



Protect applications running on Azure NetApp Files

SnapCenter software

NetApp
February 20, 2026

Table of Contents

- Protect applications running on Azure NetApp Files 1
 - Protect applications running on Azure NetApp Files 1
 - Limitations 1
 - Install SnapCenter and create credentials 1
 - Install SnapCenter on Azure Virtual Machine 1
 - Create the Azure credential in SnapCenter 3
 - Configure the Azure storage account 3
 - Create the credential to add the plug-in host 4
- Protect SAP HANA databases 5
 - Add hosts and install SnapCenter plug-in for SAP HANA database 5
 - Add SAP HANA database 5
 - Create backup policies for SAP HANA databases 6
 - Create resource groups and attach SAP HANA backup policies 6
 - Back up SAP HANA databases running on Azure NetApp Files 7
 - Back up SAP HANA resource groups 8
 - Restore and recover SAP HANA databases 9
 - Clone SAP HANA database backup 10
- Protect Microsoft SQL Server databases 11
 - Add hosts and install SnapCenter plug-in for SQL Server database 11
 - Create backup policies for SQL Server databases 11
 - Create resource groups and attach SQL backup policies 13
 - Back up SQL Server databases running on Azure NetApp Files 14
 - Back up SQL Server resource groups 15
 - Restore and recover SQL Server databases 15
 - Clone SQL Server database backup 16
- Protect Oracle databases 18
 - Add hosts and install SnapCenter plug-in for Oracle database 18
 - Create backup policies for Oracle databases 18
 - Create resource groups and attach Oracle backup policies 19
 - Back up Oracle databases running on Azure NetApp Files 20
 - Back up Oracle resource groups 21
 - Restore and recover Oracle databases 21
 - Clone Oracle database backup 24

Protect applications running on Azure NetApp Files

Protect applications running on Azure NetApp Files

SnapCenter supports protecting your applications such as Oracle, SQL, and SAP HANA residing on Azure NetApp Files. Beginning with 6.0.1 release, SnapCenter supports the Azure NetApp Files backup feature that expands the data protection capabilities of Azure NetApp Files by providing fully managed backup solution for long-term recovery, archive, and compliance.

Azure NetApp Files is a premium storage solution that can be expensive for long-term backup retention. To optimize costs, you can move the backups from the Azure NetApp Files storage to an Azure object store. Beginning with SnapCenter 6.0.1, you can back up and clone applications residing on Azure NetApp Files to Azure Blob Storage (object store). You can maintain two copies of your data, volume snapshot copies on Azure NetApp Files storage for short term recovery and another copy on Azure Blob Storage for long term recovery.

When a policy with Azure NetApp Files backup is enabled and associated with an resource, SnapCenter handles the creation of volume snapshots and backing it to Azure Blob Storage. SnapCenter creates the Backup Vault and enable backup for the volume. If you have enabled backup for the volume, SnapCenter utilizes the existing vault.

Limitations

- Object storage functionalities for FAS, ASA, or AFF ONTAP and Amazon FSx for NetApp ONTAP storage systems is not supported.
- Oracle and SAP HANA mount and catalog workflows are not supported for object storage backups but are supported for snapshots.
- Oracle PDB clones are not supported for object storage backups but are supported for snapshots.
- Backup verification from object storage, REST API support, clone lifecycle management from object storage, and reporting features for object storage backups are not supported.
- Restoring from backups on Azure Blob Storage to Azure NetApp Files is not supported. You can use the clone option, alternatively.
- Clone split is not supported.

Install SnapCenter and create credentials

Install SnapCenter on Azure Virtual Machine

You can download the SnapCenter software from the NetApp Support site and install the software on the Azure virtual machine.

Before you begin

- Ensure that the Azure Windows virtual machine meets the requirements for SnapCenter Server installation. For information, see [Requirements to install SnapCenter Server](#).
- If you are new to Azure NetApp Files and do not have an existing NetApp account, ensure that you have

registered so that you can access the SnapCenter Software. For information, see [Register to access the SnapCenter software](#).

Steps

1. Download the SnapCenter Server installation package from [NetApp Support Site](#).
2. Initiate the SnapCenter Server installation by double-clicking the downloaded .exe file.

After you initiate the installation, all the pre-checks are performed and if the minimum requirements are not met appropriate error or warning messages are displayed. You can ignore the warning messages and proceed with installation; however, errors should be fixed.

3. Review the pre-populated values required for the SnapCenter Server installation and modify if required.

You do not have to specify the password for MySQL Server repository database. During SnapCenter Server installation the password is auto generated.



The special character “%” is not supported in the custom path for the repository database. If you include “%” in the path, installation fails.

4. Click **Install Now**.

If you have specified any values that are invalid, appropriate error messages will be displayed. You should re-enter the values, and then initiate the installation.



If you click the **Cancel** button, the step that is being executed will be completed, and then start the rollback operation. The SnapCenter Server will be completely removed from the host.

However, if you click **Cancel** when "SnapCenter Server site restart" or "Waiting for SnapCenter Server to start" operations are being performed, installation will proceed without cancelling the operation.

Register the product to enable support

If you are new to NetApp and do not have an existing NetApp account, you should register the product to enable support.

Steps

1. After installing SnapCenter, navigate to **Help > About**.
2. In the *About SnapCenter* dialog box, make a note of the SnapCenter Instance, a 20 digit number that starts with 971.
3. Click <https://register.netapp.com>.
4. Click **I am not a registered NetApp Customer**.
5. Specify your details to register yourself.
6. Leave the NetApp Reference SN field blank.
7. Select **SnapCenter** from the Product Line drop-down.
8. Select the billing provider.
9. Enter the 20-digit SnapCenter instance ID.
10. Click **Submit**.

Create the Azure credential in SnapCenter

You should create the Azure credential in SnapCenter to access the Azure NetApp account.

Before you begin

- Ensure that you have created the service principal in Azure.
- Ensure that you have the tenant ID, client ID, and secret key associated with the service principal available with you.
 - The Tenant ID can be found in the Azure Portal on the Overview page for Entra ID.
 - The Client ID is also known as the Application ID for the Enterprise App/Service Principal. This can be found in the Azure Portal on the Overview page for the Enterprise Application you've created to act as the Service Principal for SnapCenter.
 - The Client Secret Key is also known as the Secret Value. You can create this client secret in the Azure Portal by navigating to **App registrations** under Entra ID. After selecting the Enterprise App that you created, navigate to **Certificates & secrets** and then **New client secret**.



The secret value can only be accessed when created. You cannot access it later.

- The Service Principal needs to be granted permissions. The role of Contributor will allow the service principal to perform required actions in Azure. This can be granted from the Access Control (IAM) page within the Subscription page.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, specify the following information required to create the credential.

For this field...	Do this...
Credential Name	Enter a name for the credential.
Authentication Mode	Select Azure Credential from the drop-down list.
Tenant ID	Enter the tenant ID.
Client ID	Enter the client ID.
Client Secret Key	Enter the client secret key.

5. Click **OK**.

Configure the Azure storage account

You should configure the Azure storage account in SnapCenter.

The Azure storage account contains details about the subscription ID, Azure credential, and Azure NetApp account.



The standard licenses and capacity-based license is not required for Azure NetApp Files.

Steps

1. In the left navigation pane, click **Storage Systems**.
2. In the Storage Systems page, select **Azure NetApp Files** and click **New**.
3. Select the credential, subscription ID, and NetApp account from the respective drop-down lists.
4. Click **Submit**.


Create the credential to add the plug-in host

SnapCenter uses credentials to authenticate users for SnapCenter operations.

You should create credentials for installing SnapCenter plug-ins and additional credentials for performing data protection operations.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, specify the following information required to create the credential.

For this field...	Do this...
Credential Name	Enter a name for the credential.
Authentication Mode	Select the authentication mode from the drop-down list.
Authentication Type	Select either Password Based or SSH Key Based (only for Linux host).
Username	Specify the username.
Password	If you have selected Password based authentication, specify the password.
SSH Private Key	If you have selected SSH Key Based authentication, specify the private key.
Use sudo privileges	Select the Use sudo privileges check box if you are creating credentials for a non-root user.  This is applicable only for Linux users.

5. Click **OK**.

Protect SAP HANA databases

Add hosts and install SnapCenter plug-in for SAP HANA database

You must use the SnapCenter Add Host page to add hosts, and then install the plug-ins packages. The plug-ins are automatically installed on the remote hosts.

Before you begin

- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.
- If you are installing on the centralized host, ensure that the SAP HANA client software is installed on that host and open the required ports on the SAP HANA database host to run the HDB SQL queries remotely.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected.
3. Click **Add**.
4. In the Hosts page, perform the following actions:
 - a. In the Host Type field, select the host type.
 - b. In the Host name field, enter the fully qualified domain name (FQDN) or the IP address of the host.
 - c. In the Credentials field, enter the credential that you created.
5. In the Select Plug-ins to Install section, select the plug-ins to install.
6. (Optional) Click **More Options** and specify the details.
7. Click **Submit**.
8. If host type is Linux, verify the fingerprint, and then click **Confirm and Submit**.

In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.

9. Monitor the installation progress.

Add SAP HANA database

You should add the SAP HANA database manually.

About this task

Resources need to be added manually if the plug-in is installed on a centralized server. If the SAP HANA plug-in is installed on the HANA database host, then the HANA system is discovered automatically.



Automatic discovery is not supported for HANA multi-host configuration, they must be added through centralized plug-in only.

Steps

1. In the left navigation pane, select the SnapCenter Plug-in for SAP HANA Database from the drop-down list, and then click **Resources**.
2. In the Resources page, click **Add SAP HANA Database**.
3. In the Provide Resource Details page, perform the following actions:
 - a. Enter the resource type either as Single Container, Multitenant Database Container, or Non-data Volume.
 - b. Enter the SAP HANA system name.
 - c. Enter the system ID (SID).
 - d. Select the plug-in host.
 - e. Enter the key to connect to the SAP HANA system.
 - f. Enter the username for whom the HDB Secure User Store Key is configured.
4. In the Provide Storage Footprint page, select **Azure NetApp Files** as the storage type.
 - a. Select the Azure NetApp account.
 - b. Select the capacity pool and the associated volumes.
 - c. Click **Save**.
5. Review the summary, and then click **Finish**.

Create backup policies for SAP HANA databases

Before you use SnapCenter to back up SAP HANA database resources, you must create a backup policy for the resource or resource group that you want to back up.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.
4. In the Name page, enter the policy name and description.
5. In the Policy type page, perform the following steps:
 - a. Select **Azure NetApp Files** as the storage type.
 - b. Select **File-Based** if you want to perform an integrity check of the database.
 - c. Select **Snapshot Based** if you want to create a backup using Snapshot technology.
6. In the Snapshot and backup page, perform the following steps:
 - a. Select the frequency of the scheduled backup.
 - b. Specify the retention settings.
 - c. If you want to enable Azure NetApp Files backup, select **Enable backup** and specify the retention settings.
7. Review the summary and click **Finish**.

Create resource groups and attach SAP HANA backup policies

A resource group is the container to which you must add resources that you want to back


up and protect.

A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field...	Do this...
Name	Enter a name for the resource group.
Tags	Enter one or more labels that will help you later search for the resource group.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot name.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.
5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
6. In the Policies page, perform the following steps:
 - a. Select one or more policies from the drop-down list.
 - b. In the Configure Schedules column, click  for the policy you want to configure.
 - c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.
7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
8. Review the summary, and then click **Finish**.

Back up SAP HANA databases running on Azure NetApp Files

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, filter resources from the **View** drop-down list based on resource type.
3. Select the resource that you want to back up.
4. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.
5. In the Application Settings page, do the following:

- a. Select the **Backups** arrow to set additional backup options.
- b. Select the **Scripts** arrow to run pre and post commands for quiesce, Snapshot, and unquiesce operations.
- c. Select the **Custom Configurations** arrow, and then enter the custom value pairs required for all jobs using this resource.
- d. Select the **Snapshot Copy Tool > SnapCenter without File System Consistency** to create Snapshots.

The **File System Consistency** option is applicable only for applications running on Windows hosts.

6. In the Policies page, perform the following steps:
 - a. Select one or more policies from the drop-down list.
 - b. Select in the Configure Schedules column for the policy for which you want to configure a schedule.
 - c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.

policy_name is the name of the policy that you selected.
7. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.

You must also specify the sender and receiver email addresses, and the subject of the email. SMTP must also be configured in **Settings > Global Settings**.

8. Review the summary, and then select **Finish**.
9. Select **Back up Now**.
10. In the Backup page, perform the following steps:
 - a. If multiple policies are associated with the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.
11. Select **Backup**.
12. Monitor the operation progress by clicking **Monitor > Jobs**.

Back up SAP HANA resource groups

A resource group is a collection of resources on a host. A backup operation on the resource group is performed on all resources defined in the resource group.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:

- a. If multiple policies are associated with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.

If the policy selected for the on-demand backup is associated with a backup schedule, the on-demand backups will be retained based on the retention settings specified for the schedule type.

- b. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor > Jobs**.

Restore and recover SAP HANA databases


You can restore and recover data from the backups.

About this task

For Auto discovered HANA systems, if the **Complete Resource** option is selected, then restore is performed using Single File snapshot restore technology. If the **Fast Restore** check box is selected, then Volume Revert technology is used.

For manually added resources, Volume Revert technology is always used.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.
3. Select the resource or select a resource group and then select a resource in that group.
4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or vaulted) storage systems.
5. In the Primary backup(s) table, select the backup that you want to restore from, and then click .
6. In the Restore Scope page, select **Complete Resource**.

All the configured data volumes of the SAP HANA database are restored.

7. For Auto discovered HANA systems, in the Recovery scope page, perform the following actions:
 - a. Select **Recover to most recent state** if you want to recover as close as possible to the current time.
 - b. Select **Recover to point in time** if you want to recover to the specified point in time.
 - c. Select **Recover to specified data backup** if you want to recover to a specific data backup.
 - d. Select **No recovery** if you do not want to recover now.
 - e. Specify the log backup locations.
 - f. Specify the backup catalog location.
8. In the Pre ops page, enter pre restore and unmount commands to run before performing a restore job.
9. In the Post ops page, enter mount and post restore commands to run after performing a restore job.
10. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.


You must also specify the sender and receiver email addresses and the subject of the email. SMTP must also be configured on the **Settings > Global Settings** page.

11. Review the summary, and then click **Finish**.
12. Monitor the operation progress by clicking **Monitor > Jobs**.

Clone SAP HANA database backup

You can use SnapCenter to clone a SAP HANA database using the backup of the database. The clones created are thick clones and are created on the parent capacity pool.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, filter resources from the **View** drop-down list based on resource type.
3. Select the resource or resource group.
4. From the Manage Copies view, select **Backups** from the primary storage system.
5. Select the data backup from the table, and then click .
6. In the Location page, perform the following actions:
 - a. Select the host that has the SAP HANA plug-in installed for managing the cloned HANA system.

It can be a centralized plug-in host or HANA system host.



If the HANA plug-in is installed on a centralized host that manages HANA databases on other hosts, while creating or deleting clones, SnapCenter intentionally skips host side operations (mount or unmount filesystem) as the target server is a centralized host. You should use custom pre or post clone scripts to perform mount and unmount operations.

- b. Enter the SAP HANA SID to clone from the existing backups.
 - c. Enter IP addresses or the host names on which the cloned volumes will be exported.
 - d. If the SAP HANA database ANF volumes are configured in a manual QOS capacity pool, specify the QOS for the cloned volumes.

If QOS for the cloned volumes is not specified, the QOS of the source volume will be used. If the automatic QOS capacity pool is used, the QOS value specified will be ignored.
7. In the Scripts page, perform the following steps:
 - a. Enter the commands for pre clone or post clone that should be run before or after the clone operation, respectively.
 - b. Enter the mount command to mount a file system to a host.

If the source HANA system is auto discovered and the clone target host plug-in is installed on the SAP HANA host, then SnapCenter automatically unmounts the existing HANA data volumes on the clone target host and mounts the newly cloned HANA data volumes.
 8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
 9. Review the summary, and then click **Finish**.
 10. Monitor the operation progress by clicking **Monitor > Jobs**.



Clone Split is disabled for ANF clones because ANF clone is already an independent volume created from the selected Snapshot.

Protect Microsoft SQL Server databases

Add hosts and install SnapCenter plug-in for SQL Server database

SnapCenter supports data protection of SQL instances on SMB shares on Azure NetApp Files. The standalone and availability group (AG) configurations are supported.

You must use the SnapCenter Add Host page to add hosts, and then install the plug-ins package. The plug-ins are automatically installed on the remote hosts.

Before you begin

- You must be a user that is assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.
- When installing a plug-in on a Windows host, if you specify a credential that is not built-in, or if the user belongs to a local workgroup user, you must disable UAC on the host.

Steps

1. In the left navigation pane, select **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Select **Add**.
4. In the Hosts page do the following:
 - a. In the Host Type field, select the host type.
 - b. In the Host name field, enter the fully qualified domain name (FQDN) or the IP address of the host.
 - c. In the Credentials field, enter the credential that you created.
5. In the **Select Plug-ins to Install** section, select the plug-ins to install.
6. (Optional) Click **More Options** and specify the details.
7. Select **Submit**.
8. Select **Configure log directory** and in the Configure host log directory page, enter the SMB path of the host log directory, and click **Save**.
9. Click **Submit** and monitor the installation progress.

Create backup policies for SQL Server databases

You can create a backup policy for the resource or the resource group before you use SnapCenter to back up SQL Server resources, or you can create a backup policy at the time you create a resource group or backup a single resource.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Click **New**.

4. In the Name page, enter the policy name and description.
5. In the Policy type page, perform the following steps:
 - a. Select **Azure NetApp Files** as the storage type.
 - b. Select the backup type.
 - i. Select **Full Backup and Log Backup** if you want to back up database files and transaction logs.
 - ii. Select **Full Backup** if you want to back up only the database files.
 - iii. Select **Log Backup** if you want to back up only the transaction logs.
 - iv. Select **Copy Only Backup** if you want to back up your resources by using another application.
 - c. In the Availability Group Settings section, perform the following actions:
 - i. Select Backup on preferred backup replica if you want to back up only on the replica.
 - ii. Select primary AG replica or the secondary AG replica for the backup.
 - iii. Select the backup priority.
6. In the Snapshot and backup page, perform the following steps:
 - a. Select the frequency of the scheduled backup.
 - b. Specify the retention settings depending on the backup type selected.
 - c. If you want to enable Azure NetApp Files backup, select **Enable backup** and specify the retention settings.
7. In the Verification page, perform the following steps:
 - a. In the Run verification for following backup schedules section, select the schedule frequency.
 - b. In the Database consistency check options section, perform the following actions:
 - i. Select **Limit the integrity structure to physical structure of the database (PHYSICAL_ONLY)** to limit the integrity check to the physical structure of the database and to detect torn pages, checksum failures, and common hardware failures that impact the database.
 - ii. Select **Suppress all information messages (NO_INFOMSGS)** to suppress all informational messages.

Selected by default.
 - iii. Select **Display all reported error messages per object (ALL_ERRORMSGs)** to display all the reported errors per object.
 - iv. Select **Do not check nonclustered indexes (NOINDEX)** if you do not want to check nonclustered indexes.

The SQL Server database uses Microsoft SQL Server Database Consistency Checker (DBCC) to check the logical and physical integrity of the objects in the database.
 - v. Select **Limit the checks and obtain the locks instead of using an internal database Snapshot copy (TABLOCK)** to limit the checks and obtain locks instead of using an internal database Snapshot.
 - c. In the **Log Backup** section, select **Verify log backup upon completion** to verify the log backup upon completion.
 - d. In the **Verification script settings** section, enter the path and the arguments of the prescript or postscript that should be run before or after the verification operation, respectively.

8. Review the summary and click **Finish**.

Create resource groups and attach SQL backup policies

A resource group is the container to which you must add resources that you want to back up and protect.

A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field...	Do this...
Name	Enter a name for the resource group.
Tags	Enter one or more labels that will help you later search for the resource group.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot name.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.
5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
6. In the Policies page, perform the following steps:
 - a. Select one or more policies from the drop-down list.
 - b. In the Configure Schedules column, click for the policy you want to configure.
 - c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.
 - d. Select the Microsoft SQL Server scheduler.
7. In the Verification page, perform the following steps:
 - a. Select the verification server.
 - b. Select the policy for which you want to configure your verification schedule, and then click .
 - c. Either select **Run verification after backup** or **Run scheduled verification**.
 - d. Click **OK**.
8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
9. Review the summary, and then click **Finish**.

Back up SQL Server databases running on Azure NetApp Files

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Before you begin

You should create a load balancer, if the Azure Windows Failover Cluster does not have a cluster IP assigned or if it is not reachable from SnapCenter. The load balancer's IP should be configured and reachable from the SnapCenter Server.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, select **Database**, **Instance**, or **Availability Group** from the View drop-down list.
3. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.
4. In the Policies page, perform the following steps:
 - a. Select one or more policies from the drop-down list.
 - b. Select in the Configure Schedules column for the policy for which you want to configure a schedule.
 - c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.
policy_name is the name of the policy that you selected.
 - d. Select **Use Microsoft SQL Server scheduler**, and then select the scheduler instance from the **Scheduler Instance** drop-down list that is associated with the scheduling policy.
5. In the Verification page, perform the following steps:
 - a. Select the verification server.
 - b. Select the policy for which you want to configure your verification schedule, and then click .
 - c. Either select **Run verification after backup** or **Run scheduled verification**.
 - d. Click OK.
6. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
7. Review the summary, and then click **Finish**.
8. Select **Back up Now**.
9. In the Backup page, perform the following steps:
 - a. If multiple policies are associated with the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.
 - b. Select **Verify after backup**.
 - c. Select **Backup**.
10. Monitor the operation progress by clicking **Monitor > Jobs**.

Back up SQL Server resource groups

You can back up the resource groups that consist of multiple resources. A backup operation on the resource group is performed on all resources defined in the resource group.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:
 - a. If multiple policies are associated with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.
 - b. After backup, select **Verify** to verify the on-demand backup.
 - c. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor > Jobs**.


Restore and recover SQL Server databases

You can use SnapCenter to restore SQL Server databases. Restoring a database is a multiphase process that copies all the data and log pages from a specified SQL Server backup to a specified database.

Before you begin

- Configure the target instance with an active directory user who belongs to the SMB ADactive directory domain and has permissions to set the file permissions appropriately.
- Configure the credentials at SnapCenter instance level.
- If the SQL database exists on CIFS share, then the RunAS account and SQL service account (domain account) of the plug-in host need to be part of the BUILTIN\Administrators group of the CIFS server in ONTAP.
- SMB configurations do not support SQL authentication for the target instance. Configure the target instance in SnapCenter with an active directory user who has the required permissions.
- If the SnapCenter Plug-in service account is not an active directory user, ensure that a user with full control over the source volumes is available when restoring to an alternate host.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database or the resource group from the list.
4. From the Manage Copies view, select **Backups** from storage system.
5. Select the backup from the table, and then click the  icon.
6. In the Restore Scope page, select one of the following options:

- a. Select **Restore the database to the same host where the backup was created** to restore the database to the same SQL server.
 - b. Select **Restore the database to an alternate host** if you want to restore the database to a different SQL server in the same or another host where backups are taken.
7. In the Recovery Scope page, select one of the following options:
- a. Select **None** when you need to restore only the full backup without any logs.
 - b. Select **All log backups** up-to-the-minute backup restore operation to restore all the available log backups after the full backup.
 - c. Select **By log backups** to perform a point-in-time restore operation, which restores the database based on backup logs until the backup log with the selected date.
 - d. Select **By specific date until** to specify the date and time after which transaction logs are not applied to the restored database.
 - e. If you have selected **All log backups**, **By log backups**, or **By specific date until** and the logs are located at a custom location, select **Use custom log directory**, and then specify the log location.
8. In the Pre-Ops and Post Ops page, specify the required details.
9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
10. Review the summary, and then click **Finish**.
11. Monitor the restore process by using the **Monitor > Jobs** page.

Clone SQL Server database backup

You can use SnapCenter to clone a SQL database using the backup of the database. The clones created are thick clones and are created on the parent capacity pool.


About this task

You should ensure that the target instance for clone is configured with an active directory user who belongs to the SMB ADActive directory domain and has permissions to set the file permissions appropriately. You should configure the credentials in SnapCenter at instance level.

The SQL authentication for target instance will not be supported for SMB configurations. The target instance should be configured in SnapCenter with the active directory user having the required permissions.

If the SnapCenter Plug-in services service account is not an active directory user then while performing clone, the user who have the full control over the source volumes is required so that it can be impersonated and perform the required operation.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database or resource group.
4. From the **Manage Copies** view page, select the backup from primary storage system.
5. Select the backup, and then select .
6. In the **Clone Options** page, provide all the required details.
7. In the Location page, select a storage location to create a clone.

If the SQL Server database ANF volumes are configured in a manual QOS capacity pool, specify the QOS for the cloned volumes.


If QOS for the cloned volumes is not specified, the QOS of the source volume will be used. If the automatic QOS capacity pool is used, the QOS value specified will be ignored.

8. In the Logs page, select one of the following options:
 - a. Select **None** if you want to clone only the full back up without any logs.
 - b. Select **All log backups** if you want to clone all the available log backups dated after the full backup.
 - c. Select **By log backups until** if you want to clone the database based on the backup logs that were created up to the backup log with the selected date.
 - d. Select **By specific date until** if you do not want to apply the transaction logs after the specified date and time.
9. In the **Script** page, enter the script timeout, path, and the arguments of the prescript or postscript that should be run before or after the clone operation, respectively.
10. In the **Notification** page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
11. Review the summary, and then select **Finish**.
12. Monitor the operation progress by selecting **Monitor > Jobs**.

Perform Clone Lifecycle

Using SnapCenter, you can create clones from a resource group or database. You can either perform on-demand clone or you can schedule recurring clone operations of a resource group or database. If you clone a backup periodically, you can use the clone to develop applications, populate data, or recover data.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database or resource group.
4. From the **Manage Copies** view page, select the backup from primary storage system.
5. Select the backup, and then select .
6. In the **Clone Options** page, provide all the required details.
7. In the Location page, select a storage location to create a clone.

If the SQL Server database ANF volumes are configured in a manual QOS capacity pool, specify the QOS for the cloned volumes.

If QOS for the cloned volumes is not specified, the QOS of the source volume will be used. If the automatic QOS capacity pool is used, the QOS value specified will be ignored.

8. In the **Script** page, enter the script timeout, path, and the arguments of the prescript or postscript that should be run before or after the clone operation, respectively.
9. In the Schedule page, perform one of the following actions:
 - Select **Run now** if you want to execute the clone job immediately.
 - Select **Configure schedule** when you want to determine how frequently the clone operation should

occur, when the clone schedule should start, on which day the clone operation should occur, when the schedule should expire, and whether the clones must be deleted after the schedule expires.

10. In the **Notification** page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
11. Review the summary, and then select **Finish**.
12. Monitor the operation progress by selecting **Monitor > Jobs**.

Protect Oracle databases

Add hosts and install SnapCenter plug-in for Oracle database

You can use the Add Host page to add hosts, and then install the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX. The plug-ins are automatically installed on the remote hosts.

You can add a host and install plug-in packages either for an individual host or for a cluster. If you are installing the plug-in on a cluster (Oracle RAC), the plug-in is installed on all the nodes of the cluster. For Oracle RAC One Node, you should install the plug-in on both active and passive nodes.

Steps

1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected.
3. Click **Add**.
4. In the Hosts page, perform the following actions:
 - a. In the Host Type field, select the host type.
 - b. In the Host name field, enter the fully qualified domain name (FQDN) or the IP address of the host.
 - c. In the Credentials field, enter the credential that you created.
5. In the Select Plug-ins to Install section, select the plug-ins to install.
6. (Optional) Click **More Options** and specify the details.
7. Click **Submit**.
8. Verify the fingerprint, and then click **Confirm and Submit**.

In a cluster setup, you should verify the fingerprint of each of the nodes in the cluster.

9. Monitor the installation progress.

Create backup policies for Oracle databases

Before you use SnapCenter to back up Oracle database resources, you must create a backup policy for the resource or the resource group that you want to back up.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Policies**.
3. Select Oracle Database from the drop-down list.

4. Click **New**.
5. In the Name page, enter the policy name and description.
6. In the Policy type page, perform the following steps:
 - a. Select **Azure NetApp Files** as the storage type.
 - b. Select the backup type as either online or offline backup.
 - c. If you want to catalog backup using Oracle Recovery Manager (RMAN), select **Catalog backup with Oracle Recovery Manager (RMAN)**.
 - d. If you want to prune archive logs after backup, select **Prune archive logs after backup**.
 - e. Specify the delete archive log settings.
7. In the Snapshot and backup page, perform the following steps:
 - a. Select the frequency of the scheduled backup.
 - b. Specify the retention settings.
 - c. If you want to enable Azure NetApp Files backup, select **Enable backup** and specify the retention settings.
8. In the Script page, enter the path and the arguments of the prescript or postscript that you want to run before or after the backup operation, respectively.
9. In the Verification page, select the backup schedule for which you want to perform the verification operation and enter the path and the arguments of the prescript or postscript that you want to run before or after the verification operation, respectively.
10. Review the summary and click **Finish**.

Create resource groups and attach Oracle backup policies

A resource group is the container to which you must add resources that you want to back up and protect.



A resource group enables you to back up all the data that is associated with a given application simultaneously. A resource group is required for any data protection job. You must also attach one or more policies to the resource group to define the type of data protection job that you want to perform.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, click **New Resource Group**.
3. In the Name page, perform the following actions:

For this field...	Do this...
Name	Enter a name for the resource group.
Tags	Enter one or more labels that will help you later search for the resource group.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot name.



For this field...	Do this...
Archive log file destination	Specify the destinations of the archive log files.

4. In the Resources page, select a host name from the **Host** drop-down list and resource type from the **Resource Type** drop-down list.
5. Select the resources from the **Available Resources** section, and then click the right arrow to move them to the **Selected Resources** section.
6. In the Policies page, perform the following steps:
 - a. Select one or more policies from the drop-down list.
 - b. In the Configure Schedules column, click  for the policy you want to configure.
 - c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then click **OK**.
7. In the Verification page, perform the following steps:
 - a. Select the verification server.
 - b. Select the policy for which you want to configure your verification schedule, and then click *  .
 - c. Either select **Run verification after backup** or **Run scheduled verification**.
 - d. Click **OK**.
8. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
9. Review the summary, and then click **Finish**.

Back up Oracle databases running on Azure NetApp Files

If a resource is not yet part of any resource group, you can back up the resource from the Resources page.

Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resource page, select **Database** from the View drop-down list.
3. In the Resource page, select **Use custom name format for Snapshot copy**, and then enter a custom name format that you want to use for the Snapshot name.
4. In the Policies page, perform the following steps:
 - a. Select one or more policies from the drop-down list.
 - b. Select  in the Configure Schedules column for the policy for which you want to configure a schedule.
 - c. In the Add schedules for policy *policy_name* dialog box, configure the schedule, and then select **OK**.
5. In the Verification page, perform the following steps:
 - a. Select the verification server.
 - b. Select the policy for which you want to configure your verification schedule, and then click  .

- c. Either select **Run verification after backup** or **Run scheduled verification**.
 - d. Click OK.
6. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
7. Review the summary, and then click **Finish**.
8. Select **Back up Now**.
9. In the Backup page, perform the following steps:
 - a. If multiple policies are associated with the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.
 - b. Click **Backup**.
10. Monitor the operation progress by clicking **Monitor > Jobs**.

Back up Oracle resource groups

You can back up the resource groups that consist of multiple resources. A backup operation on the resource group is performed on all resources defined in the resource group.


Steps

1. In the left navigation pane, select **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.
3. In the Resource Groups page, select the resource group that you want to back up, and then select **Back up Now**.
4. In the Backup page, perform the following steps:
 - a. If multiple policies are associated with the resource group, from the **Policy** drop-down list, select the policy that you want to use for backup.
 - b. Select **Backup**.
5. Monitor the operation progress by selecting **Monitor > Jobs**.

Restore and recover Oracle databases

In the event of data loss, you can use SnapCenter to restore data from one or more backups to your active file system and then recover the database.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database or the resource group from the list.
4. From the Manage Copies view, select **Backups** from the primary storage system.
5. Select the backup from the table, and then click  .
6. In the Restore Scope page, perform the following tasks:
 - a. Select RAC if you have selected a backup of a database in RAC environment.

- b. Perform the following actions:
 - i. Select **All Datafiles** if you want to restore only the database files.
 - ii. Select **Tablespaces** if you want to restore only the tablespaces.
 - iii. Select **Redo log files** if you want to restore the redo log files of the Data Guard standby or Active Data Guard standby databases.
 - iv. Select **Pluggable databases** and specify the PDBs you want to restore.
 - v. Select **Pluggable database (PDB) tablespaces**, and then specify the PDB and the tablespaces of that PDB that you want to restore.
 - vi. Select **Restore the database to the same host where the backup was created** if you want to restore the database to the same SQL server where the backups are taken.
 - vii. Select **Restore the database to an alternate host** if you want the database to be restored to a different SQL server in the same or different host where backups are taken.
 - viii. Select **Change database state if needed for restore and recovery** to change the state of the database to the state required to perform restore and recovery operations.
 - ix. Select **Force in place restore** if you want to perform in-place restore in the scenarios where new datafiles are added after backup or when LUNs are added, deleted, or re-created to an LVM disk group.
7. In the Recovery Scope page, select one of the following options:
 - a. Select **All Logs** if you want to recover to the last transaction.
 - b. Select **Until SCN (System Change Number)** if you want to recover to a specific SCN.
 - c. Select **Date and Time** if you want to recover to a specific date and time.
 - d. Select **No recovery** if you do not want to recover.
 - e. Select **Specify external archive log locations** if you want to specify the location of the external archive log files.
8. In the Pre-Ops and Post Ops page, specify the required details.
9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
10. Review the summary, and then click **Finish**.
11. Monitor the operation progress by clicking **Monitor > Jobs**.


Restore and recover tablespaces using point-in-time recovery

You can restore a subset of tablespaces that have been corrupted or dropped without impacting the other tablespaces in the database. SnapCenter uses RMAN to perform point-in-time recovery (PITR) of the tablespaces.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database of type single instance (multitenant).
4. From the Manage Copies view, select **Backups** from the storage system.

If the backup is not cataloged, you should select the backup and click **Catalog**.

5. Select the catalogued backup, and then click .
6. In the Restore Scope page, perform the following tasks:
 - a. Select **RAC** if you have selected a backup of a database in RAC environment.
 - b. Select **Tablespaces** if you want to restore only the tablespaces.
 - c. Select **Change database state if needed for restore and recovery** to change the state of the database to the state required to perform restore and recovery operations.
7. In the Recovery Scope page, select one of the following options:
 - a. Select **Until SCN (System Change Number)** if you want to recover to a specific SCN.
 - b. Select **Date and Time** if you want to recover to a specific date and time.
8. In the Pre-Ops and Post Ops page, specify the required details.
9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
10. Review the summary, and then click **Finish**.
11. Monitor the restore process by using the **Monitor > Jobs** page.


Restore and recover pluggable database using point-in-time recovery

You can restore and recover a pluggable database (PDB) that has been corrupted or dropped without impacting the other PDBs in the container database (CDB). SnapCenter uses RMAN to perform point-in-time recovery (PITR) of the PDB.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database of type single instance (multitenant).
4. From the Manage Copies view, select **Backups** from the storage system.

If the backup is not cataloged, you should select the backup and click **Catalog**.

5. Select the catalogued backup, and then click .
6. In the Restore Scope page, perform the following tasks:
 - a. Select **RAC** if you have selected a backup of a database in RAC environment.
 - b. Depending on whether you want to restore the PDB or tablespaces in a PDB, perform one of the actions:
 - Select **Pluggable databases (PDBs)** if you want to restore a PDB.
 - Select **Pluggable database (PDB) tablespaces** if you want to restore tablespaces in a PDB.
7. In the Recovery Scope page, select one of the following options:
 - a. Select **Until SCN (System Change Number)** if you want to recover to a specific SCN.
 - b. Select **Date and Time** if you want to recover to a specific date and time.
8. In the Pre-Ops and Post Ops page, specify the required details.
9. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want


to send the emails.

10. Review the summary, and then click **Finish**.
11. Monitor the restore process by using the **Monitor > Jobs** page.

Clone Oracle database backup

You can use SnapCenter to clone an Oracle database using the backup of the database. The clones created are thick clones and are created on the parent capacity pool.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database.
4. From the Manage Copies view page, select the backup from primary storage system.
5. Select the Data backup, and then click .
6. In the Name page, select whether you want to clone a database (CDB or non CDB) or clone a pluggable database (PDB).
7. In the Locations page, specify the required details.

If the Oracle database ANF volumes are configured in a manual QOS capacity pool, specify the QOS for the cloned volumes.

If QOS for the cloned volumes is not specified, the QOS of the source volume will be used. If the automatic QOS capacity pool is used, the QOS value specified will be ignored.

8. In the Credentials page, perform one of the following:
 - a. For Credential name for sys user, select the Credential to be used for defining the sys user password of the clone database.
 - b. For ASM Instance Credential name, select **None** if OS authentication is enabled for connecting to the ASM instance on the clone host.

Otherwise, select the Oracle ASM credential configured with either “sys” user or a user having “sysasm” privilege applicable to the clone host.

9. In the Pre-Ops page specify the path and arguments of the prescripts and in the Database Parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.
10. In the Post-Ops page, **Recover database** and **Until Cancel** are selected by default to perform recovery of the cloned database.
 - a. If you select **Until Cancel**, SnapCenter performs recovery by mounting the latest log backup having the unbroken sequence of archive logs after that data backup that was selected for cloning.
 - b. If you select **Date and time**, SnapCenter recovers the database up to a specified date and time.
 - c. If you select **Until SCN**, SnapCenter recovers the database up to a specified SCN.
 - d. If you select **Specify external archive log locations**, SnapCenter identifies and mounts optimal number of log backups based on the specified SCN or the selected date and time.
 - e. By default, **Create new DBID** check box is selected to generate a unique number (DBID) for the cloned

database differentiating it from the source database.


Clear the check box if you want to assign the DBID of the source database to the cloned database. In this scenario, if you want to register the cloned database with the external RMAN catalog where the source database is already registered, the operation fails.

- f. Select **Create tempfile for temporary tablespace** check box if you want to create a tempfile for the default temporary tablespace of the cloned database.
 - g. In **Enter sql entries to apply when clone is created**, add the sql entries that you want to apply when the clone is created.
 - h. In **Enter scripts to run after clone operation**, specify the path and the arguments of the postscript that you want to run after the clone operation.
11. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
 12. Review the summary, and then select **Finish**.
 13. Monitor the operation progress by selecting **Monitor > Jobs**.

Clone a pluggable database

You can clone a pluggable database (PDB) to a different or same target CDB on the same host or alternate host. You can also recover the cloned PDB to a desired SCN or date and time.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the View list.
3. Select the database of type single instance (multitenant).
4. From the Manage Copies view page, select the backup from primary storage system.
5. Select the backup, and then click .
6. In the Name page, select **PDB Clone** and specify the other details.
7. In the Locations page, specify the required details.
8. In the Pre-Ops page specify the path and arguments of the prescripts and in the Database Parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.
9. In the Post-Ops page, **Until Cancel** is selected by default to perform recovery of the cloned database.
 - a. If you select **Until Cancel**, SnapCenter performs recovery by mounting the latest log backup having the unbroken sequence of archive logs after that data backup that was selected for cloning.
 - b. If you select **Date and time**, SnapCenter recovers the database up to a specified date and time.
 - c. If you select **Specify external archive log locations**, SnapCenter identifies and mounts optimal number of log backups based on the specified SCN or the selected date and time.
 - d. By default, **Create new DBID** check box is selected to generate a unique number (DBID) for the cloned database differentiating it from the source database.

Clear the check box if you want to assign the DBID of the source database to the cloned database. In this scenario, if you want to register the cloned database with the external RMAN catalog where the source database is already registered, the operation fails.

- e. Select **Create tempfile for temporary tablespace** check box if you want to create a tempfile for the default temporary tablespace of the cloned database.
 - f. In **Enter sql entries to apply when clone is created**, add the sql entries that you want to apply when the clone is created.
 - g. In **Enter scripts to run after clone operation**, specify the path and the arguments of the postscript that you want to run after the clone operation.
10. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the emails.
 11. Review the summary, and then select **Finish**.
 12. Monitor the operation progress by selecting **Monitor > Jobs**.

Copyright information

Copyright © 2026 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.