



Back up and restore cloud-native applications data

BlueXP backup and recovery

NetApp
April 18, 2024

Table of Contents

- Back up and restore cloud-native applications data 1
 - Protect your cloud-native applications data 1
 - Back up cloud-native Oracle databases 4
 - Back up cloud-native SAP HANA databases 17
 - Back up cloud-native SQL Server databases using REST APIs 26
 - Restore cloud-native Oracle databases 38
 - Restore cloud-native SAP HANA databases 40
 - Restore Microsoft SQL Server database 42
 - Clone cloud-native Oracle databases 45
 - Refresh SAP HANA target system 53
 - Manage protection of cloud-native application data 55

Back up and restore cloud-native applications data

Protect your cloud-native applications data

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

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

Architecture

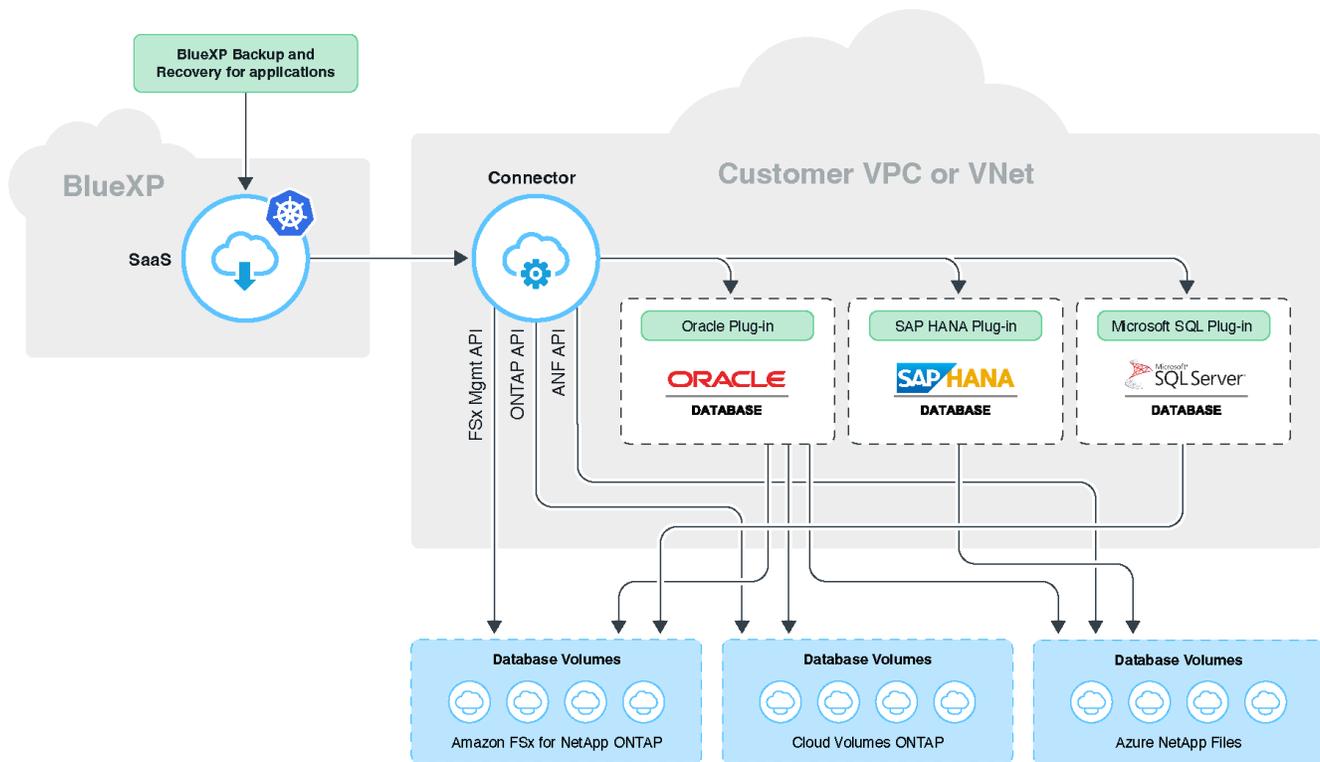
The BlueXP backup and recovery for applications architecture includes the following components.

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

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

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

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



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

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



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



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

Protect Oracle databases

Features

- Add host and deploy plug-in

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

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

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

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

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

Limitations

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

Protect SAP HANA databases

Features

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

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

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

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

Limitations

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

Protect Microsoft SQL Server database

Features

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

Limitations

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

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

- Restore to alternate location is not supported

Back up cloud-native Oracle databases

Quick start

Get started quickly by following these steps.

1

Verify support for your configuration

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



Azure NetApp Files does not support SAN environment.

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

2

Sign up to BlueXP

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

3

Log into BlueXP

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

4

Manage your BlueXP account

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

Configure FSx for ONTAP

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

Create FSx for ONTAP working environment

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

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

Create a Connector

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

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

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



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

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

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

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

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

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

Configure Cloud Volumes ONTAP

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

Create Cloud Volumes ONTAP working environment

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

Create a Connector

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

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

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

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

Configure Azure NetApp Files

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

Create Azure NetApp Files working environment

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

Create a connector

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

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

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

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

Install SnapCenter Plug-in for Oracle and add database hosts

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

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

Prerequisites

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

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

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

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

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

Add host from UI using SSH option

Steps

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

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

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

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

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

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

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

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

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

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

Root user is not supported.

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

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

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

- e. Select the Connector.

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



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

- g. Click **Next**.

6. In the Configuration page, perform the following:

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

7. Review the details and click **Discover Applications**.

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

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

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

steps to install the plug-in.

Before you begin

Ensure that the SSH connection to the Connector is enabled.

Steps

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

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

The service account (*SnapCenter-account- \langle accountid \rangle*) is used to run the scheduled backup operations.

You should never delete the service account.

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

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

Ensure that the Connector can communicate with the database host using the FQDN or IP address.
 - c. Specify the plug-in port.

Default port is 8145.
 - d. Specify the non-root (sudo) user using which the plug-in package will be copied to the host.
 - e. Select the Connector.
 - f. Select the check box to confirm that the plug-in is installed on the host.
 - g. Click **Next**.
6. In the Configuration page, perform the following:
 - a. Configure sudo access for the SnapCenter user in the Oracle database host by logging in to the Linux machine running Oracle database.
 - b. Copy the text displayed in BlueXP UI.
 - c. Create the */etc/sudoers.d/snapcenter* file on the Linux machine and paste the copied text.
 - d. In the BlueXP UI, select the checkbox and click **Next**.

7. Log into the Connector VM.

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

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

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

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

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

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

For example:

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

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

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

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

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

Steps

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

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

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

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

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

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

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

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

- c. Specify the plug-in port.

Default port is 8145.

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

6. In the Configuration page, perform the following:

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

7. Log into the Connector VM.

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

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

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

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

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

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

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

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

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

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

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

Configure Oracle database credentials

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

Steps

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

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

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

3. Click **Configure**.

Upgrade SnapCenter Plug-in for Oracle Database

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

Before you begin

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

Steps

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

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

Back up cloud-native Oracle databases

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

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

Create policy to protect Oracle database

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

Steps

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

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

Optionally, you can also specify the arguments.

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

The post script supports a set of environment variables.

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

8. Click **Create**.

Configure RMAN catalog repository

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

Before you begin

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

Steps

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

Create a backup of the Oracle Database

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



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

Steps

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

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

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

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



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

Create on-demand backup of the Oracle database

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

Steps

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

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

Limitations

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

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

Back up cloud-native SAP HANA databases

Quick start

Get started quickly by following these steps.

1

Verify support for your configuration

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

2

Sign up to BlueXP

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

3

Log into BlueXP

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



Manage your BlueXP account

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

Configure Azure NetApp Files

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

Create Azure NetApp Files working environment

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

Create a connector

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

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

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

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

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

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

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

Prerequisites

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

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

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

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

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



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

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

Configure sudo for SnapCenter user

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

Steps

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


```
sudo docker exec -it cloudmanager_scs_cloud curl -X GET 'http://127.0.0.1/deploy/downloadLinuxPlugin'
```
3. Copy the contents of **sudoer.txt** located at: `/var/lib/docker/volumes/service-manager_2_cloudmanager_scs_cloud_volume/_data/${sudo docker ps|grep -Po "cloudmanager_scs_cloud:.? "}|sed -e 's/ *$//|cut -f2 -d":')/sc-linux-host-plugin`
4. Log into the SAP HANA system host using root user account.
5. Configure sudo access for the non-root user by copying the text copied in the step 3 to `/etc/sudoers.d/snapcenter` file.

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

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

Install the plug-in using script

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

Before your begin

Ensure that the SSH connection to the Connector is enabled.

Steps

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

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

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

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

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

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

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

Install the plug-in manually

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

Steps

1. Log in to Connector VM.

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

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

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

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

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

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

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

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

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

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

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

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

9. Restart SPL: `systemctl restart spl`

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

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

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

Upgrade SnapCenter Plug-in for SAP HANA Database

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

Before you begin

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

Steps

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

2. Run the following script.

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

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

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

Add SAP HANA database hosts

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

Steps

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

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

You can delete or modify the user store keys.

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

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

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

c. Select **Next**.

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

Select the check boxes to confirm.

9. Select **Next**.

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

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

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

b. Select the required volumes.

c. Select **Add Storage**.

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

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

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

Add Non-Data Volumes

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

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

Steps

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

Add Global Non-Data Volumes

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

Steps

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

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

Back up cloud-native SAP HANA databases

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

Create a policy to protect SAP HANA database

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

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

Prescripts and postscripts

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

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

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

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

Environmental variables

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

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

Create a backup of the SAP HANA Database

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

Before you begin

You should have added the SAP HANA database hosts.

[Add SAP HANA database hosts](#)

About this task

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

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

Steps

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

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

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

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



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

Create on-demand backup of the SAP HANA database

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

Steps

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

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

- c. Click **Create Backup**.

Back up cloud-native SQL Server databases using REST APIs

Quick start

Get started quickly by following these steps.

1

Verify support for your configuration

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

NAS configuration is not supported.

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

2

Sign up to BlueXP

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

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

3

Log into BlueXP

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

4

Manage your BlueXP account

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

Configure FSx for ONTAP

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

Create FSx for ONTAP working environment

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

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

Create a Connector

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

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

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



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

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

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

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

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

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

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

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

Install the SnapCenter Plug-in for SQL Server

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

Before you begin

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

Steps

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

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

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

2. Copy `snapcenter_service_windows_host_plugin.exe` from the connector to each of the MSSQL Server database hosts either using `scp` or other alternate methods.
3. Install the plug-in.

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

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

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

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

Add the SQL Server database host

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

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

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

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameters

Name	Type	Required
addr	string	True
connector_id	string	True
plugin_type	string	True
install_method	string	True
plugin_port	number	True
username	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

View the added SQL Server database hosts

You can run this API to view all the added SQL Server database hosts.

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

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

Response

If the API is executed successfully, response code 200 is displayed.

Example:

```

{
  "num_records": 1,
  "total_records": 1,
  "records": [
    {
      "id": "85bd4603-08f7-45f4-ba8e-a0b1e2a0f4d0",
      "addr": "scspa2722211001.rtp.openenglab.netapp.com",
      "status": "Running",
      "connector_id": "fBf8Iwbp4BscBfD02qBwWm6I03gGAesRclients",
      "plugin_port": 8145,
      "plugins": [
        {
          "type": "mssql"
        }
      ],
      "os_type": "windows",
      "platform": "onprem",
      "username": "administrator",
      "operating_mode": "production"
    }
  ],
  "_links": {
    "next": {}
  }
}

```

Discover the database instances

You can run this API and enter the host ID to discover all the MSSQL instances.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/instances/discovery'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/MSSQL%20Instances/MSSQLInstancesDiscoveryRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameter

Name	Type	Required
host_id	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

View the discovered database instances

You can run this API to view all the discovered database instances.

'GET snapcenter.cloudmanager.cloud.netapp.com/api/mssql/instances'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/MSSQL%20Instances/GetMSSQLInstancesRequest>

Response

If the API is executed successfully, response code 200 is displayed.

Example:

```

{
  "num_records": 2,
  "total_records": 2,
  "records": [
    {
      "id": "953e66de-10d9-4fd9-bdf2-bf4b0eaabfd7",
      "name": "scspa2722211001\\NAMEDINSTANCE1",
      "host_id": "85bd4603-08f7-45f4-ba8e-a0b1e2a0f4d0",
      "status": "Running",
      "auth_mode": 0,
      "version": "",
      "is_clustered": false,
      "is_credentials_configured": false,
      "protection_mode": ""
    },
    {
      "id": "18e1b586-4c89-45bd-99c8-26268def787c",
      "name": "scspa2722211001",
      "host_id": "85bd4603-08f7-45f4-ba8e-a0b1e2a0f4d0",
      "status": "Stopped",
      "auth_mode": 0,
      "version": "",
      "is_clustered": false,
      "is_credentials_configured": false,
      "protection_mode": ""
    }
  ],
  "_links": {
    "next": {}
  }
}

```

Configure the database instance credentials

You can run this API to validate and set credentials for the database instances.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql//api/mssql/credentials-configuration'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/MSSQL%20Instances/ConfigureCredentialRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameter

Name	Type	Required
host_id	string	True
instance_ids	string	True
username	string	True
password	string	True
auth_mode	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

Back up cloud-native Microsoft SQL Server databases

You can create scheduled or on-demand backups by assigning the policies that you created.

Create backup policy

You can run this API to create the backup policy.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/backup/policies'

For more information, refer to: https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/MSSQL%20Backup%20Policies/MSSQLBackupPolicyService_CreateMSSQLBackupPolicy

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameters

Name	Type	Required
name	string	True
backup_type	string	True
copy_only_backup	string	False
is_system_defined	string	False
backup_name_format	string	True
schedule_type	string	True
start_time	number	True
hours_interval	number	True
minutes_interval	number	True
retention_type	string	True
retention_count	number	True
end_time	number	True

Response

If the API is executed successfully, response code 201 is displayed.

Example:

```
{
  "_links": {
    "self": {
      "href": "/api/resourcelink"
    }
  }
}
```

Assign policy to SQL database instance

You can run this API to assign policy to SQL database instance.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/instances/{id}/policy-assignment'

Where, *id* is MSSQL instance ID obtained by running the discover database instance API. For more information, refer to [Discover the database instances](#).

Array of IDs is the input here. For example:

```
[
  "c9f3e68d-1f9c-44dc-b9af-72a9dfc54320"
]
```

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Policy%20Assignment/PostMSSQLInstanceAssignPolicyRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resource/link"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

Create an on-demand backup

You can run this API to create an on-demand backup.

'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/backups'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Backups/CreateMSSQLBackupRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameters

Name	Type	Required
 This is ID of the MSSQL database instance.	string	True
resource_type	string	True
policy_id	string	True
schedule_type	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```
{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourcelink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}
```

View the backups

You can run these APIs to list view all the backups and also to view details of a particular backup.

'GET snapcenter.cloudmanager.cloud.netapp.com/api/mssql/backups'

'GET snapcenter.cloudmanager.cloud.netapp.com/api/mssql/backups/{id}'

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Backups/MSSQLGetBackupsRequest>

Response

If the API is executed successfully, response code 200 is displayed.

Example:

```

{
  "total_records": 1,
  "num_records": 1,
  "records": [
    {
      "backup_id": "602d7796-8074-43fc-a178-eee8c78566ac",
      "resource_id": "a779578d-cf78-46f3-923d-b9223255938c",
      "backup_name":
"Hourly_policy2_scspa2722211001_NAMEDINSTANCE1_2023_08_08_07_02_01_81269_0
",
      "policy_name": "policy2",
      "schedule_type": "Hourly",
      "start_time": "2023-08-08T07:02:10.203Z",
      "end_time": "0001-01-01T00:00:00Z",
      "backup_status": "success",
      "backup_type": "FullBackup"
    }
  ],
  "_links": {
    "next": {}
  }
}

```

Restore cloud-native Oracle databases

Restore cloud-native Oracle databases to original location

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

Before you begin

If the Oracle 21c database is in STARTED state, the restore operation fails. You should run the following command to restore the database successfully.

```

cp -f <ORACLE_HOME>/jdbc/lib/ojdbc8.jar
/opt/NetApp/snapcenter/spl/plugins/sco/lib/ojdbc8-8.jar

```

Steps

1. Click **...** corresponding to the database that you want to restore and click **Restore**.
2. Select the restore point to which the database should be restored and click **Restore to original location**.
3. In the Restore Scope section, perform the following actions:

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

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

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

In Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP SAN layout, if SnapCenter Plug-in for Oracle finds any foreign files other than Oracle data files on the ASM diskgroup, connect and copy restore method is performed. The foreign files could be one or more of the following types:

- Parameter
- Password
- archive log
- online log
- ASM parameter file.

The **Force in-place restore** option overrides the foreign files of type parameter, password, and archive log. You should use the latest backup when **Force in-place restore** option is selected.

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

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

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

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

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

Restore cloud-native Oracle databases to alternate location

In the event of data loss, you can restore the Oracle database to alternate location only on Azure NetApp Files. The alternate location can be on a different host or on the same host.

Before you begin

- If the Oracle 21c database is in STARTED state, the restore operation fails. You should run the following command to restore the database successfully.

```
cp -f <ORACLE_HOME>/jdbc/lib/ojdbc8.jar  
/opt/NetApp/snapcenter/spl/plugins/sco/lib/ojdbc8-8.jar
```
- You should ensure that the Oracle version on the alternate host is same as that of the original host.

About this task

While initiating the restore operation, you are not allowed to modify the configurations except the Oracle home, maximum volume throughput, Oracle SID, and database credentials.

Full recovery is enabled by default with *Until cancel* set to true.

Archive log mode is turned off by default for the restored database. You can enable archive log mode and keep the archive logs on the NetApp volume if required.

Steps

1. Click **...** corresponding to the database that you want to restore and click **Restore**.
2. Select the restore point to which the database should be restored and click **Restore to alternate location > Next**.
3. In the Configuration page, specify the details of the alternate location, SID, Oracle_Home, database credentials, and storage throughput.

For the database credential, if the OS User authentication is disabled, you should provide a password for the sys user to connect to the restored database on the same or target host.

4. Click **Next**, review the details and click **Restore**.

The progress of the restore operation can be viewed in the Job Monitor page. After the job is completed, click **Refresh Discovery** to view the restored database. However, you cannot protect the database that is restored to alternate location.

Restore cloud-native SAP HANA databases

In the event of data loss, you can restore the data and non-data files and then recover the database.

Before you begin

- The SAP HANA system must be in a stopped state.
- If the SAP HANA system is up and running, you can provide a prescript to stop the system.

About this task

- If you enable the ANF backups on a volume, Single File SnapRestore operation is performed.
- For Non-Data Volumes and Global Non-Data Volumes, connect and copy restore operation is performed.
 - The Quality of Service (QoS) values for connect and copy restore operation are picked up from the source volumes of Non-Data Volumes or Global Non-Data Volumes.



QoS is applicable only for capacity pools of type "Manual".

Steps

1. Click [...](#) corresponding to the database that you want to restore and click **View Details**.
2. Click [...](#) corresponding to the data backup that you want to restore and click **Restore**.
3. In the **Restore System** page, enter the scripts. [Prescripts and postscripts](#).

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

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

4. Click **Restore**.

What's next

After restoring, manually recover the SAP HANA system or provide a postscript, which performs the SAP HANA system recovery.

Restore Non-Data Volume

About this task

For connect and copy restore operation, go to Microsoft Azure portal, select the volume, click **Edit**, and enable **Hide snapshot path**.

Steps

1. In the **Applications** page, select Non-Data Volume from the drop-down box.
2. Click [...](#) corresponding to the backup that you want to restore and click **Restore**.

Restore Global Non-Data Volume

About this task

For connect and copy restore operation, go to Microsoft Azure portal, select the volume, click **Edit**, and enable **Hide snapshot path**.

Steps

1. In the **Applications** page, click on the Global Non-Data Volume that you want to restore.
2. Click **...** corresponding to the Global Non-Data Volume that you want to restore and click **Restore**.

Restore Microsoft SQL Server database

You can restore Microsoft SQL Server database to the same host. You should first get list of databases and then restore the database.

View the list of databases

You can run this API to view the list of databases.

```
'GET snapcenter.cloudmanager.cloud.netapp.com/api/mssql/databases'
```

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#!/MSSQL%20Databases/GetMSSQLDatabasesRequest>

Response

If the API is executed successfully, response code 200 is displayed.

Example:

```

{
  "num_records": 3,
  "total_records": 3,
  "records": [
    {
      "id": "348901e5-aeaa-419f-88b1-80240de3b1fe",
      "name": "DB4",
      "hostname": "scspa2722211001.rtp.openenglab.netapp.com",
      "size": 0.078125,
      "instance_id": "454c8413-5351-41fc-88bf-f20fb050ec87",
      "instance": "scspa2722211001\\NAMEDINSTANCE1",
      "db_status": "Normal",
      "db_access": "eUndefined",
      "db_type": "User",
      "recovery_mode": "Full"
    },
    {
      "id": "c79d33ab-7322-4ed6-92f5-51ad7a6944e0",
      "name": "DB5",
      "hostname": "scspa2722211001.rtp.openenglab.netapp.com",
      "size": 0.078125,
      "instance_id": "454c8413-5351-41fc-88bf-f20fb050ec87",
      "instance": "scspa2722211001\\NAMEDINSTANCE1",
      "db_status": "Normal",
      "db_access": "eUndefined",
      "db_type": "User",
      "recovery_mode": "Full"
    },
    {
      "id": "40d6f35a-f4fb-48bc-8e0a-0ac93ddf0888",
      "name": "model",
      "hostname": "scspa2722211001.rtp.openenglab.netapp.com",
      "size": 0.015625,
      "instance_id": "454c8413-5351-41fc-88bf-f20fb050ec87",
      "instance": "scspa2722211001\\NAMEDINSTANCE1",
      "db_status": "Normal",
      "db_access": "eUndefined",
      "db_type": "System",
      "recovery_mode": "Full"
    }
  ],
  "_links": {
    "next": {}
  }
}

```

Restore and recover the MSSQL database

You can run this API to restore the MSSQL database.

```
'POST snapcenter.cloudmanager.cloud.netapp.com/api/mssql/databases/{id}/restore'
```

Where, *id* is MSSQL database ID obtained by running the view database API. For more information, refer to [View the list of databases](#).

For more information, refer to: <https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/#/MSSQL%20Database%20Restore/RestoreMSSQLDatabaseRequest>

This API creates a job that can be tracked from the **Job Monitor** tab in the BlueXP UI.

Parameters

Name	Type	Required
backup_id	string	True
overwrite_database	bool	True
retain_replication_settings	bool	False
recovery_mode	string <i>The 3 supported strings are Operational, NonOperational, and ReadOnly.</i>	True
undo_file_directory	string	True
restore_type	string	True

Response

If the API is executed successfully, response code 202 is displayed.

Example:

```

{
  "job": {
    "_links": {
      "self": {
        "href": "/api/resourceLink"
      }
    },
    "uuid": "3fa85f64-5717-4562-b3fc-2c963f66afa6"
  }
}

```

Clone cloud-native Oracle databases

Clone concepts and requirements

You can clone an Oracle database residing on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP using the backup of the database either to the source database host or to an alternate host. You can clone the backup from primary storage systems.

Before cloning the database, you should understand the clone concepts and ensure that all the requirements are met.

Requirements for cloning an Oracle database

Before cloning an Oracle database, you should ensure that prerequisites are completed.

- You should have created a backup of the database.
You should have successfully created online data and log backup for the cloning operation to succeed.
- In the `asm_diskstring` parameter, you should configure:
 - `AFD:*` if you are using ASMFD
 - `ORCL:*` if you are using ASMLIB
 - `/dev/<exact_device_location>` if you are using ASMUDEV
- If you are creating the clone on an alternate host, the alternate host should meet the following requirements:
 - The plug-in should be installed on the alternate host.
 - Oracle software should be installed on the alternate host.
 - The clone host should be able to discover LUNs from storage if you are cloning a database residing on iSCSI SAN storage.
If you are cloning to an alternate host, then make sure that an iSCSI session is established between the storage and the alternate host.
 - If the source database is an ASM database:
 - The ASM instance should be up and running on the host where the clone will be performed.
 - The ASM diskgroup should be provisioned prior to the clone operation if you want to place archive log files of the cloned database in a dedicated ASM diskgroup.

- The name of the data diskgroup can be configured but ensure that the name is not used by any other ASM diskgroup on the host where the clone will be performed.
- Data files residing on the ASM diskgroup are provisioned as part of the clone workflow.

Limitations

- Cloning of databases residing on Azure NetApp Files is not supported.
- Cloning of databases residing on Qtree is not supported.
- Backing up a cloned database is not supported.
- If daily automatic backups are enabled on Amazon FSx for NetApp ONTAP, the cloned volumes on Amazon FSx for NetApp ONTAP cannot be deleted from BlueXP UI because FSx would have created backups on the cloned volumes.
You should delete the cloned volumes after deleting all the backups for the volume from FSx UI and then delete the clones from the BlueXP UI using force option.

Clone methods

You can create clone either using the basic method or using the clone specification file.

Clone using basic method

You can create the clone with the default configurations based on the source database and the selected backup.

- The database parameters, home, and OS user are defaulted to the source database.
- The data file paths are named based on the naming scheme selected.
- The pre-script, post-script, and SQL statements cannot be specified.
- The recovery option is by default **until cancel** and it uses the log backup associated with the data backup for recovery

Clone using specification file

You can define the configurations in the clone specification file and use it to clone the database. You can download the specification file, modify it to your requirement, and then upload the file. [Learn more](#).

The different parameters defined in the specification file and that can be modified are as follows:

Parameter	Description
control_files	<p>Location of control files for the clone database.</p> <p>The number of control files will be same as source database.</p> <p>If you want to override the control file path, you can provide a different control file path. The file system or the ASM diskgroup should exist on the host.</p>

Parameter	Description
redo_logs	Location, size, redo group number of redo logs. A minimum of two redo log groups are required to clone the database. If you want to override the redo log file path, you can customize the redo log file path to a different file system than that of the source database. The file system or the ASM diskgroup should exist on the host.
oracle_version	Version of Oracle on the target host.
oracle_home	Oracle home on the target host.
enable_archive_log_mode	Controls the archive log mode for the clone database
database_parameters	Database parameters for the cloned database
sql_statements	The SQL statements to be executed on the database after cloning
os_user_detail	Oracle OS user on the target clone database
database_port	Port used for communicating with the database if OS authentication is disabled on the host.
asm_port	Port used for communicating with ASM database if credentials are provided in the create clone input.
skip_recovery	Does not perform recovery operation.
until_scn	Recovers the database up to the specified system change number (scn).
until_time	Recovers the database up to the specified date and time. The accepted format is <i>mm/dd/yyyy hh:mm:ss</i> .
until_cancel	Recovers by mounting the log backup associated with the data backup that was selected for cloning. The cloned database is recovered till the missing or corrupt log file.
log_paths	Additional locations of archive log paths to be used for recovering the cloned database.

Parameter	Description
source_location	Location of the diskgroup or mount point on the source database host.
clone_location	Location of the diskgroup or mount point that needs to be created on the target host corresponding to the source location.
location_type	It can be either ASM_Diskgroup Or mountpoint. The values are auto-populated at the time of downloading the file. You should not edit this parameter.
pre_script	Script to be executed on the target host before creating the clone.
post_script	Script to be executed on the target host after creating the clone.
path	Absolute Path of the script on the clone host. You should store the script either in <code>/var/opt/snapcenter/spl/scripts</code> or in any folder inside this path.
timeout	The timeout time specified for the script running on the target host.
arguments	Arguments specified for the scripts.

Clone naming scheme

Clone naming scheme defines what will be the location of the mount points and name of the diskgroups of the cloned database. You can either select **Identical** or **Auto-generated**.

Identical naming scheme

If you select the clone naming scheme as **Identical**, the location of mount points and the name of the diskgroups of the cloned database will be same as the source database.

For example, if the mount point of the source database is `/netapp_sourcedb/data_1 , +DATA1_DG`, for the cloned database the mount point remains the same for both NFS and ASM on SAN.

- Configurations like number and path of control files and redo files will be same as source.



If the redo logs or control file paths are located on the non-data volumes, then the user should have provisioned the ASM diskgroup or mountpoint in the target host.

- Oracle OS user and Oracle version will be same as source database.
- Clone storage volume name will be in the following format `sourceVolNameSCS_Clone_CurrentTimeStampNumber`.

For example, if the volume name on the source database is *sourceVolName*, the cloned volume name will be *sourceVolNameSCS_Clone_1661420020304608825*.



The *CurrentTimeStampNumber* provides the uniqueness in volume name.

Auto-generated naming scheme

If you select the cloning scheme as **Auto-generated**, the location of mount points and the name of the diskgroups of the cloned database will be appended with a suffix.

- If you have selected the basic clone method, the suffix will be the **Clone SID**.
- If you have selected the specification file method, the suffix will be the **Suffix** that was specified while downloading the clone specification file.

For example, if the mount point of the source database is */netapp_sourcedb/data_1* and the **Clone SID** or the **Suffix** is *HR*, then the mount point of the cloned database will be */netapp_sourcedb/data_1_HR*.

- Number of control files and redo log files will be same as the source.
- All redo log files and control files will be located on one of the cloned data mount points or data ASM diskgroups.
- Clone storage volume name will be in the following format `sourceVolNameSCS_Clone_CurrentTimeStampNumber`.

For example, if the volume name on the source database is *sourceVolName*, the cloned volume name will be *sourceVolNameSCS_Clone_1661420020304608825*.



The *CurrentTimeStampNumber* provides the uniqueness in volume name.

- The format of the NAS mount point will be *SourceNASMountPoint_suffix*.
- The format of the ASM diskgroup will be *SourceDiskgroup_suffix*.



If the number of characters in the clone diskgroup is greater than 25 then it will have *SC_HashCode_suffix*.

Database parameters

The value of the following database parameters will be same as that of the source database irrespective of the clone naming scheme.

- `log_archive_format`
- `audit_trail`
- `processes`
- `pga_aggregate_target`
- `remote_login_passwordfile`

- undo_tablespace
- open_cursors
- sga_target
- db_block_size

The value of the following database parameters will be appended with a suffix based on the clone SID.

- audit_file_dest = {sourcedatabase_parametervalue}_suffix
- log_archive_dest_1 = {sourcedatabase_oraclehome}_suffix

Supported predefined environment variables for clone specific prescript and postscript

You can use the supported predefined environment variables when you execute the prescript and postscript while cloning a database.

- SC_ORIGINAL_SID specifies the SID of the source database.
This parameter will be populated for application volumes. Example: NFSB32
- SC_ORIGINAL_HOST specifies the name of the source host.
This parameter will be populated for application volumes. Example: asmrac1.gdl.englab.netapp.com
- SC_ORACLE_HOME specifies the path of the target database's Oracle home directory.
Example: /ora01/app/oracle/product/18.1.0/db_1
- SC_BACKUP_NAME specifies the name of the backup.
This parameter will be populated for application volumes. Examples:
 - If the database is not running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
 - If the database is running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1, RG2_scspr2417819002_07-21-2021_12.16.48.9267_1, RG2_scspr2417819002_07-22-2021_12.16.48.9267_1
- SC_ORIGINAL_OS_USER specifies the operating system owner of the source database.
Example: oracle
- SC_ORIGINAL_OS_GROUP specifies the operating system group of the source database.
Example: oinstall
- SC_TARGET_SID specifies the SID of the cloned database.
For PDB clone workflow, the value of this parameter will not be predefined. This parameter will be populated for application volumes.
Example: clonedb
- SC_TARGET_HOST specifies the name of the host where the database will be cloned.
This parameter will be populated for application volumes. Example: asmrac1.gdl.englab.netapp.com
- SC_TARGET_OS_USER specifies the operating system owner of the cloned database.
For PDB clone workflow, the value of this parameter will not be predefined. Example: oracle
- SC_TARGET_OS_GROUP specifies the operating system group of the cloned database.
For PDB clone workflow, the value of this parameter will not be predefined. Example: oinstall
- SC_TARGET_DB_PORT specifies the database port of the cloned database.
For PDB clone workflow, the value of this parameter will not be predefined. Example: 1521

Supported delimiters

- @ is used to separate data from its database name and to separate the value from its key.
Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
- | is used to separate the data between two different entities for SC_BACKUP_NAME parameter.
Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
- , is used to separate set of variables for the same key.
Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1, RG2_scspr2417819002_07-21-2021_12.16.48.9267_1, RG2_scspr2417819002_07-22-2021_12.16.48.9267_1

Clone cloud-native Oracle databases

You can clone an Oracle database residing on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP using the backup of the database either to the source database host or to an alternate host.

You might clone databases for the following reasons:

- To test functionality that must be implemented using the current database structure and content during application development cycles.
- To populate data warehouses using data extraction and manipulation tools.
- To recover data that was mistakenly deleted or changed.

Before you begin

You should understand the clone concepts and ensure that all the requirements are met. [Learn more.](#)

Steps

1. Click **...** corresponding to the database that you want to clone and click **View Details**.
2. Click **...** corresponding to the data backup and click **Clone**.
3. In the Clone Details page, select one of the clone options.
4. Depending on the option selected, perform the following actions:

If you have selected...	Do this...
<p>Basic</p>	<ol style="list-style-type: none"> 1. Select the clone host. <p>If you want to create the clone on an alternate host, select the host having the same version of Oracle and OS as that of the source database host.</p> <ol style="list-style-type: none"> 2. Specify the SID of the clone. 3. Select the clone naming scheme. <p>If the database is cloned to the source host, the clone naming scheme will be auto-generated. If the database is cloned to an alternate host, clone naming scheme will be identical.</p> <ol style="list-style-type: none"> 4. Specify the Oracle home path. 5. (Optional) Specify the database credentials. <ul style="list-style-type: none"> ◦ Database credential: If the OS User authentication is disabled, you should provide a password for the sys user to connect to the cloned database on the same or target host. ◦ ASM credential: If the OS user authentication is disabled on the target host, you should provide a sysasm privileged user's credentials to connect to the ASM instance on the target host. <div style="display: flex; align-items: center; margin-top: 20px;"> <div style="text-align: center; margin-right: 10px;">  </div> <div> <p>Ensure that the listener is up and running on the target host.</p> </div> </div> <ol style="list-style-type: none"> 6. Click Next. 7. Click Clone.

If you have selected...	Do this...
<p>Specification file</p>	<ol style="list-style-type: none"> 1. Click Download File to download the specification file. 2. Select the clone naming scheme. If you select, Auto-generated, you should specify the suffix. 3. Edit the specification file as per the requirement and upload it by clicking the Browse button. 4. Select the clone host. If you want to create the clone on an alternate host, select the host having the same version of Oracle and OS as that of the source database host. 5. Specify the SID of the clone. 6. (Optional) Specify the database credentials. <ul style="list-style-type: none"> ◦ Database credential: If the OS User authentication is disabled, you should provide a password for the sys user to connect to the cloned database on the same or target host. ◦ ASM credential: If the OS user authentication is disabled on the target host, you should provide a sysasm privileged user's credentials to connect to the ASM instance on the target host. <div style="display: flex; align-items: center; margin-top: 10px;"> <div style="text-align: center; margin-right: 10px;">  </div> <div> <p>Ensure that the listener is up and running on the target host.</p> </div> </div> 7. Click Next. 8. Click Clone.

5. Click  adjacent to **Filter By** and select **Clone options > Clones** to view the clones.

Refresh SAP HANA target system

You can perform a refresh of a SAP HANA target system with the data of a SAP HANA source system. This can be used to provide the current production data into a test system. BlueXP backup and recovery allows you to select a Snapshot copy from a source system and creates a new Azure NetApp Files volume based on the Snapshot copy. Example scripts are available, which executes the required operations on the

database host to recover the SAP HANA database.

Before you begin

- You should install the SAP HANA target system before you execute the first refresh operation.
- You should add the source and target HANA systems manually into BlueXP backup and recovery.
- Ensure that the SAP HANA database version is same on source and the target system.
- You should have decided on which refresh scripts to be used. The refresh scripts are available in the solution technical report.

Automation example scripts

You can customize the refresh scripts.

- The following environmental variables are available as part of the prescript and postscript:
 - CLONED_VOLUMES_MOUNT_PATH
 - <SOURCEVOLUME>_DESTINATION
 - HANA_DATABASE_TYPE
 - TENANT_DATABASE_NAMES
- You must upgrade the plug-in to 3.0 version.
- The mount paths should be the same for the data volume on both the source and the target SAP HANA systems.
- Before the first refresh operation, ensure that the '/etc/fstab' file does not have entries for the data volumes of the target SAP HANA system.

About this task

- System refresh is supported only for multi-tenant database container HANA system.
- The existing policies will be valid after the system refresh.
- The new volumes created will have the following naming convention: <sourcevolumename>-<timestamp>
 - Timestamp format: <year><month><day>-<hour><minute><second>

For example, if the source volume is vol1, the refreshed volume name will be vol1-20230109-184501



The new volume will be placed in the same capacity pool as that of the target volumes.

- The junction path will be the same as the volume name.
- The “max throughput number” for the new volume is picked from the volume of the target system with manual Quality of Service (QoS) capacity pools.
For auto QoS capacity pools the throughput is defined by the capacity of the source volume.
- During system refresh, the auto mount and unmount of the volumes are performed using workflows instead of scripts.

Steps

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

2. In the **Applications** page, click **☰** icon to select the action corresponding to the system that you want to refresh and select **System Refresh**.
3. In the **System Refresh** page, perform the following actions:
 - a. Select source system and Snapshot copy.
 - b. (Optional) Enter Export addresses from which the new volumes can be accessed.
 - c. (Optional) Enter Maximum storage throughput (MIBs).
 - d. Enter prescript, postscript, and on failure script paths.
On failure script is executed only when the system refresh operation fails.
 - e. Click **Refresh**.

Manage protection of cloud-native application data

Monitor jobs

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



The scheduled jobs will be listed in the BlueXP Job monitor page after a delay of 5 minutes (maximum) from the job completion time.

For more information, refer to [Monitor job status](#).

Maintenance of Oracle database hosts

An admin can manually put the database hosts in maintenance mode to perform maintenance tasks on the hosts. During upgrade, the hosts are automatically put to maintenance mode and after upgrade, the hosts are automatically switched to production mode.

When the hosts are put in maintenance mode, the on-demand operations fail and the scheduled jobs are skipped.



You cannot verify if any jobs are running for the resources on the hosts before putting the hosts in maintenance mode.

Steps

1. In the BlueXP UI, click **Protection > Backup and recovery > Applications**
2. Select **Oracle** as the application type.
3. Click **Settings > Hosts**.
4. Perform one of the following actions:

If you...	Do this...
Want to put the host in maintenance mode	Click ☰ corresponding to the host and select Enable maintenance mode .

If you...	Do this...
Want to bring the host out of maintenance mode	Click  corresponding to the host that is under maintenance and select Disable maintenance mode .

Audit data

When you either run an API directly or use the UI to make the API call to any of the externally exposed APIs of the BlueXP backup and recovery for applications, the request details such as headers, role, request body, and API information will be logged in the BlueXP timeline and the audit entries are retained in the timeline forever. The status and error response of the API call are also audited post operation completion. In the case of asynchronous API responses like jobs, the job id also gets logged as part of response.

BlueXP backup and recovery for applications log the entries such as host IP, request body, operation name, who triggered, some headers, and the operation state of the API.

View backup details

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

Steps

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



The agent ID is associated to the Connector. If a Connector that was used during registering the SAP HANA host no longer exists, the subsequent backups of that application will fail because the agent ID of the new Connector is different. You should modify the Connector id in the host. For information, see [Update the Connector Details](#).

Delete clone

You can delete a clone if you no longer require.

Steps

1. Click  adjacent to **Filter By** and select **Clone options > Clone parents**.
2. Click  corresponding to the application and click **View Details**.
3. In the Database Details page, click  adjacent to **Filter By** and select **Clone**.
4. Click  corresponding to the clone that you want to delete and click **Delete**.
5. (Optional) Select the **force delete** checkbox.

Update the Connector Details

You should deploy a new Connector, if the Connector that was used during registering the application host no longer exists or is corrupted. After deploying the new connector, you should run the **connector-update** API to update the Connector details for all hosts registered using the old connector.

After updating the Connector details for Oracle or SAP HANA hosts, perform the following to ensure that the Connector details were updated successfully.

Steps

1. Login into BlueXP Connector VM and perform the following steps:

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

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

b. Obtain the base mount path.

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

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

2. Log in to the plug-in host and perform the following steps:

a. Navigate to /var/opt/snapcenter/spl/etc and run the keytool command to import the certificate.pem file.

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

b. Restart SPL: `systemctl restart spl`

c. Perform one of the following:

If you are on...	Do this...
Oracle database host	<ol style="list-style-type: none">1. Ensure that all the prerequisites are met.2. Click Backup and recovery > Applications3. Click ... corresponding to the application and click View Details.4. Modify Connector ID.

If you are on...	Do this...
SAP HANA database host	<ol style="list-style-type: none"> 1. Ensure that all the prerequisites are met. 2. Run the following command: <div data-bbox="867 275 1485 976" style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <pre>curl --location --request PATCH 'https://snapcenter.cloudmanager .cloud.netapp.com/api/saphana/ho sts/connector/update' \ --header 'x-account-id: <CM account-id>' \ --header 'Authorization: Bearer token' \ --header 'Content-Type: application/json' \ --data-raw '{ "old_connector_id": "Old connector id that no longer exists", "new_connector_id": "New connector Id"}'</pre> </div> <p>Connector details will get updated successfully if all the hosts have SnapCenter Plug-in for SAP HANA service installed and running and also if they are all reachable from the new Connector.</p>

Configure CA signed certificate

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

Configure CA signed certificate for BlueXP Connector

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

Steps

1. Perform the following steps on the Connector to use the CA certificate as the client certificate when the Connector is connecting with the plug-in.
 - a. Login to Connector.
 - b. Run the following command to get the `<base_mount_path>`:

```
sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs
sudo docker volume inspect | grep Mountpoint
```
 - c. Delete all the existing files located at `<base_mount_path>/client/certificate` in the Connector.

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

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

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

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

2. Perform the following steps on the plug-in host to validate the certificate sent by the Connector.

- a. Log in to the plug-in host.
- b. Copy the `certificate.pem` and certificates for all the intermediate CA and root CA from the Connector to the plug-in host at `/var/opt/snapcenter/spl/etc/`.



The format of the Intermediate CA and root CA certificate should be in `.crt` format.

- c. Navigate to `/var/opt/snapcenter/spl/etc` and run the `keytool` command to import the `certificate.pem` file.

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

- d. Import the root CA and intermediate certificates.

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



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

- e. Restart SPL: `systemctl restart spl`

Configure CA signed certificate for the plug-in

The CA certificate should have the same name as registered in Cloud Backup for the plug-in host.

Steps

1. Perform the following steps on the plug-in host to host the plug-in using the CA certificate.

- a. Navigate to the folder containing the SPL's keystore `/var/opt/snapcenter/spl/etc`.
- b. Create the PKCS12 format of the certificate having both certificate and key with alias `splkeystore`.

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

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

- c. Add the CA certificate created in the above step.

```
keytool -importkeystore -srckeystore certificate.p12 -srcstoretype pkcs12 -destkeystore keystore.jks -deststoretype JKS -srcalias splkeystore -destalias splkeystore -noprompt
```

- d. Verify the certificates.

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

- e. Restart SPL: `systemctl restart spl`
2. Perform the following steps on the Connector so that the Connector can verify the plug-in's certificate.
 - a. Log in to the Connector as non-root user.
 - b. Run the following command to get the `<base_mount_path>`:

```
sudo docker volume ls | grep scs_cloud_volume | awk {'print $2'} | xargs
sudo docker volume inspect | grep Mountpoint
```
 - c. Copy the the root CA and intermediate CA files under the server directory.

```
cd <base_mount_path>
mkdir server
```

The CA files should be in pem format.
 - d. Connect to the `cloudmanager_scs_cloud` and modify the **enableCACert** in `config.yml` to **true**.

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

```
sudo docker restart cloudmanager_scs_cloud
```

Access REST APIs

The REST APIs to protect the applications to cloud is available at:
<https://snapcenter.cloudmanager.cloud.netapp.com/api-doc/>.

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

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.