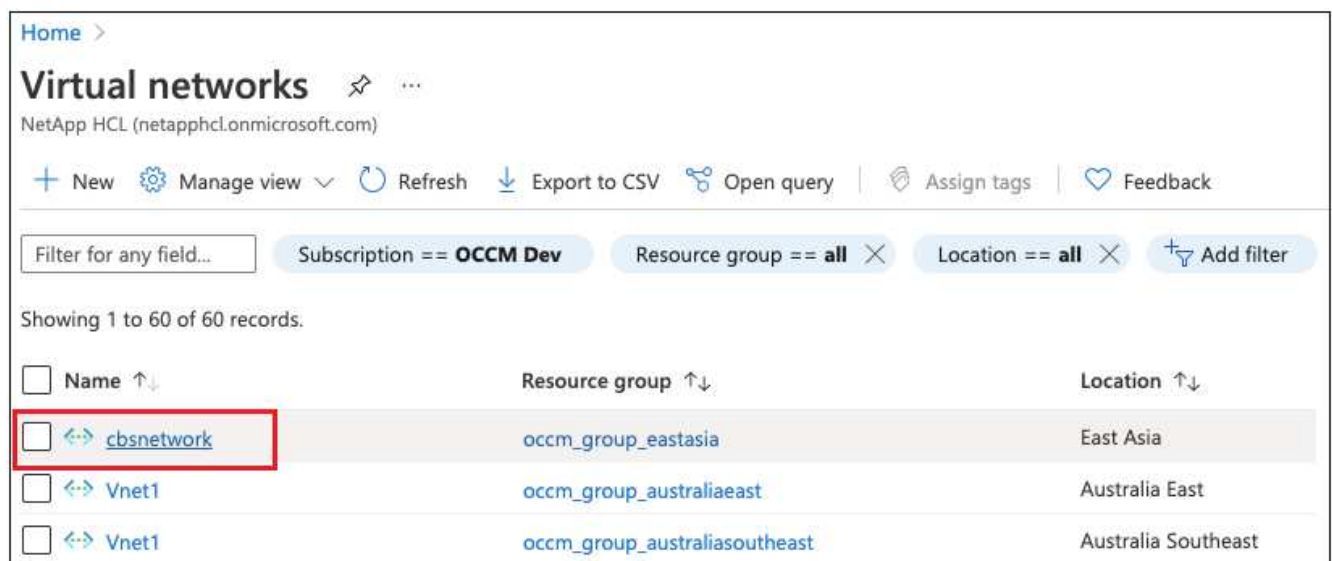
These steps are required only when you are [backing up Cloud Volumes ONTAP data to Azure Blob storage](#).

Just follow the steps below to set up your configuration in this manner.

Set up VNet peering between accounts

Note that if you want BlueXP to manage your Cloud Volumes ONTAP system in a different account/region, then you need to setup VNet peering. VNet peering is not required for storage account connectivity.

1. Log in to the Azure portal and from home, select Virtual Networks.
2. Select the subscription you are using as subscription 1 and click on the VNet where you want to set up peering.



3. Select **cbsnetwork** and from the left panel, click on **Peerings**, and then click **Add**.


The screenshot shows the 'Add peering' configuration page in the Azure portal. It contains the following sections:

- Subscription ***: A dropdown menu with 'OCCM Automation' selected.
- Virtual network ***: A dropdown menu with 'cbse2evnet' selected.
- Traffic to remote virtual network**: Two radio button options. 'Allow (default)' is selected, and 'Block all traffic to the remote virtual network' is unselected.
- Traffic forwarded from remote virtual network**: Two radio button options. 'Allow (default)' is selected, and 'Block traffic that originates from outside this virtual network' is unselected.
- Virtual network gateway or Route Server**: Three radio button options. 'None (default)' is selected, 'Use this virtual network's gateway or Route Server' is unselected, and 'Use the remote virtual network's gateway or Route Server' is unselected.

At the bottom of the form is a blue button labeled 'Add'.

4. Enter the following information on the Peering page and then click **Add**.

- Peering link name for this network: you can give any name to identify the peering connection.
- Remote virtual network peering link name: enter a name to identify the remote VNet.
- Keep all the selections as default values.
- Under subscription, select the subscription 2.
- Virtual network, select the virtual network in subscription 2 to which you want to set up the peering.


cbsnetwork | Peerings

Virtual network

«
+ Add
↻ Refresh

Overview

Activity log

Access control (IAM)

Tags

Diagnose and solve problems

Settings

Address space

Connected devices

Subnets

DDoS protection

Firewall

Security

DNS servers

Peerings

| Name | Peering status | Peer |
|------------|----------------|------------|
| cbsnetwork | Connected | cbse2evnet |

- Perform the same steps in subscription 2 VNet and specify the subscription and remote VNet details of subscription 1.

Subscription * ⓘ

OCCM Dev

Virtual network *

cbsnetwork

Traffic to remote virtual network ⓘ

☒ Allow (default)
 ☐ Block all traffic to the remote virtual network

Traffic forwarded from remote virtual network ⓘ


☒ Allow (default)
 ☐ Block traffic that originates from outside this virtual network

Virtual network gateway or Route Server ⓘ

☐ Use this virtual network's gateway or Route Server
 ☐ Use the remote virtual network's gateway or Route Server
 ☒ None (default)

Add

The peering settings are added.



Create a private endpoint for the storage account

Now you need to create a private endpoint for the storage account. In this example, the storage account is created in subscription 1 and the Cloud Volumes ONTAP system is running in subscription 2.




You need network contributor permission to perform the following action.

```

{
  "id": "/subscriptions/d333af45-0d07-4154-943dc25fbbce1b18/providers/Microsoft.Authorization/roleDefinitions/4d97b98b-1d4f-4787-a291-c67834d212e7",
  "properties": {
    "roleName": "Network Contributor",
    "description": "Lets you manage networks, but not access to them.",
    "assignableScopes": [
      "/"
    ],
    "permissions": [
      {
        "actions": [
          "Microsoft.Authorization/*/read",
          "Microsoft.Insights/alertRules/*",
          "Microsoft.Network/*",
          "Microsoft.ResourceHealth/availabilityStatuses/read",
          "Microsoft.Resources/deployments/*",
          "Microsoft.Resources/subscriptions/resourceGroups/read",
          "Microsoft.Support/*"
        ],
        "notActions": [],
        "dataActions": [],
        "notDataActions": []
      }
    ]
  }
}

```

1. Go to the Storage account > Networking > Private endpoint connections and click + **Private endpoint**.



2. In the Private Endpoint *Basics* page:

- Select subscription 2 (where the BlueXP Connector and Cloud Volumes ONTAP system are deployed) and the resource group.
- Enter an endpoint name.
- Select the region.

Create a private endpoint

1 Basics 2 Resource 3 Configuration 4 Tags 5 Review + create

Use private endpoints to privately connect to a service or resource. Your private endpoint must be in the same region as your virtual network, but can be in a different region from the private link resource that you are connecting to. [Learn more](#)

Project details

Subscription * ⓘ OCCM Dev

Resource group * ⓘ cbsoccmdevcvo-rg [Create new](#)

Instance details

Name * cbse2e ✓

Region * (Asia Pacific) East Asia

3. In the *Resource* page, select Target sub-resource as **blob**.

Create a private endpoint ...

✓ Basics **2 Resource** 3 Configuration 4 Tags 5 Review + create

Private Link offers options to create private endpoints for different Azure resources, like your private link service, a SQL server, or an Azure storage account. Select which resource you would like to connect to using this private endpoint. [Learn more](#)

Subscription OCCM Dev (d333af45-0d07-4154-943d-c25fbbce1b18)

Resource type Microsoft.Storage/storageAccounts

Resource test150521

Target sub-resource * ⓘ

4. In the Configuration page:

- Select the virtual network and subnet.
- Click the **Yes** radio button to "Integrate with private DNS zone".

Create a private endpoint ...

✓ Basics ✓ Resource **3 Configuration** 4 Tags 5 Review + create

Networking

To deploy the private endpoint, select a virtual network subnet. [Learn more](#)

Virtual network * ⓘ

Subnet * ⓘ

ⓘ If you have a network security group (NSG) enabled for the subnet above, it will be disabled for private endpoints on this subnet only. Other resources on the subnet will still have NSG enforcement.

Private DNS integration

To connect privately with your private endpoint, you need a DNS record. We recommend that you integrate your private endpoint with a private DNS zone. You can also utilize your own DNS servers or create DNS records using the host files on your virtual machines. [Learn more](#)

Integrate with private DNS zone ☒ Yes ☐ No

| Configuration name | Subscription | Private DNS zone |
|---------------------------|--------------|-----------------------------------|
| privatelink-blob-core-... | OCCM Dev | privatelink.blob.core.windows.net |

Review + create < Previous Next: Tags >

5. In the Private DNS zone list, ensure that the Private Zone is selected from the correct Region, and click **Review + Create**.

| Configuration name | Subscription | Private DNS zone |
|---------------------------|--------------|---|
| privatelink-blob-core-... | OCCM Dev | privatelink.blob.core.windows.net |
| | | <input type="text" value="Filter private DNS zones"/> <ul style="list-style-type: none"> occm_group_centralus privatelink.blob.core.windows.net occm_group_eastus privatelink.blob.core.windows.net occm_group_eastus2 privatelink.blob.core.windows.net |

Now the storage account (in subscription 1) has access to the Cloud Volumes ONTAP system which is running in subscription 2.

6. Retry enabling BlueXP backup and recovery on the Cloud Volumes ONTAP system and this time it should be successful.

Restore BlueXP backup and recovery data in a dark site

When using BlueXP backup and recovery in a site with no internet access, known as *private mode*, the BlueXP backup and recovery configuration data is backed up to the StorageGRID or ONTAP S3 bucket where your backups are being stored. If you have an issue with the BlueXP Connector host system in the future, you can deploy a new Connector and restore the critical BlueXP backup and recovery data.

Note that when you use BlueXP backup and recovery in a SaaS environment where the BlueXP Connector is deployed at your cloud provider, or on your own host system that has internet access, all the important BlueXP backup and recovery configuration data is backed up and protected in the cloud. If you have an issue with the Connector, just create a new Connector and add your working environments and the backup details are automatically restored.

There are 2 types of data that are backed up:

- BlueXP backup and recovery database - contains a listing of all the volumes, backup files, backup policies, and configuration information.
- Indexed Catalog files - contains detailed indexes that are used for Search & Restore functionality that make your searches very quick and efficient when looking for volume data that you want to restore.

This data is backed up once per day at midnight, and a maximum of 7 copies of each file are retained. If the Connector is managing multiple on-premises ONTAP working environments, the BlueXP backup and recovery files will be located in the bucket of the working environment that was activated first.



No volume data is ever included in the BlueXP backup and recovery database or Indexed Catalog files.

Restore BlueXP backup and recovery data to a new Connector

If your on-premises Connector has a catastrophic failure, you'll need to install a new Connector, and then restore the BlueXP backup and recovery data to the new Connector.

There are 4 tasks you'll need to perform to return your BlueXP backup and recovery system to a working state:

- Install a new BlueXP Connector
- Restore the BlueXP backup and recovery database
- Restore the Indexed Catalog files
- Rediscover all of your on-prem ONTAP systems and StorageGRID systems to the BlueXP UI

Once you verify that your system is back in a working order, we recommend that you create new backup files.

What you'll need

You'll need to access the most recent database and index backups from the StorageGRID or ONTAP S3 bucket where your backup files are being stored:

- BlueXP backup and recovery MySQL database file

This file is located in the following location in the bucket `netapp-backup-
<GUID>/mysql_backup/`, and it is named `CBS_DB_Backup_<day>_<month>_<year>.sql`.

- Indexed Catalog backup zip file

This file is located in the following location in the bucket `netapp-backup-
<GUID>/catalog_backup/`, and it is named `Indexed_Catalog_DB_Backup_<db_name>_<day>_<month>_<year>.zip`.

Install a new Connector on a new on-premises Linux host

When installing a new BlueXP Connector, make sure you download the same release of software as you had installed on the original Connector. Periodic changes to the BlueXP backup and recovery database structure may make newer software releases incompatible with the original database backups. You can [upgrade the Connector software to the most current version after restoring the Backup database](#).

1. [Install the BlueXP Connector on a new on-premises Linux host](#)
2. Log into BlueXP using the admin user credentials that you just created.

Restore the BlueXP backup and recovery database

1. Copy the MySQL backup from the backup location to the new Connector host. We'll use the example file name "CBS_DB_Backup_23_05_2023.sql" below.
2. Copy the backup into the MySQL docker container using the following command:

```
docker cp CBS_DB_Backup_23_05_2023.sql ds_mysql_1:/. 
```

3. Enter the MySQL container shell using the following command:

```
docker exec -it ds_mysql_1 sh
```

4. In the container shell, deploy the "env".
5. You'll need the MySQL DB password, so copy the value of the key "MYSQL_ROOT_PASSWORD".
6. Restore the BlueXP backup and recovery MySQL DB using the following command:

```
mysql -u root -p cloud_backup < CBS_DB_Backup_23_05_2023.sql
```

7. Verify that the BlueXP backup and recovery MySQL DB has been restored correctly using the following SQL commands:

```
mysql -u root -p cloud_backup
```

Enter the password.

```
mysql> show tables;  
mysql> select * from volume;
```

Check if the volumes that are shown are the same as those that existed in your original environment.

Restore the Indexed Catalog files

1. Copy the Indexed Catalog backup zip file (we'll use the example file name "Indexed_Catalog_DB_Backup_catalogdb1_23_05_2023.zip") from the backup location to the new Connector host in the "/opt/application/netapp/cbs" folder.
2. Unzip the "Indexed_Catalog_DB_Backup_catalogdb1_23_05_2023.zip" file using the following command:

```
unzip Indexed_Catalog_DB_Backup_catalogdb1_23_05_2023.zip -d catalogdb1
```

3. Run the **ls** command to make sure that the folder "catalogdb1" has been created with the subfolders "changes" and "snapshots" underneath.

Discover your ONTAP clusters and StorageGRID systems

1. [Discover all the on-prem ONTAP working environments](#) that were available in your previous environment. This includes the ONTAP system you have used as an S3 server.
2. [Discover your StorageGRID systems](#).

Set up the StorageGRID environment details

Add the details of the StorageGRID system associated with your ONTAP working environments as they were set up on the original Connector setup using the [BlueXP APIs](#).

You'll need to perform these steps for each ONTAP system that is backing up data to StorageGRID.

1. Extract the authorization token using the following oauth/token API.

```
curl 'http://10.193.192.202/oauth/token' -X POST -H 'User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10.15; rv:100101 Firefox/108.0' -H 'Accept: application/json' -H 'Accept-Language: en-US,en;q=0.5' -H 'Accept-Encoding: gzip, deflate' -H 'Content-Type: application/json' -d '{"username":admin@netapp.com,"password":"Netapp@123","grant_type":"password"}'> '
```

This API will return a response like the following. You can retrieve the authorization token as shown below.

```
{"expires_in":21600,"access_token":"eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjZlIn0eyJzdWIiOiJvY2NtYXV0aHwxiwiYXVkJjpbImh0dHBzOi8vYXBpLmNsb3VkLm5ldGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm5ldGFwcC5jb20vZnVsbF9uYW1lIjoiYWRTaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS9lbWFPbCI6ImFkbWluQG5ldGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwcm9maWx1IiwiaWF0IjoxNjcyNzMDIzLCJleHAiOiE2NzI3NTc2MjMsImZlcyI6Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9CJtRpRDY23PokyLgl1f67bmgnMcYxdCvBOY-ZUYWzhrWbbY_hqUH4T-114v_pNDsPyNDyWqHaKizThdjyHYHxm56vTz_Vdn4NqjaBDPwN9KAnC6Z88WA1cJ4WRQqj5ykODNDmrv5At_f9HHp0-xVMYHqywZ4nNFAlMvAh4xEsc5jfoKOZc-IOQdWm4F4LHpMzs4qFzCYthTuSKLYtqSTUrZB81-o-ipvrOqSoliwIeHXZJJV-UsWun9daNgiYd_wX-4WWJViGEnDzzwOKfUoUoe1Fg3ch--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 eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjZlIn0eyJzdWIiOiJvY2NtYXV0aHwxiwiYXVkJjpbImh0dHBzOi8vYXBpLmNsb3VkLm5ldGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm5ldGFwcC5jb20vZnVsbF9uYW1lIjoiYWRTaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS9lbWFPbCI6ImFkbWluQG5ldGFwcC5jb20iLCJzY29wZSI6Im9wZW5pZCBwcm9maWx1IiwiaWF0IjoxNjcyNzMDIzLCJleHAiOiE2NzI3NDQzMTMsImZlcyI6Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdu_kN-fLWpdJJX98HODwPpVUitLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-ye0fH9gr4XgkdsWjWcNvw2rRkFzjHpWrETgfgAMkZcAukV4DHuxogHWh6-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 eyJhbGciOiJSUzI1NiIsInR5cCI6IkpXVCIsImtpZCI6IjJlMGFiZjRiIn0eyJzdWIiOiJvY2NtYXV0aHwxIiwiaXVkaW9uImh0dHBzOi8vYXBpLmNsb3VkLm5ldGFwcC5jb20iXSwiaHR0cDovL2Nsb3VkLm5ldGFwcC5jb20vZnVsbF9uYW11IjoieWRtaW4iLCJodHRwOi8vY2xvdWQubmV0YXBwLmNvbS9lbWVpY29wZSI6Im9wZW5pZCBwcmaWx1IiwiaWF0IjoxNjcyNzIyNzEzNDQzMjM5Imh0dHA6Ly9vY2NtYXV0aDo4NDIwLyJ9X_cQF8xttD0-S7sU2uph2cdu_kN-fLWpdJJX98HODwPpVUitLcxV28_sQhuopjWobozPelNISf7KvMqcoXc5kLDyX-yE0fH9gr4XgkdsWjWcNvw2rRkFzjHpWrETgfqAMkZcAukV4DHuxogHWh6-DggB1NgPZT8A_szHinud5W0HJ9c4AaT0zC-sp81GaqMahPf0KcFVybBL4krOewgKHGfo_7ma_4mF39B1LCj7Vc2XvUd0wCaJvDMjwp19-KbZqmmBX9vDnYp7SSxC1hHJRdstcFgJLdJHtowweNH2829KsjEGBTtcBd08SvIDtctNH_GAxwSgMT3zUfwaOimPw' \
> --header 'x-agent-id: vB_1xShPpBtUosjD7wfBlLIhqDgIPA0wclients' \
> -d '{ "storage-server" : "sr630ip15.rtp.eng.netapp.com:10443", "access-key": "2ZMYOAVAS5E70MCNH9", "secret-password": "uk/6ikd4LjlXQOFnzSzP/T0zR4ZQlG0w1xgWsB" }'
```

Verify BlueXP backup and recovery settings

1. Select each ONTAP working environment and click **View Backups** next to the Backup and recovery service in the right-panel.

You should be able to see all the backups that have been created for your volumes.

2. From the Restore Dashboard, under the Search & Restore section, click **Indexing Settings**.

Make sure that the working environments which had Indexed Cataloging enabled previously remain enabled.

3. From the Search & Restore page, run a few catalog searches to confirm that the Indexed Catalog restore has been completed successfully.

Restart the BlueXP backup and recovery service

There may be situations where you'll need to restart the BlueXP backup and recovery service.

BlueXP backup and recovery functionality is built into the BlueXP Connector. You'll need to follow different initial steps to restart the service depending on whether you deployed the Connector in the cloud or whether you installed the Connector manually on a Linux system.

Steps

1. Connect to the Linux system that the Connector is running on.

| Connector location | Procedure |
|---------------------|---|
| Cloud deployment | Follow the instructions for connecting to the Connector Linux virtual machine depending on the cloud provider you're using. |
| Manual installation | Log in to the Linux system. |

2. Enter the command to restart the service.

| Connector location | Command |
|---|---|
| Cloud deployment | <code>docker restart cloudmanager_cbs</code> |
| Manual installation with internet access | <code>docker restart cloudmanager_cbs</code> |
| Manual installation without internet access | <code>docker restart ds_cloudmanager_cbs_1</code> |

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.