



Back up and restore applications data

Cloud Backup

NetApp
December 07, 2022

Table of Contents

- Back up and restore applications data 1
 - Back up and restore on-premises applications data 1
 - Back up and restore cloud native applications data 14

Back up and restore applications data

Back up and restore on-premises applications data

Protect your on-premises applications data

You can integrate Cloud Backup for Applications with BlueXP (formerly Cloud Manager) and on-premises SnapCenter to back up the application consistent Snapshots from on-premises ONTAP to cloud. When required you can restore from cloud to on-premises SnapCenter Server.

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



You should be using SnapCenter Software 4.6 or later.

For more information about Cloud Backup for Applications, refer to:

- [Application aware backup with Cloud Backup and SnapCenter](#)
- [Cloud Backup for Applications podcast](#)

Requirements

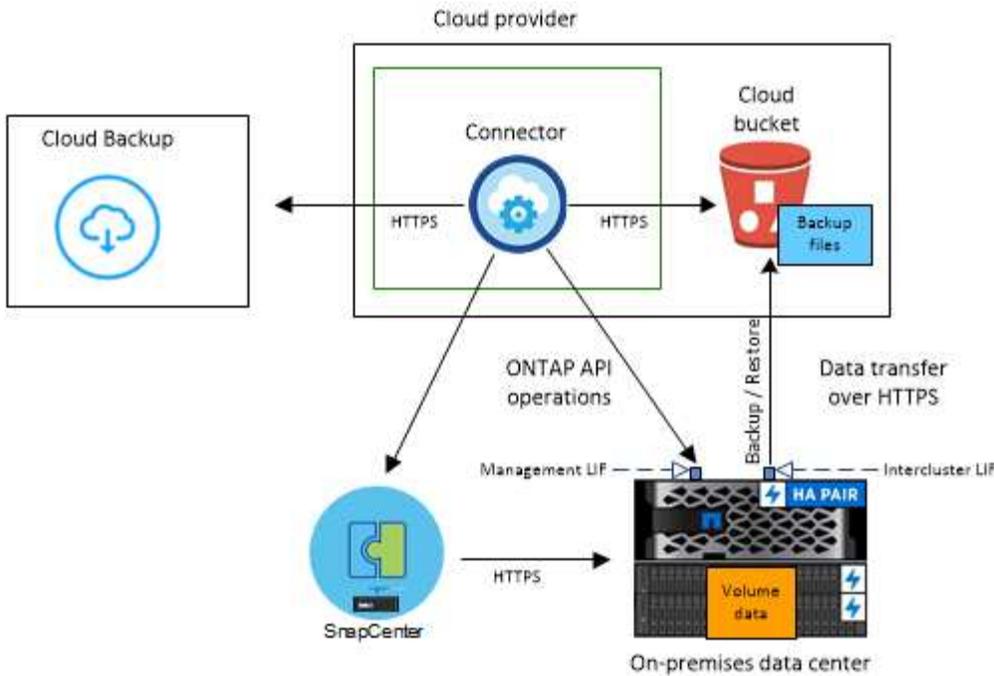
Read the following requirements to make sure that you have a supported configuration before you start backing up application data to cloud services.

- ONTAP 9.8 or later
- BlueXP 3.9
- SnapCenter Server 4.6 or later
You should be using SnapCenter Server 4.7 if you want to use the following features:
 - protect backups from on-premises secondary storage
 - protect SAP HANA applications
 - protect Oracle and SQL applications that are on VMware environment
 - mount backups
 - deactivate backups
 - unregister SnapCenter Server
- At least one backup per application should be available in SnapCenter Server
- At least one daily, weekly, or monthly policy in SnapCenter with no label or same label as that of the Cloud Backup for Applications policy in BlueXP.

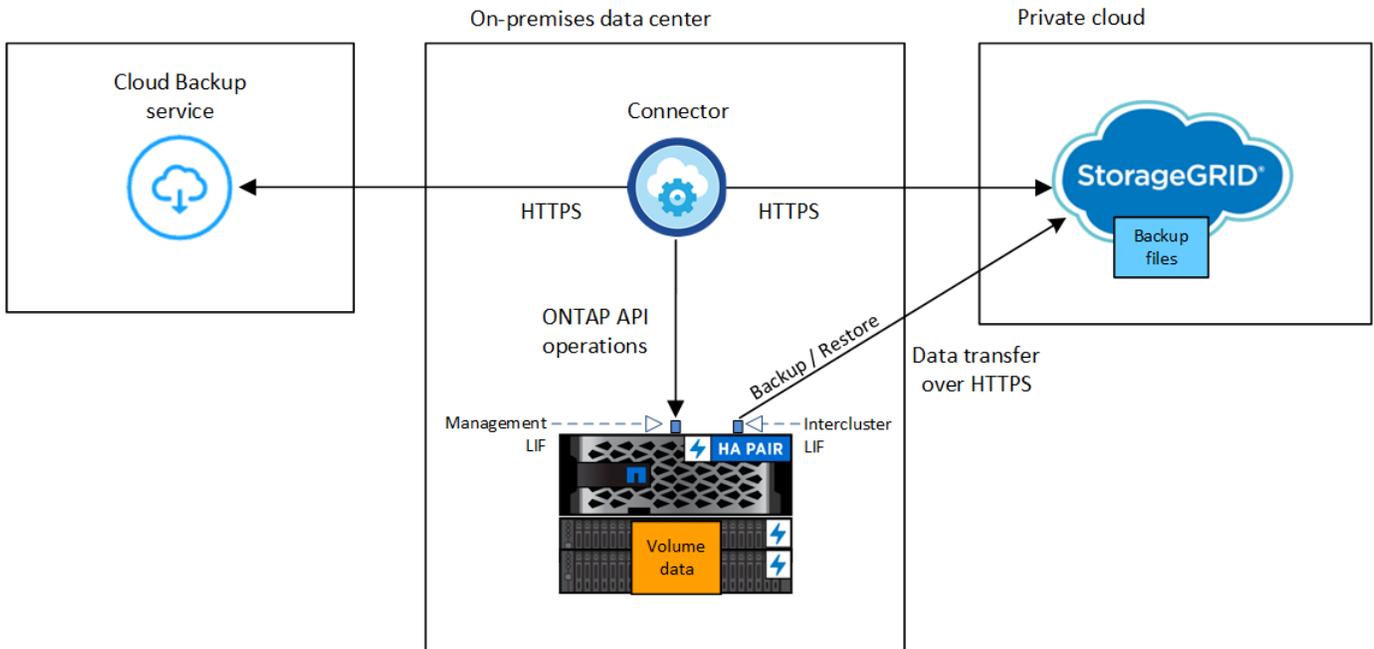


Cloud Backup for Applications does not support protection of applications that are on SVMs which were added using FQDN or IP address.

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



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



Register SnapCenter Server

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

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. From the **Settings** drop-down, click **SnapCenter Servers**.

3. Click **Register SnapCenter Server**.

4. Specify the following details:

- a. In the SnapCenter Server field, specify the FQDN or IP address of the SnapCenter Server host.
- b. In the Port field, specify the port number on which the SnapCenter Server is running.

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

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

The tags are comma separated.

- d. In the Username and Password field, specify the credentials of the user with SnapCenterAdmin role.

5. Click **Register**.

After you finish

Click **Backup & Restore > Applications** to view all the applications that are protected using the registered SnapCenter Server host.

By default, the applications are automatically discovered every day midnight. You can configure the schedule to discover the applications.



For SQL Server databases, the Application Name column displays the name in *application_name (instance name)* format.

The supported applications and their configurations are:

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

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

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

The following databases will not be displayed:

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

Create a policy to back up applications

You can either use one of the pre-canned policies or create a custom policy to back up the application data to cloud. You can create policies if you do not want to edit the pre-canned policies.

The pre-canned policies are:

Policy Name	Label	Retention Value
1 Year Daily LTR	Daily	366
5 Years Daily LTR	Daily	1830
7 Year Weekly LTR	Weekly	370
10 Year Monthly LTR	Monthly	120

Steps

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

You can move backups from object store to archival storage only if you are using ONTAP 9.10.1 or later and Amazon Web Services or Azure as the cloud provider. You should configure the archive access tier for each cloud provider.

7. Click **Create**.

You can edit, copy, and delete the customized policies.



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

Back up on-premises applications data to Amazon Web Services

You can back up the applications data from ONTAP to Amazon Web Services by integrating Cloud Backup for Applications with BlueXP and on-premises SnapCenter.

You can protect one or more applications simultaneously to the cloud using a single policy.



You can protect only one application at a time if you are using the BlueXP GUI. However, if you are using REST APIs, you can protect multiple applications simultaneously.

Steps

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

Configure the ONTAP cluster that hosts the SVM on which the application is running. After adding the working environment for one of the applications, it can be reused for all the other applications residing on the same ONTAP cluster.

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

Cloud Backup for Applications supports only cluster admin.

- c. Click **Add Working Environment**.
5. Select **Amazon Web Services** as the cloud provider.
 - a. Specify the AWS account.
 - b. In the AWS Access Key field, specify the key.
 - c. In the AWS Secret Key field, specify the password.
 - d. Select the region where you want to create the backups.
 - e. Specify the IP space.
 - f. Select the archive tier.

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

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

Back up on-premises applications data to Microsoft Azure

You can back up the applications data from ONTAP to Microsoft Azure by integrating Cloud Backup for Applications with BlueXP and on-premises SnapCenter.

You can protect one or more applications simultaneously to the cloud using a single policy.



You can protect only one application at a time if you are using the BlueXP GUI. However, if you are using REST APIs, you can protect multiple applications simultaneously.

Steps

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

Configure the ONTAP cluster that hosts the SVM on which the application is running. After adding the working environment for one of the applications, it can be reused for all the other applications residing on the same ONTAP cluster.

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

Cloud Backup for Applications supports only cluster admin.

- c. Click **Add Working Environment**.
5. Select **Microsoft Azure** as the cloud provider.
 - a. Specify the Azure subscription ID.
 - b. Select the region where you want to create the backups.
 - c. Either create a new resource group or use an existing resource group.
 - d. Specify the IP space.
 - e. Select the archive 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 click **Activate Backup**.

Back up on-premises applications data to Google Cloud Platform

You can back up the applications data from ONTAP to Google Cloud Platform by integrating Cloud Backup for Applications with BlueXP and on-premises SnapCenter.

You can protect one or more applications simultaneously to the cloud using a single policy.



You can protect only one application at a time if you are using the BlueXP GUI. However, if you are using REST APIs, you can protect multiple applications simultaneously.

Steps

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

Configure the ONTAP cluster that hosts the SVM on which the application is running. After adding the working environment for one of the applications, it can be reused for all the other applications residing on the same ONTAP cluster.

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

Cloud Backup for Applications supports only cluster admin.

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

5. Select **Google Cloud Platform** as the cloud provider.

- a. Select the Google Cloud Project where you want the Google Cloud Storage bucket to be created for backups..
- b. In the Google Cloud Access Key field, specify the key.
- c. In the Google Cloud Secret Key field, specify the password.
- d. Select the region where you want to create the backups.
- e. Specify the IP space.

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

Back up on-premises applications data to StorageGRID

You can back up the applications data from ONTAP to StorageGRID by integrating Cloud Backup for Applications with BlueXP and on-premises SnapCenter.

You can protect one or more applications simultaneously to StorageGRID using a single policy.



You can protect only one application at a time if you are using the BlueXP GUI. However, if you are using REST APIs, you can protect multiple applications simultaneously.

What you will need

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

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

Steps

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

Configure the ONTAP cluster that hosts the SVM on which the application is running. After adding the working environment for one of the applications, it can be reused for all the other applications residing on

the same ONTAP cluster.

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

Cloud Backup for Applications supports only cluster admin.

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

5. Select **StorageGRID**.

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

Enter the details in the format FQDN:PORT.

- b. In the Access Key field, specify the key.
- c. In the Secret Key field, specify the password.
- d. Specify the IP space.

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

Manage protection of applications

You can manage protection of applications by performing different operations from the BlueXP UI.

View policies

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

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

The associated applications are listed.



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

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

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running `Get-Help command_name`. Alternatively, you can also refer the [SnapCenter Software Cmdlet Reference Guide](#).

View backups on cloud

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

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



The time taken for the backups to be listed depends on ONTAP's default replication schedule (maximum of 1 hour) and BlueXP (maximum of 6 hours).

- For Oracle databases, both data and log backups, SCN number for each backup, end date for each backup are listed. You can select only the data backup and restore the database to on-premises SnapCenter Server.
- For Microsoft SQL Server databases, only the full backups and the end date for each backup are listed. You can select the backup and restore the database to on-premises SnapCenter Server.
- For Microsoft SQL Server instance, backups are not listed instead only the databases under that instance is listed.
- For SAP HANA databases, only the data backups and the end date for each backup are listed. You can select the backup and perform mount operation.



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

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

The information regarding the parameters that can be used with the cmdlet and their descriptions can be obtained by running `Get-Help command_name`. Alternatively, you can also refer the [SnapCenter Software Cmdlet Reference Guide](#).

Database layout change

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

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

The new volumes are discovered.

4. Click **...** corresponding to the application and click **Refresh Protection** to enable cloud protection for the new volume.

If a storage volume is removed from the application after configuring the cloud service, for new backups SnapCenter Server will only label the snapshots on which the application is residing. If the removed volume is not used by any other applications, then you should manually delete the object store relationship. If you update the application inventory, it will contain the current storage layout of the application.

Policy or resource group change

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

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

Unregister SnapCenter Server

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

Monitor Jobs

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

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

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

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

Set IP space of the primary working environment

If you want to restore or mount a backup that was moved to object store from secondary storage, you should add the primary working environment details and set the IP space.

Steps

1. In BlueXP UI, click **Storage** > **Canvas** > **My Working Environments** > **Add Working Environment**.
2. Specify the primary working environment details and click **Add**.
3. Click **Backup and recovery** > **Volumes**.
4. Click **...** corresponding to any of the volumes and click **Details**.
5. Click **...** corresponding to the backup and click **Restore**.
6. In the wizard, select the newly added primary working environment as the destination.
7. Specify the IP space.

Configure CA Certificates

If you have CA certificates, you should manually copy the root CA certificates to the connector machine.

However, if you do not have CA certificates, you can proceed without configuring CA certificates.

Steps

1. Copy the certificate to the volume that can be accessed from the docker agent.

```
° cd /var/lib/docker/volumes/cloudmanager_snapcenter_volume/_data/mkdir
  sc_certs
° chmod 777 sc_certs
```

2. Copy the RootCA certificate files to the above folder on the connector machine.

```
cp <path on connector>/<filename>
/var/lib/docker/volumes/cloudmanager_snapcenter_volume/_data/sc_certs
```

3. Copy the CRL file to the volume which can be accessed from the docker agent.

- `cd /var/lib/docker/volumes/cloudmanager_snapcenter_volume/_data/mkdir sc_crl`
- `chmod 777 sc_crl`

4. Copy the CRL files to the above folder on the connector machine.

```
cp <path on connector>/<filename>  
/var/lib/docker/volumes/cloudmanager_snapcenter_volume/_data/sc_crl
```

5. After copying the certificates and CRL files, restart the Cloud Backup for Apps service.

- `sudo docker exec cloudmanager_snapcenter sed -i 's/skipSCCertValidation: true/skipSCCertValidation: false/g' /opt/netapp/cloudmanager-snapcenter-agent/config/config.yml`
- `sudo docker restart cloudmanager_snapcenter`

Restore applications data

Restore Oracle database

You can only restore the Oracle database to the same SnapCenter Server host, same SVM, or to the same database host. For a RAC database, the data will be restored to the on-premises node where the backup was created.



Restore of secondary backups through primary is supported.

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



Single File Restore (SFR) is not supported.

What you will need

If you want to restore a backup that was moved to object store from secondary storage, you should add the primary working environment details and set the IP space. For more information, see [Set IP space of the primary working environment](#).

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. In the **Filter By** field, select the filter **Type** and from the drop-down select **Oracle**.
3. Click **View Details** corresponding to the database that you want to restore and click **Restore**.
4. On the Restore Type page, perform the following actions:
 - a. Select **Database State** if you want to change the state of the database to the state required to perform restore and recovery operations.

The various states of a database from higher to lower are open, mounted, started, and shutdown. You must select this check box if the database is in a higher state but the state must be changed to a lower state to perform a restore operation. If the database is in a lower state but the state must be changed to a higher state to perform the restore operation, the database state is changed automatically even if you

do not select the check box.

If a database is in the open state, and for restore the database needs to be in the mounted state, then the database state is changed only if you select this check box.

- b. Select **Control files** if you want to restore control file along with full database.
 - c. If the snapshot is in archival storage, specify the priority to restore your data from the archival storage.
5. On the Recovery Scope page, perform the following actions:
- a. Specify the recovery scope.

If you...	Do this...
Want to recover to the last transaction	Select All Logs .
Want to recover to a specific System Change Number (SCN)	Select Until SCN (System Change Number) .
Want to recover to a specific data and time	Select Date and Time . You must specify the date and time of the database host's time zone.
Do not want to recover	Select No recovery .
Want to specify any external archive log locations	If the archive logs are not present in the AFS, you should specify the location that contains the archive logs required for recovery.

- b. Select the check box if you want to open the database after recovery.

In a RAC setup, only the RAC instance that is used for recovery is opened after recovery.

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

Restore SQL Server database

You can restore SQL Server database either to the same host or to the alternate host. Recovery of log backups and reseed of availability groups are not supported.



IMPORTANT: Restore of secondary backups through primary is supported.



Single File Restore (SFR) is not supported.

What you will need

If you want to restore a backup that was moved to object store from secondary storage, you should add the primary working environment details and set the IP space. For more information, see [Set IP space of the primary working environment](#).

Steps

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

Option	Description
Restore the database to the same host where the backup was created	Select this option if you want to restore the database to the same SQL server where the backups are taken.
Restore the database to an alternate host	<p>Select this option if you want the database to be restored to a different SQL server in the same or different host where backups are taken.</p> <p>Select a host name, provide a database name (optional), select an instance, and specify the restore paths.</p> <div data-bbox="873 894 927 947" style="border: 1px solid #ccc; border-radius: 50%; width: 30px; height: 30px; display: flex; align-items: center; justify-content: center; margin: 10px 0;"> i </div> <p data-bbox="987 856 1455 982">The file extension provided in the alternate path must be same as the file extension of the original database file.</p> <p data-bbox="841 1035 1463 1129">If the Restore the database to an alternate host option is not displayed in the Restore Scope page, clear the browser cache.</p>

6. If the snapshot is in archival storage, specify the priority to restore your data from the archival storage.
7. On the **Pre Restore Options** page, select one of the following options:
 - Select **Overwrite the database with same name during restore** to restore the database with the same name.
 - Select **Retain SQL database replication settings** to restore the database and retain the existing replication settings.
8. On the **Post Restore Options** page, to specify the database state for restoring additional transactional logs, select one of the following options:
 - Select **Operational, but unavailable** if you are restoring all of the necessary backups now.

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

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

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

If the Undo directory option is enabled, more transaction logs are restored. If the restore operation for the transaction log is unsuccessful, the changes can be rolled back. The SQL Server documentation contains more information.

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

Mount application backups

SnapCenter does not support restore of Oracle and HANA backups to alternate host. Thus, Cloud Backup for Applications allows you to mount the Oracle and HANA backups to the given host.

What you will need

If you want to mount a backup that was moved to object store from secondary storage, you should add the primary working environment details and set the IP space. For more information, see [Set IP space of the primary working environment](#).

Steps

1. In BlueXP UI, click **Protection > Backup and recovery > Applications**.
2. In the Filter By field, select **Type** and from the drop-down either select **SAP HANA** or **Oracle**.
3. Click **...** corresponding to the protected application and select **View Details**.
4. Click **...** corresponding to the backup and select **Mount**.
 - a. Specify one of the following:
 - i. For NAS environment, specify the FQDN or IP address of the host to which alternate volumes restored from object store are to be exported.
 - ii. For SAN environment, specify the initiators of the host to which LUNs of alternate volume restored from object store are to be mapped.
 - b. Specify the suffix that will be added to the alternate volume name.
 - c. If the snapshot is in archival storage, specify the priority to retrieve your data from the archival storage.
 - d. Click **Mount**.

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

For information on how to bring up the SAP HANA database see, [TR-4667: Automating SAP HANA System Copy and Clone Operations with SnapCenter](#).

Back up and restore cloud native applications data

Protect your cloud native applications data

Cloud Backup for Applications is a SaaS based service that provides data protection capabilities for applications running on NetApp Cloud Storage. Cloud Backup for

Applications enabled within NetApp BlueXP (formerly Cloud Manager) offers efficient, application consistent, policy-based backup, and restore of Oracle databases residing on Amazon FSx for NetApp ONTAP.

Architecture

The Cloud Backup for Applications architecture includes the following components.

- Cloud Backup for Applications is a set of data protection services hosted as a SaaS service by NetApp and is based on the BlueXP SaaS platform.

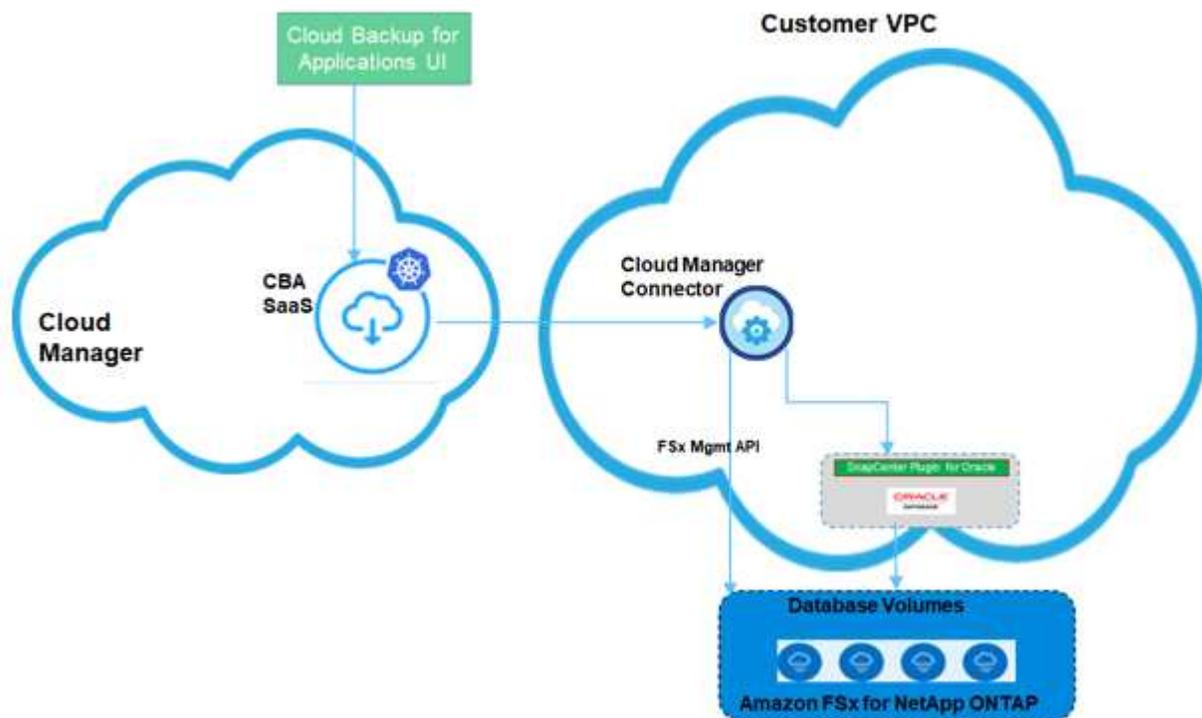
It orchestrates the data protection workflows for applications residing on NetApp Cloud Storage.

- Cloud Backup for Applications UI is integrated with BlueXP UI.

The Cloud Backup for Applications UI offers multiple storage and data management capabilities.

- BlueXP Connector is a component from BlueXP that runs in your cloud network and interacts with Amazon FSx storage file systems and SnapCenter Plug-in for Oracle running on Oracle database hosts.
- SnapCenter Plug-in for Oracle is a component that runs on each Oracle database hosts and interacts with the Oracle databases running on the host while performing data protection operations.

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



For any user-initiated request, the Cloud Backup for Applications UI communicates with the BlueXP SaaS which upon validating the request processes the same. If the request is to run a workflow such as a backup or restore, the SaaS service initiates the workflow and where required, forwards the call to the BlueXP Connector. The Connector then communicates with Amazon FSx for NetApp ONTAP and SnapCenter Plug-in for Oracle as part of running the workflow tasks.

The Connector can be deployed in the same VPC as that of the Oracle databases, or in a different one. If

the Connector and Oracle databases are on different network, you should establish a network connectivity between them.



Cloud Backup for Applications infrastructure is resilient to availability zone failures within a region. It now supports regional failures by failing over to a new region and this failover involves a downtime of around 2 hours.

Supported configurations

- Operating System:
 - RHEL 7.5 or later and 8.x
 - OL 7.5 or later and 8.x
- Storage System: Amazon FSx for ONTAP
- Storage layouts: NFS v3 and v4.1 (dNFS is supported) and iSCSI with ASM (ASMFD, ASMLib and ASMUdev)
- Applications: Oracle Standard and Oracle Enterprise – Standalone (legacy and multitenant – CDB and PDB)
- Oracle versions: 12cR2, 18c, and 19c

Features

- Auto-discovery of Oracle databases
- Backing up Oracle databases residing on Amazon FSx for NetApp ONTAP
 - Full (data + control + archive log files) backup
 - On-demand backup
 - Scheduled backup based on the system-defined or custom policies

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

- Retaining backups based on the policy
- Restoring complete Oracle database (data files + control file) from the specified backup
- Restoring data files only and control files only from the specified backup
- Recovering Oracle database with until SCN, until time, all available logs, and no recovery options
- Monitoring backups and other jobs
- Displaying the protection summary on the dashboard
- Sending alerts through email

Limitations

- Does not support Oracle versions 11g and 21c
- Does not support mount, clone, catalog, and verification operations on backups
- Does not support Oracle on RAC and Data Guard
- Backup limitations:

- Does not support online data or log only backups
- Does not support offline backups
- Does not support backing up of Oracle database residing on recursive mount points
- Does not support consistency group Snapshots for Oracle databases residing on Multiple ASM disk groups with overlap of FSx volumes
- If your Oracle databases are configured on ASM, ensure your SVM names are unique across the FSx systems. If you have same SVM name across FSx systems, back up of Oracle databases residing on those SVMs are not supported.
- Restore limitations:
 - Does not support granular restores, for example restoring of tablespaces and PDBs
 - Supports only in-place restore of Oracle databases on NAS and SAN Layouts
 - Does not support restore of control file only or data files + control file of Oracle databases on SAN layouts
 - In SAN layout, restore operation fails if SnapCenter Plug-in for Oracle finds any foreign files other than Oracle data files on the ASM diskgroup. The foreign files could be one or more of the following types:
 - Parameter
 - Password
 - archive log
 - online log
 - ASM parameter file.

You should select Force in-place restore checkbox to override the foreign files of type parameter, password, and archive log.



If there are other types of foreign files, restore operation fails and the database cannot be recovered. If you have other type of foreign files, you should delete or move them to a different location before performing restore operation.

The failure message due to the presence of foreign files are not displayed in the job page in UI due to a known issue. Check the connector logs if there is a failure during SAN pre-restore stage to know the cause of the issue.

Prerequisites

You should have access to BlueXP, created a BlueXP account, created the working environment and a Connector, and deployed the SnapCenter Plug-in for Oracle.

Access BlueXP

You should [log into BlueXP](#), and then set up a [NetApp account](#).

Create FSx for ONTAP working environment

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

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

Create a Connector

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

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

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



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

After creating the connector, click **Storage > Canvas > My Working Environments > Add Working Environment** and follow the prompts to add the working environment. Ensure that there is connectivity from the connector to the Oracle database hosts and FSx working environment. The connector should be able to connect to the cluster management IP address of the FSx working environment.



After creating the Connector, click **Connector > Manage Connectors**; select the Connector name and copy the Connector ID.

Deploy SnapCenter Plug-in for Oracle

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



Ensure that JAVA 8 is installed on each of the Oracle database hosts and the JAVA_HOME variable is set appropriately.

Plug-in deployment using SSH key based authentication

If SSH key based authentication is enabled on the Oracle host, you can perform the following steps to deploy the plug-in. Before performing the steps, ensure that the SSH connection to the Connector is enabled.

1. Log into the Connector VM as non root user.

2. Obtain the base mount path.

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

3. Deploy the plug-in.

```
sudo <base_mount_path>/scripts/oracle_plugin_copy_and_install.sh --host  
<host_name> --sshkey <ssh_key_file> --username <user_name> --port <ssh_port>  
--pluginport <plugin_port> --installdir <install_dir>
```

- `host_name` is the name of the Oracle host and this is a mandatory parameter.
- `ssh_key_file` is SSH key used to connect to the Oracle host and this is a mandatory parameter.
- `user_name`: User with SSH privileges on the Oracle host and this is an optional parameter. Default value is `ec2-user`.
- `ssh_port`: SSH port on the Oracle host and this is an optional parameter. Default value is 22
- `plugin_port`: Port used by the plug-in and this is an optional parameter. Default value is 8145
- `install_dir`: Directory where the plug-in will be deployed and this is an optional parameter. Default value is `/opt`.

For example: `sudo /var/lib/docker/volumes/service-manager-2_cloudmanager_scs_cloud_volume/_data/scripts/oracle_plugin_copy_and_install.sh --host xxx.xx.x.x --sshkey /keys/netapp-ssh.ppk`

Manual deployment of the plug-in

If SSH key based authentication is not enabled on the Oracle host, you should perform the following manual steps to deploy the plug-in.

1. Log into the Connector VM.

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

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

3. Obtain the base mount path.

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

4. Obtain the binary path of the downloaded plug-in.

```
sudo ls <base_mount_path> $(sudo docker ps|grep -Po  
"cloudmanager_scs_cloud:.*? " | sed -e 's/ *$//' | cut -f2 -d":")/sc-linux-host  
-plugin/snapcenter_linux_host_plugin_scs.bin
```

5. Copy `snapcenter_linux_host_plugin_scs.bin` to each of the Oracle database hosts either using `scp` or other alternate methods.

6. On the Oracle database host, run the following command to enable execute permissions for the binary.

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

7. Deploy the Oracle plug-in as a root user.

```
./snapcenter_linux_host_plugin_scs.bin -i silent
```

8. Copy `certificate.p12` from `<base_mount_path>/client/certificate/` path of the Connector VM to `/var/opt/snapcenter/spl/etc/` on the plug-in host.

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

```
keytool -v -importkeystore -srckeystore certificate.p12 -srcstoretype PKCS12  
-destkeystore keystore.jks -deststoretype JKS -srcstorepass snapcenter  
-deststorepass snapcenter -srccalias agentcert -destalias agentcert -noprompt
```

b. Restart SPL: `systemctl restart spl`

Back up cloud native applications data

Discover the applications

You should discover the databases on the host to assign policies and create backups.

What you will need

- You should have created the FSx for ONTAP working environment and the Connector.
- Ensure that the Connector has connectivity to the FSx for ONTAP working environment and Oracle database hosts.
- Ensure that the BlueXP user has the “Account Admin” role.
- You should have deployed the SnapCenter Plug-in for Oracle. [Learn more](#).

Steps

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

A service account with *SnapCenter System* role is created to perform scheduled data protection operations for all the users in this account.

- Click **Account > Manage Account > Members** to view the service account.



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

4. In the Specify Host Details page, enter the details of the Oracle database host, select the check box to confirm that the plug-in is installed on the host, and click **Discover**.
 - Displays all the databases on the host. If OS authentication is disabled for the database, you should configure database authentication by clicking **Configure**. For more information, refer to [Configure Oracle database credentials](#).
 - Click **Manage Application**, select **Add** to add a new host, **Refresh** to discover new databases, or **Remove** to remove a database host.
 - Click **Settings** and select **Policies** to view the pre-canned policies. Review the pre-canned policies and if you want you can either edit them to meet your requirement or create a new policy.

Configure Oracle database credentials

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

Steps

1. If OS authentication is disabled for the database, you should configure database authentication by clicking **Configure**.
2. Specify the username, password, and the port details either in the Database Settings or ASM Settings section.

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

3. Click **Configure**.

Create policy

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

Steps

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

Back up the cloud native application data

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



When creating ASM diskgroups for Oracle, ensure that there are no common volumes across diskgroups. Each diskgroup needs to have dedicated volumes.

Steps

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

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

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

The backups will be created as per the schedule defined in the policy.



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

Create on-demand backups

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

Steps

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

Find more information

After restoring a large database (250 GB or more), if you perform a full online backup on the same database the operation might fail with the following error:

```
failed with status code 500, error
{"error\":{\"code\":\"app_internal_error\",\"message\":\"Failed to create
snapshot. Reason: Snapshot operation not allowed due to clones backed by
snapshots. Try again after sometime.
```

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

Manage protection of cloud native application data

Monitor jobs

You can monitor the status of the jobs that have been initiated in your working environments. This allows you to see the jobs that have completed successfully, those currently in progress, and those that have failed so you can diagnose and fix any problems.

You can view a list of all the operations and their status. Each operation, or job, has a unique ID and a status. The status can be:

- Successful
- In Progress
- Queued
- Warning
- Failed

Steps

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

You can click the name of a job to view details corresponding to that operation. If you are looking for specific job, you can:

- use the time selector at the top of the page to view jobs for a certain time range
- enter a part of the job name in the Search field
- sort the results by using the filter in each column heading

View backup details

You can view total number of backups created, policies used for creating backups, database version, and agent ID.

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



The agent ID is associated to the Connector. If a Connector that was used during registering the Oracle database host no longer exists, the subsequent backups of that application will fail because the agent ID of the new Connector is different. You should run the **connector-update** API to modify the agent ID.

Update the Connector Details

If a Connector that was used during registering the Oracle database host no longer exists or is corrupted in AWS, you should deploy a new connector. After deploying the new connector, you should run the **connector-update** API to update the Connector details.

```
curl --location --request PATCH
'https://snapcenter.cloudmanager.cloud.netapp.com/api/oracle/databases/connector-update' \
--header 'x-account-id: <CM account-id>' \
--header 'x-agent-id: <connector Agent ID >' \
--header 'Authorization: Bearer token' \
--header 'Content-Type: application/json' \
--data-raw '{
"old_connector_id": "Old connector id that no longer exist",
"new_connector_id": "New connector Id"
}
```

After updating the Connector details, you should connect to the Oracle database host and perform the following steps:

1. Obtain the plug-in information running on the Oracle database host.
`rpm -qi netapp-snapcenter-plugin-oracle`
2. Uninstall the plug-in.
`sudo /opt/NetApp/snapcenter/spl/installation/plugins/uninstall`
3. Verify that the plug-in is uninstalled successfully.
`rpm -qi netapp-snapcenter-plugin-oracle`

After uninstalling the plug-in, you can deploy the plug-in. [Learn more](#).

Configure CA signed certificate

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

Configure CA signed certificate for client certificate authentication

The connector uses a self-signed certificate to communicate with plug-in. The self-signed certificate is imported to the keystore by the installation script. You can perform the following steps to replace the self-signed certificate with CA signed certificate.

What you will need

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

```
sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs sudo
```

```
docker volume inspect | grep Mountpoint
```

Steps

1. Login to Connector.
2. Delete all the existing files located at `<base_mount_path>/client/certificate` in the Connector virtual machine.
3. Copy the CA signed certificate and key file to the `<base_mount_path>/client/certificate` in the Connector virtual machine.

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

4. Create the PKCS12 format of the certificate with the name `certificate.p12` and keep at `<base_mount_path>/client/certificate`.
5. Copy the `certificate.p12` and certificates for all the intermediate ca and root ca to the plug-in host at `/var/opt/snapcenter/spl/etc/`.
6. Log in to the plug-in host.
7. Navigate to `/var/opt/snapcenter/spl/etc` and run the `keytool` command to import the `certificate.p12` file.

```
keytool -v -importkeystore -srckeystore certificate.p12 -srcstoretype PKCS12 -destkeystore keystore.jks -deststoretype JKS -srcstorepass snapcenter -deststorepass snapcenter -srccalias agentcert -destalias agentcert -noprompt
```
8. Import the root CA and intermediate certificates.

```
keytool -import -trustcacerts -keystore keystore.jks -storepass snapcenter -alias trustedca -file <certificate.crt>
```



The `certfile.crt` refers to the certificates of root CA as well as intermediate CA.

9. Restart SPL:

```
systemctl restart spl
```

Configure CA signed certificate for server certificate of plug-in

The CA certificate should have the exact name of the Oracle plug-in host with which the Connector virtual machine communicates.

What you will need

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

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

Steps

1. Perform the following steps on the plug-in host:
 - a. Navigate to the folder containing the SPL's keystore `/var/opt/snapcenter/spl/etc`.
 - b. Create the PKCS12 format of the certificate having both certificate and key with alias `splkeystore`.
 - c. Add the CA certificate.

```
keytool -importkeystore -srckeystore <CertificatePathToImport> -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS -srccalias splkeystore
```

```
-destalias splkeystore -noprompt
```

d. Verify the certificates.

```
keytool -list -v -keystore keystore.jks
```

e. Restart SPL: `systemctl restart spl`

2. Perform the following steps on the Connector:

a. Log in to the Connector as non-root user.

b. Copy the entire chain of CA certificates to the persistent volume located at `<base_mount_path>/server`.

Create the server folder if it does not exist.

c. Connect to the `cloudmanager_scs_cloud` and modify the **enableCACert** in `config.yml` to **true**.

```
sudo docker exec -t cloudmanager_scs_cloud sed -i 's/enableCACert:
false/enableCACert: true/g' /opt/netapp/cloudmanager-scs-
cloud/config/config.yml
```

d. Restart `cloudmanager_scs_cloud` container.

```
sudo docker restart cloudmanager_scs_cloud
```

Access REST APIs

The REST APIs to protect the applications to cloud is available [here](#).

You should obtain the user token with federated authentication to access the REST APIs. For information to obtain the user token, refer to [Create a user token with federated authentication](#).

Restore cloud native application data

In the event of data loss, you can restore the data files, control files, or both and then recover the database.

Steps

1. Click **...** corresponding to the database that you want to restore and click **View Details**.
2. Click **...** corresponding to the data backup that you want to use for restoring and click **Restore**.
3. In the Restore Scope section, perform the following actions:

If you...	Do this...
Want to restore only the data files	Select All Data Files .
Want to restore only the control files	Select Control Files
Want to restore both data files and control files	Select All Data Files and Control Files .



Restore of datafiles with control files or only control files are not supported for iSCSI on ASM layout.

You can also select **Force in-place restore** checkbox.

The **Force in-place restore** option overrides the spfile, password file, and archive log files from the data files diskgroup. You should use the latest backup when **Force in-place restore** option is selected.

4. In the Recovery Scope section, perform the following actions:

If you...	Do this...
Want to recover to the last transaction	Select All Logs .
Want to recover to a specific System Change Number (SCN)	Select Until System Change Number and specify the SCN.
Want to recover to a specific date and time	Select Date and Time .
Do not want to recover	Select No recovery .

For the selected recovery scope, in the **Archive Log Files Locations** field you can optionally specify the location that contains the archive logs required for recovery.

Select the check box if you want to open the database in READ-WRITE mode after recovery.

5. Click **Next** and review the details.
6. Click **Restore**.

Copyright information

Copyright © 2022 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.